"The Missing Manual series is simply the most intelligent and usable series of guidebooks..."
—KEVIN KELLY, CO-FOUNDER OF *WIRED*

Excel 2010

the missing manual®

The book that should have been in the box®

O'REILLY®

Matthew MacDonald

OFFICE APPLICATIONS/SPREADSHEETS/MS EXCEL

Answers found here!

Excel 2010: The Missing Manual. Excel, the world's most popular spreadsheet program, has the muscle to analyze heaps of data. Beyond basic number-crunching, Excel 2010 has many impressive features that are hard to find, much less master—especially from online help pages. This Missing Manual clearly explains how everything works with a unique and witty style to help you learn quickly.

the missing manual®

The book that should have been in the box®

The important stuff you need to know

- **Navigate with ease.** Master Excel's tabbed toolbar and its new backstage view
- **Perform a variety of calculations.** Write formulas for rounding numbers, calculating mortgage payments, and more
- **Organize your data.** Search, sort, and filter huge amounts of information
- **Illustrate trends.** Bring your data to life with charts and graphics—including miniature charts called Sparklines
- **Examine your data.** Summarize information and find hidden patterns with pivot tables and slicers
- **Share your spreadsheets.** Use the Excel Web App to collaborate with colleagues online
- **Rescue lost data.** Restore old versions of data and find spreadsheets you forgot to save

Matthew MacDonald is a science and technology writer with well over a dozen books to his name, including *Access 2010: The Missing Manual*, *Creating a Web Site: The Missing Manual*, *Your Brain: The Missing Manual*, and *Your Body: The Missing Manual*.

US $39.99 CAN $49.99

ISBN: 978-1-449-38235-3

Free online edition for 45 days with purchase of this book. Details on last page.

O'REILLY®

missingmanuals.com

Excel 2010

THE MISSING MANUAL

The book that
should have been
in the box®

Excel 2010

THE MISSING MANUAL

Matthew MacDonald

POGUE PRESS™
O'REILLY®
Beijing • Cambridge • Farnham • Köln • Sebastopol • Taipei • Tokyo

Excel 2010: The Missing Manual

Matthew MacDonald

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North,
Sebastopol, CA 95472.

O'Reilly Media books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles: *http://my.safaribooksonline.com.* For more information, contact our corporate/institutional sales department: 800-998-9938 or *corporate@oreilly.com*.

June 2010: First Edition.

This book uses a durable and flexible lay-flat binding.

ISBN: 9781449382353

[CS]

Table of Contents

The Missing Credits . . . . **xix**

Introduction **1**

Part One: Worksheet Basics

Chapter 1: Creating Your First Spreadsheet **19**

Creating a Basic Worksheet 19
- Starting a New Workbook 21
- Adding the Column Titles 22
- Adding Data 24
- Editing Data 25

Moving Around the Grid 26
- Shortcut Keys 26
- The Go To Feature 28

A Tour of the Excel Window 29
- The Tabs of the Ribbon 29
- The Formula Bar 31
- The Status Bar 32
- Going Backstage 36
- Excel Options 37

Saving Files 37
- The Excel File Format 39
- Sharing Your Spreadsheet with Excel 2007 41
- Saving Your Spreadsheet for Excel 2003 43
- Saving a Spreadsheet in Other Formats 44
- Saving Your Spreadsheet As a PDF 44
- Saving Your Spreadsheet As an HTML File 47
- Saving Your Spreadsheet with a Password 50
- Disaster Recovery 53

Opening Files . . . 55
Opening Recent Documents . . . 56
Protected View . . . 57
Opening Files—with a Twist . . . 58
Working with Multiple Open Spreadsheets . . . 60

Chapter 2: Adding Information to Worksheets . . . 63
Adding Different Types of Data . . . 63
How Excel Identifies Text . . . 66
How Excel Identifies Numbers . . . 67
How Excel Identifies Dates and Times . . . 68
Regional Dating . . . 70
Handy Timesavers . . . 71
AutoComplete . . . 72
AutoCorrect . . . 73
AutoFill . . . 75
AutoFit . . . 78
Dealing with Change: Undo, Redo, and AutoRecover . . . 79
Undo and Redo . . . 79
AutoRecover . . . 82
Spell Check . . . 85
Spell Checking Options . . . 87
Adding Hyperlinks . . . 90
Creating a Link to a Web Page or Document . . . 91
Creating a Link to a Worksheet Location . . . 93

Chapter 3: Moving Data . . . 95
Selecting Cells . . . 95
Moving Cells Around . . . 102
Adding and Moving Columns or Rows . . . 113

Chapter 4: Managing Worksheets . . . 117
Worksheets and Workbooks . . . 118
Adding and Removing Worksheets . . . 119
Hiding Worksheets . . . 121
Naming and Rearranging Worksheets . . . 121
Moving Worksheets from One Workbook to Another . . . 123
Grouping Worksheets . . . 125
Find and Replace . . . 127
The Basic Find . . . 128
Find All . . . 129
More Advanced Searches . . . 130
Finding Formatted Cells . . . 133
Finding and Replacing Values . . . 134

Chapter 5: Formatting Cells 137
Formatting Cell Values 138
Changing the Cell Value Format 139
Formatting Numbers 141
Formatting Dates and Times 146
Special Formats for Special Numbers 149
Custom Formats 150
Formatting Cell Appearance 158
Alignment and Orientation 159
Fonts and Color 162
Borders and Fills 167
Drawing Borders by Hand 169

Chapter 6: Smart Formatting Tricks 173
The Format Painter 173
Styles and Themes 174
Custom Styles 176
Modifying Styles 178
Transferring Styles Between Workbooks 179
Themes: A Package of Styles 180
Modifying Themes 183
Conditional Formatting 186
The Basics of Conditional Formatting 187
Highlighting Specific Values 187
Using Multiple Rules 192

Chapter 7: Viewing and Printing Worksheets 195
Controlling Your View 195
Zooming 196
Viewing Distant Parts of a Spreadsheet at Once 198
Freezing Columns or Rows 201
Hiding Data 204
Saving View Settings 205
Viewing Multiple Workbooks at Once 207
Printing 210
How to Print an Excel File 211
Page Layout View: A Better Print Preview 217
Creating Headers and Footers 220
Controlling Pagination 222
Page Breaks 222
Scaling 223
Page Break Preview: A Bird's-Eye View of Your Worksheet 224

Part Two: Formulas and Functions

Chapter 8: Building Basic Formulas. 227

Creating a Basic Formula 227
Excel's Order of Operations 229
Cell References 231
How Excel Formats Cells That Contain Cell References 232
Functions 234
Using a Function in a Formula 234
Using Cell References with a Function 236
Using Cell Ranges with a Function 237
Excel Compatibility Functions 239
Formula Errors 240
Logical Operators 243
Formula Shortcuts. 245
Point-and-Click Formula Creation 245
Point-and-Click Formula Editing 246
The Formulas Tab. 247
Using the Insert Function Button. 249
Copying Formulas. 252
Absolute Cell References 254
Partially Fixed References 256
Referring to Other Worksheets. 258
Referring to Other Workbooks 258

Chapter 9: Math and Statistical Functions 263

Rounding Numbers 263
ROUND(), ROUNDDOWN(), ROUNDUP(): Rounding Numbers 264
MROUND(), CEILING(), and FLOOR(): More Rounding Functions 265
INT() and TRUNC(): Chopping Off Non-Whole Numbers. 266
EVEN() and ODD(): Rounding Up to Even or Odd Values. 268
Groups of Numbers. 268
SUM(): Summing Up Numbers 268
COUNT(), COUNTA(), and COUNTBLANK(): Counting Items in a List 269
MAX() and MIN(): Finding Maximum and Minimum Values 271
LARGE(), SMALL(), and RANK(): Ranking Numbers. 272
AVERAGE() and MEDIAN(): Finding Average or Median Values 274
PERCENTILE() and PERCENTRANK(): Advanced Ranking Functions 275
FREQUENCY(): Putting Numbers into Grouped Ranges. 277
MODE(): Finding Numbers That Frequently Occur in a List. 279
General Math Functions. 279
PRODUCT(), FACT(), POWER(), and SQRT(): Products, Factorials,
Powers, and Square Roots 280
QUOTIENT() and MOD(): Higher Division. 281
ABS() and SIGN(): Absolute Value and Determining a Number's Sign 281
RAND() and RANDBETWEEN(): Generating Random Numbers 282

GCD() and LCM(): Greatest and Least Common Denominator 284
COMBIN() and PERMUT(): Figuring Combinations and Permutations 285
Trigonometry and Advanced Math 286
Advanced Statistics 289

Chapter 10: Financial Functions 295
The World of Finance 295
Financial Functions 296
FV(): Future Value 297
PV(): Present Value 300
PMT(), PPMT(), and IPMT(): Calculating the Number of Payments You Need to Make 302
NPER(): Figuring Out How Much Time You'll Need to Pay Off a Loan or Meet an Investment Target 304
RATE(): Figuring the Interest Rate You Need to Achieve Future Value 305
NPV() and IRR(): Net Present Value and Internal Rate of Return 306
Depreciation. 309
Other Financial Functions. 312

Chapter 11: Manipulating Dates, Times, and Text 315
Manipulating Text 315
CONCATENATE(): Joining Strings of Text Together 316
LEFT(), MID(), and RIGHT(): Copying Portions of a Text String 316
LEN(), FIND(), and SEARCH(): Counting Characters in a String. 317
UPPER(), LOWER(), and PROPER(): Changing Capitalization 319
TRIM() and CLEAN(): Removing Unwanted Spaces and Non-Printing Characters . . 320
SUBSTITUTE(): Replacing One Sequence of Characters with Another 320
TEXT(), VALUE(), FIXED(), and DOLLAR(): Converting Text to Numbers and Vice Versa 321
Other Text Functions 323
Manipulating Dates and Times 324
Math with Dates and Times. 325
Using Dates and Times with Ordinary Functions 325
Date and Time Functions 326
TODAY() and NOW(): Inserting the Current Date and Time. 327
DATE() and TIME(): Dates and Times in Calculations 327
DAY(), MONTH(), and YEAR(): More Date Calculations 329
HOUR(), MINUTE(), SECOND(): More Time Calculations. 330
WEEKDAY(): Determining the Day of the Week 330
DATEDIF(): Calculating the Difference Between Dates 332
DATEVALUE() and TIMEVALUE(): Converting Dates and Times into Serial Numbers 333
DAYS360(): Finding Out the Number of Days Between Two Dates 335
EDATE(): Calculating Future Dates 335
YEARFRAC(): Calculating the Percentage of a Year Between Two Dates 336
EOMONTH(): Finding the Last Day of Any Month. 337

NETWORKDAYS(): Counting the Number of Business Days . . . 337
WORKDAY(): Figuring Out When Days Will Fall in the Future . . . 338
WEEKNUM(): Figuring Out in Which Week a Date Falls . . . 338

Chapter 12: Lookup, Reference, and Information Functions . . . 339

The Basic Lookup . . . 340
VLOOKUP(): Vertical Lookups . . . 341
HLOOKUP(): Horizontal Lookups . . . 344
Advanced Lookups . . . 344
MATCH(): Finding the Position of Items in a Range . . . 345
INDEX(): Retrieving the Value from a Cell . . . 346
Performing a "Left Lookup" . . . 346
Performing a Double Lookup . . . 347
OFFSET(): Shifting Cell References . . . 349
Other Reference and Lookup Functions . . . 350
INDIRECT() and ADDRESS(): Working with Cell References Stored As Text . . . 351
TRANSPOSE(): Changing Rows into Columns and Vice Versa . . . 353
The HYPERLINK() Function: Creating a Dynamic Link . . . 355
Information Functions . . . 356
The "IS" Functions: Checking the Value Inside a Cell . . . 356
TYPE() and ERROR.TYPE(): Finding a Value's Data Type or Error Type . . . 357
INFO() and CELL(): Gathering Info About Your Computer and Your Worksheet's Cells . . . 358
Tutorial: Generating Invoices from a Product Catalog . . . 359

Chapter 13: Advanced Formula Writing and Troubleshooting . . . 363

Conditions in Formulas . . . 363
IF(): Building Conditional Formulas . . . 364
COUNTIF(): Counting Only the Cells You Specify . . . 365
SUMIF(): Adding Only the Cells You Specify . . . 367
COUNTIFS() and SUMIFS(): Counting and Summing Using Multiple Criteria . . . 368
Descriptive Names for Cell References . . . 371
Creating and Using a Named Range . . . 372
Creating Slightly Smarter Named Ranges . . . 373
Naming Formulas and Constants . . . 375
Managing Named Ranges . . . 376
Automatically Creating Named Ranges . . . 378
Applying Names to Existing Formulas . . . 380
Variable Data Tables . . . 382
Creating a One-Variable Data Table . . . 382
Creating a Two-Variable Data Table . . . 384
Controlling Recalculation . . . 386
Solving Formula Errors . . . 387
Step-by-Step Evaluation . . . 387
Tracing Precedents and Dependents . . . 389
Error Checking . . . 392

Part Three: Organizing Worksheets

Chapter 14: Tables: List Management Made Easy 395

The Basics of Tables . . . 396
Creating a Table . . . 396
Formatting a Table . . . 399
Editing a Table . . . 402
Selecting Parts of a Table . . . 404
Sorting and Filtering a Table . . . 405
Applying a Simple Sort Order . . . 406
Sorting with Multiple Criteria . . . 407
Sorting by Color . . . 410
Filtering with the List of Values . . . 411
Creating Smarter Filters . . . 413
Dealing with Duplicate Rows . . . 415
Highlighting Duplicates . . . 416
Removing Duplicates Automatically . . . 417
Performing Table Calculations . . . 418
Dynamic Calculations . . . 419
Column Names . . . 421
Table Names . . . 421
The Total Row . . . 422
The SUBTOTAL() Function . . . 423
The Database Functions . . . 425

Chapter 15: Grouping and Outlining Data 429

Basic Data Grouping . . . 429
Creating a Group . . . 430
Nesting Groups Within Groups . . . 434
Summarizing Your Data . . . 436
Combining Data from Multiple Tables . . . 438
Grouping Timesavers . . . 440
Auto Outline . . . 441
Automatic Subtotaling . . . 441

Chapter 16: Templates 445

Understanding Templates . . . 446
Creating a New Workbook from a Template . . . 447
Downloading Templates (Method 1: From Backstage View) . . . 448
Downloading Templates (Method 2: The Office Online Website) . . . 450
Creating Templates . . . 454
Understanding Custom Templates . . . 454
Building a Custom Template . . . 456
Sharing Templates with Others . . . 458

Part Four: Charts and Graphics

Chapter 17: Creating Basic Charts 461

Charting 101 462
- Embedded and Standalone Charts 463
- Creating a Chart with the Ribbon 463
- The Chart Tools Ribbon Tabs 465

Basic Tasks with Charts 466
- Moving and Resizing a Chart 466
- Creating a Standalone Chart 468
- Editing and Adding to Chart Data 470
- Changing the Chart Type 470
- Printing Charts 471

Practical Charting 472
- Charts with Multiple Series of Numbers 473
- Controlling the Data Excel Plots on the X-Axis 473
- Data That Uses a Date or Time Scale 476
- Noncontiguous Chart Ranges 478
- Changing the Order of Your Data Series 480
- Changing the Way Excel Plots Blank Values 481

Chart Types 481
- Column 482
- Bar 484
- Line 484
- Pie 486
- Area 486
- XY (Scatter) 488
- Stock 489
- Surface 490
- Donut 491
- Bubble 492
- Radar 493

Chapter 18: Formatting and Perfecting Charts 495

Chart Styles and Layouts 495
- Chart Styles 496
- Chart Layouts 497

Adding Chart Elements 497
- Adding Titles 499
- Adding a Legend 500
- Adding Data Labels to a Series 501
- Adding Individual Data Labels 504
- Adding a Data Table 506

Selecting Chart Elements 507

Formatting Chart Elements . . . 509
Coloring the Background . . . 509
Fancy Fills . . . 510
Fancy Borders and Lines . . . 514
Formatting Data Series and Data Points . . . 516
Reusing Your Favorite Charts with Templates . . . 518
Improving Your Charts . . . 520
Controlling a Chart's Scale . . . 520
Adding a Trendline . . . 523
Adding Error Bars to Scientific Data . . . 526
Formatting 3-D Charts . . . 528
Changing the Shape of a 3-D Column . . . 529
Advanced Charting . . . 530
Exploding Slices in a Pie . . . 530
Grouping Slices in a Pie . . . 531
Gaps, Widths, and Overlays in a Column Chart . . . 532
Creating Combination Charts . . . 534

Chapter 19: Inserting Graphics. . . . 537
Adding Pictures to a Worksheet . . . 538
Inserting a Picture . . . 538
Positioning and Resizing a Picture . . . 540
Picture Touch-Up . . . 543
Compressing Pictures . . . 545
Cropping and Shaping a Picture . . . 547
Picture Borders, Effects, and Styles . . . 550
Excel's Clip Art Library . . . 551
Drawing Shapes . . . 554
Drawing a Shape . . . 556
Adding Text to a Shape . . . 559
Selecting and Arranging Shapes . . . 560
Connecting Shapes . . . 564
SmartArt . . . 565

Part Five: Advanced Data Analysis

Chapter 20: Visualizing Your Data . . . 569
Data Bars . . . 570
Editing a Formatting Rule . . . 572
Fine-Tuning Data Bars . . . 573
Color Scales . . . 578
Fine-Tuning Color Scales . . . 578
Icon Sets . . . 579
Fine-Tuning Icon Sets . . . 581

Sparklines . . . 583
Creating a Sparkline . . . 585
Changing the Axis . . . 588
Markers . . . 590

Chapter 21: Scenarios and Goal Seeking . . . 593
Using Scenarios . . . 594
Creating a New Scenario . . . 594
Managing Scenarios . . . 597
Creating a Summary Report . . . 598
Using Goal Seek . . . 599
Goal Seeking with Complex Equations . . . 603
Solver . . . 604
Understanding Solver . . . 605
Defining a Problem in Solver . . . 606
More Advanced Solver Problems . . . 612
Saving Solver Models . . . 615
Configuring Solver . . . 617

Chapter 22: Pivot Tables . . . 621
Summary Tables Revisited . . . 622
Life Without Pivot Tables . . . 622
Life with Pivot Tables . . . 624
Building Pivot Tables . . . 624
Preparing a Pivot Table . . . 626
Pivot Table Regions . . . 628
Laying Out a Pivot Table . . . 629
Formatting a Pivot Table . . . 632
Rearranging a Pivot Table . . . 633
Getting to the Source . . . 633
Multi-Layered Pivot Tables . . . 635
Hiding and Showing Details . . . 637
Fine-Tuning Pivot Table Calculations . . . 639
Changing the Type of Calculation . . . 639
Adding a Calculated Field . . . 640
Filtering a Pivot Table . . . 643
Report Filtering . . . 643
Slicers . . . 645
Group Filtering . . . 650
Pivot Charts . . . 652
Creating a Pivot Chart . . . 653
Manipulating a Pivot Chart . . . 654

Chapter 23: Analyzing Databases, XML, and Web Pages 657

Excel and Databases . . . 658
Connecting to an Access Database . . . 659
Refreshing Data . . . 661
Data Source Security . . . 664
Connecting to a SQL Server Database . . . 665
Reusing Your Database Connection . . . 668
Understanding XML . . . 671
What Is XML, Really? . . . 671
Three Rules of XML . . . 672
XML Files and Schemas . . . 675
Excel and XML . . . 676
Mapping a Simple Document . . . 677
Importing and Exporting XML . . . 681
Mapping Lists . . . 683
Gaining the Benefits of XML Mapping . . . 685
Web Queries . . . 686
The Limitations of Web Queries . . . 687
Creating a Web Query . . . 687
The Research Pane—A Web Query Alternative . . . 690

Part Six: Sharing Data with the Rest of the World

Chapter 24: Protecting Your Workbooks 693

Understanding Excel's Safeguards . . . 693
Data Validation . . . 694
Settings . . . 695
Input Message . . . 696
Error Alert . . . 698
Data Validation with Formulas and Cell References . . . 700
Data Validation with Lists . . . 702
Locked and Hidden Cells . . . 703
Protecting a Worksheet . . . 705
Protecting the Entire Workbook . . . 707
Protecting Cell Ranges (with More Passwords) . . . 709
Allowing Specific Windows Users to Edit a Range . . . 711

Chapter 25: Worksheet Collaboration 715

Your Excel Identity . . . 716
Preparing Your Workbook . . . 716
Workbook Protection . . . 717
Checking for Issues . . . 718
Document Properties . . . 719

Distributing a Document . . . 722
Sending by Email . . . 723
Uploading to the Web . . . 724
Adding Comments . . . 724
Inserting a Comment . . . 725
Showing and Hiding Comments . . . 726
Fine-Tuning Comments . . . 727
Reviewing Comments . . . 728
Printing Comments . . . 729
Tracking Changes . . . 729
Switching On Change Tracking . . . 731
Understanding the Change Log . . . 732
Highlighting Changes . . . 734
Examining the Change Log . . . 735
Accepting and Rejecting Changes . . . 737
Merging Multiple Revisions into One Workbook . . . 739
Sharing Your Workbook . . . 741
Multiple Users Without Workbook Sharing . . . 742
Turning On Workbook Sharing . . . 743
Workbook Sharing in Action . . . 745

Chapter 26: Using Excel on the Web . . . 749
Putting Your Files Online . . . 750
Introducing SkyDrive . . . 751
Your Web Workflow . . . 753
Uploading a File to SkyDrive (Using Your Browser) . . . 753
Uploading a Workbook to SkyDrive (Using Excel) . . . 756
Uploading a Workbook to a SharePoint Server . . . 759
Viewing a File in Your SkyDrive Account . . . 760
Viewing a File in Someone Else's SkyDrive Account . . . 762
Using the Excel Web App . . . 764
Supported Features . . . 765
Saving Files . . . 768
Collaboration: The Excel Web App's Specialty . . . 768
Taking a Workbook Back to Desktop Excel . . . 770
Unsupported Features . . . 771
Partially Supported Features . . . 773

Chapter 27: Exchanging Data with Other Programs . . . 775
Sharing Information in Windows . . . 775
Embedding and Linking Objects . . . 777
Exporting Charts Out of Excel . . . 778
Editing a Linked Object . . . 781
Editing an Embedded Object . . . 782
Importing Objects into Excel . . . 784

Transferring Data . . . 785
Exporting Tables of Data . . . 786
Importing Tables of Data . . . 787
Importing Text Files . . . 787

Part Seven: Programming Excel

Chapter 28: Automating Tasks with Macros . . . 793
Macros 101 . . . 793
Macro-Free and Macro-Enabled Workbooks . . . 795
The Macro Recorder . . . 795
Relative and Absolute Recording . . . 796
Where Macros Live . . . 797
Recording a Macro . . . 799
Playing a Macro . . . 802
Macro Security . . . 803
Trusted Documents . . . 804
Temporary Trust . . . 806
The Trust Center . . . 806
Setting Up a Trusted Location . . . 809
Creating Practical Macros . . . 810
Inserting a Header . . . 811
Alternating Row Formatting . . . 812
A Combined Task . . . 813
Placing a Macro on the Quick Access Toolbar . . . 814
Attaching a Macro to a Button Inside a Worksheet . . . 816

Chapter 29: Programming Spreadsheets with VBA . . . 819
The Visual Basic Editor . . . 820
The Project Window . . . 820
Modules and Macros . . . 822
Finding and Moving Macros . . . 824
Debugging a Macro . . . 824
Understanding Macro Code . . . 826
The Anatomy of a Macro . . . 826
Objects 101 . . . 828
Using Properties and Methods . . . 829
Hunting for Objects . . . 830
Exploring the VBA Language . . . 830
Entering Text in the Current Cell . . . 831
Interacting with Other Cells . . . 832
Editing Specific Cells . . . 833

Formatting Cells 834
Using Variables 834
Making Decisions 836
Repeating Actions with a Loop 838
Creating Custom Functions 839

Part Eight: Appendix

Appendix A: Customizing the Ribbon 843

Index 855

The Missing Credits

About the Author

Matthew MacDonald (author) is a science and technology writer with well over a dozen books to his name. Office geeks can follow him into the word of databases with *Access 2010: The Missing Manual*. Web fans can build an online home with him in *Creating a Web Site: The Missing Manual.* And human beings of all description can discover just how strange they really are in *Your Brain: The Missing Manual* and *Your Body: The Missing Manual*.

About the Creative Team

Brian Sawyer (editor) is an editor for O'Reilly Media's Head First division. He's also served as lead editor for the company's popular Hacks series, editor for Missing Manuals and Make: Books, and contributing editor to *Craft* magazine. When not writing or editing about technology, he uses it to help train for marathons (see Chapter 4 of *Best Android Apps*).

Nellie McKesson (production editor) is a graduate of St. John's College in Santa Fe, New Mexico. She lives in Brockton, Mass., and spends her spare time studying graphic design and making t-shirts (*www.endplasticdesigns.com*). Email: *nellie@oreilly.com*.

Alison O'Byrne (copyeditor) is a full-time freelance editor with over eight years' experience specializing in corporate and government projects for international clients. She lives with her family in Dublin, Ireland. Email: *alison@alhaus.com*. Web: *www.alhaus.com*.

Zack Barresse (technical reviewer) is a Microsoft Excel MVP and has held that title for six consecutive years. He works as a full-time fire fighter in his hometown of Boardman, Oregon, where he resides with his wife and four children.

Rhea Howard (technical reviewer) is an Operations Analyst at O'Reilly Media. She lives with her partner in Berkeley, CA.

Acknowledgements

Writing a book about a program as sprawling and complex as Excel is a labor of love (love of pain, that is). I'm deeply indebted to a whole host of people who helped out with this edition and the two previous ones. They include Nellie McKesson, Brian Sawyer, Peter Meyers, Sarah Milstein, and technical reviewers Zack Barresse and Rhea Howard. I also owe thanks to many people who worked to get this book formatted, indexed, and printed—you can meet many of them on the Missing Credits page.

Completing this book required a few sleepless nights (and many sleep-deprived days). I extend my love and thanks to my daughters Maya and Brenna, who put up with it without crying most of the time, my dear wife Faria, who mostly did the same, and our moms and dads (Nora, Razia, Paul, and Hamid), who contributed hours of babysitting, tasty meals, and general help around the house that kept this book on track. So thanks everyone—without you half of the book would still be trapped inside my brain!

—Matthew MacDonald

The Missing Manual Series

Missing Manuals are witty, superbly written guides to computer products that don't come with printed manuals (which is just about all of them). Each book features a handcrafted index; cross-references to specific page numbers (not just "see Chapter 14"); and RepKover, a detached-spine binding that lets the book lie perfectly flat without the assistance of weights or cinder blocks.

Recent and upcoming titles include:

Access 2007: The Missing Manual by Matthew MacDonald

AppleScript: The Missing Manual by Adam Goldstein

AppleWorks 6: The Missing Manual by Jim Elferdink and David Reynolds

CSS: The Missing Manual by David Sawyer McFarland

Creating Web Sites: The Missing Manual by Matthew MacDonald

David Pogue's Digital Photography: The Missing Manual by David Pogue

Dreamweaver 8: The Missing Manual by David Sawyer McFarland

Dreamweaver CS3: The Missing Manual by David Sawyer McFarland

Dreamweaver CS4: The Missing Manual by David Sawyer McFarland

eBay: The Missing Manual by Nancy Conner

Excel 2003: The Missing Manual by Matthew MacDonald

Excel 2007: The Missing Manual by Matthew MacDonald

Facebook: The Missing Manual by E.A. Vander Veer

FileMaker Pro 8: The Missing Manual by Geoff Coffey and Susan Prosser

FileMaker Pro 9: The Missing Manual by Geoff Coffey and Susan Prosser

Flash 8: The Missing Manual by E.A. Vander Veer

Flash CS3: The Missing Manual by E.A. Vander Veer and Chris Grover

Flash CS4: The Missing Manual by Chris Grover with E.A. Vander Veer

FrontPage 2003: The Missing Manual by Jessica Mantaro

Google Apps: The Missing Manual by Nancy Conner

Google SketchUp: The Missing Manual by Chris Grover

The Internet: The Missing Manual by David Pogue and J.D. Biersdorfer

iMovie 6 & iDVD: The Missing Manual by David Pogue

iMovie '08 & iDVD: The Missing Manual by David Pogue

iPhone: The Missing Manual by David Pogue

iPhoto '08: The Missing Manual by David Pogue

iPhoto '09: The Missing Manual by David Pogue and J.D. Biersdorfer

iPod: The Missing Manual, 6th Edition by J.D. Biersdorfer

iWork '09: The Missing Manual by Josh Clark

JavaScript: The Missing Manual by David Sawyer McFarland

Living Green: The Missing Manual by Nancy Conner

Mac OS X: The Missing Manual, Tiger Edition by David Pogue

Mac OS X: The Missing Manual, Leopard Edition by David Pogue

Microsoft Project 2007: The Missing Manual by Bonnie Biafore

Netbooks: The Missing Manual by J.D. Biersdorfer

Office 2004 for Macintosh: The Missing Manual by Mark H. Walker and Franklin Tessler

Office 2007: The Missing Manual by Chris Grover, Matthew MacDonald, and E.A. Vander Veer

Office 2008 for Macintosh: The Missing Manual by Jim Elferdink

PCs: The Missing Manual by Andy Rathbone

Photoshop CS4: The Missing Manual by Lesa Snider

Photoshop Elements 7: The Missing Manual by Barbara Brundage

Photoshop Elements 6 for Mac: The Missing Manual by Barbara Brundage

PowerPoint 2007: The Missing Manual by E.A. Vander Veer

QuickBase: The Missing Manual by Nancy Conner

QuickBooks 2008: The Missing Manual by Bonnie Biafore

Quicken 2008: The Missing Manual by Bonnie Biafore

Quicken 2009: The Missing Manual by Bonnie Biafore

QuickBooks 2009: The Missing Manual by Bonnie Biafore

Switching to the Mac: The Missing Manual, Tiger Edition by David Pogue and Adam Goldstein

Switching to the Mac: The Missing Manual, Leopard Edition by David Pogue

Wikipedia: The Missing Manual by John Broughton

Windows XP Home Edition: The Missing Manual, Second Edition by David Pogue

Windows XP Pro: The Missing Manual, Second Edition by David Pogue, Craig Zacker, and Linda Zacker

Windows Vista: The Missing Manual by David Pogue

Windows Vista for Starters: The Missing Manual by David Pogue

Word 2007: The Missing Manual by Chris Grover

Your Body: The Missing Manual by Matthew MacDonald

Your Brain: The Missing Manual by Matthew MacDonald

Introduction

Most people don't need much convincing to use Excel, Microsoft's premier spreadsheet software. Its overwhelming popularity, especially in the business world, makes it the obvious choice for millions of number crunchers. But despite its wide use, few people know where to find Excel's most impressive features or why they'd want to use them in the first place. *Excel 2010: The Missing Manual* fills that void, explaining everything from basic Excel concepts to the fancy tricks of the trade.

This book not only teaches you how Excel works, but also shows you how to use Excel's tools to answer real-world questions like "How many workdays are there between today and my vacation?", "How much money do I need in the bank right now to retire a millionaire?", and "Statistically speaking, who's smarter—Democrats or Republicans?" Best of all, you'll steer clear of obscure options that aren't worth the trouble to learn, while homing in on the hidden gems that will win you the undying adoration of your coworkers, your family, and your friends—or at least your accountant.

Note: This book is written with Microsoft's latest and greatest release in mind: Excel 2010. This book isn't the best choice if you're using an earlier version of Excel, because Microsoft is continually changing Excel's user interface (the "look and feel" of the program). To get the right instructions, look for a previous edition of this book, such as *Excel 2007: The Missing Manual* or *Excel 2003: The Missing Manual.*

FREQUENTLY ASKED QUESTION

Is That 2,010 or 2010?

What do Excel's version numbers mean?

Most people realize that the "2010" in Excel 2010 indicates the year, not the 2,010th release of the software. Microsoft's on-again, off-again naming policy is to leave the actual version number out of product names. So what version is Excel 2010?

If you dig around a little, you'll discover that Excel 2010 is actually Excel Version 14. But even this version number doesn't mean what you might expect. Excel 14 is actually the *eleventh* release of Excel for Windows. The first version of Excel was a Macintosh-only release. There is no Excel 6, because Microsoft felt the changes they made were so great they were entitled to jump up two version numbers at once. (As questionable as that sounds, it's a technique nearly all software makers use at some point.) And Excel 13 is also missing in action, due to superstitions about the bad-luck number 13.

What You Can Do with Excel

Excel and Word are the two powerhouses of the Microsoft Office family. While Word lets you create and edit documents, Excel specializes in letting you create, edit, and analyze *data* that's organized into lists or tables. This grid-like arrangement of information is called a *spreadsheet*. Figure I-1 shows an example.

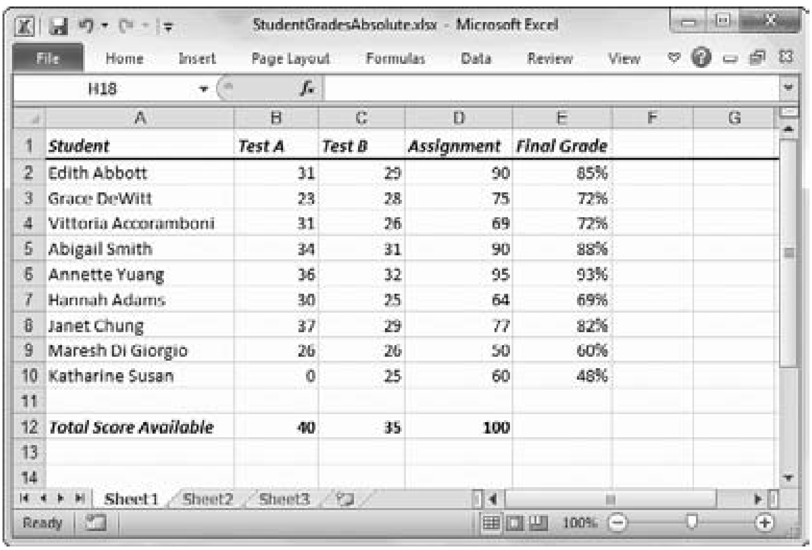

Figure I-1: *This spreadsheet lists nine students, each of whom has two test scores and an assignment grade. Using Excel formulas, it's easy to calculate the final grade for each student. And with a little more effort, you can calculate averages and medians, and determine each student's percentile. Chapter 8 looks at how to perform these calculations.*

Note: Excel shines when it comes to *numerical* data, but the program doesn't limit you to calculations. While it has the computing muscle to analyze stacks of numbers, it's equally useful for keeping track of the DVDs in your personal movie collection.

Some common spreadsheets include:

- **Business documents** like financial statements, invoices, expense reports, and earnings statements.
- **Personal documents** like weekly budgets, catalogs of your *Star Wars* action figures, exercise logs, and shopping lists.
- **Scientific data** like experimental observations, models, and medical charts.

These examples just scratch the surface. Resourceful spreadsheet gurus use Excel to build everything from cross-country trip itineraries to logs of every Kevin Bacon movie they've ever seen.

Of course, Excel really shines in its ability to help you *analyze* a spreadsheet's data. For example, once you've entered a list of household expenses, you can start crunching numbers with Excel's slick formula tools. Before long you'll have totals, subtotals, monthly averages, a complete breakdown of cost by category, and maybe even some predictions for the future. Excel can help track your investments and tell you how long until you'll have saved enough to buy that weekend house in Vegas.

The bottom line is that once you enter raw information, Excel's built-in smarts can help compute all kinds of useful figures. Figure I-2 shows a sophisticated spreadsheet that has been configured to help identify hot-selling product categories.

Note: Keen eyes will notice that neither of these examples (Figures I-1 and I-2) include the omnipresent Excel ribbon, which usually sits atop the window, stacked with buttons. That's because it's been collapsed neatly out of the way to let you focus on the spreadsheet. You'll learn how to use this trick yourself on page 30.

Excel is not just a math wizard. If you want to add a little life to your data, you can inject color, apply exotic fonts, and even create *macros* (automated sequences of steps) to help speed up repetitive formatting or editing chores. And if you're bleary-eyed from staring at rows and rows of spreadsheet numbers, you can use Excel's many chart-making tools to build everything from 3-D pie charts to more exotic scatter graphs. (See Chapter 17 to learn about all of Excel's chart types.) Excel can be as simple or as sophisticated as you want it to be.

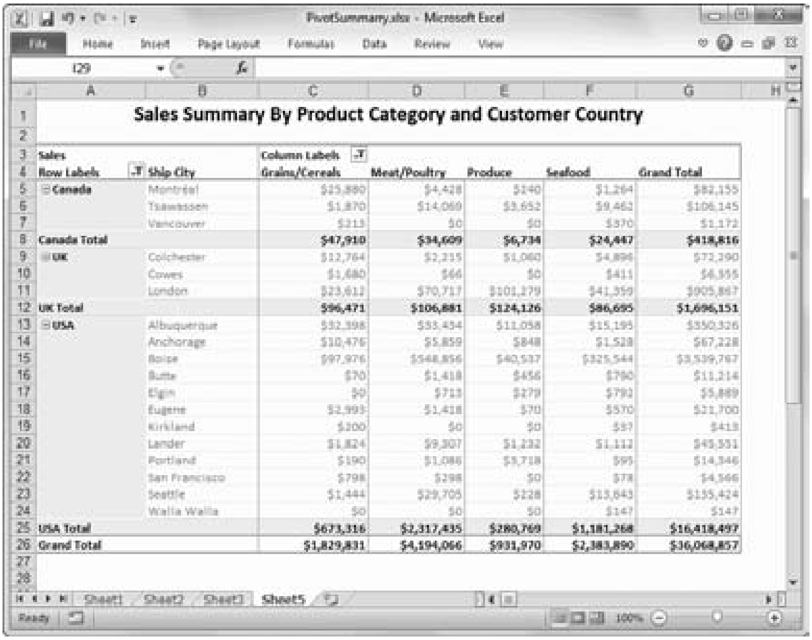

***Figure I-2:** This spreadsheet summarizes a company's total sales. The information is grouped based on where the company's customers live, and it's further divided according to product category. Summaries like these can help you spot profitable product categories and identify items popular in specific cities. This advanced example uses pivot tables, which are described in Chapter 22.*

The Modern Face of Excel

Although Microsoft is reluctant to admit it, most of Excel's core features were completed nearly 10 years ago. So what has Microsoft been doing ever since? The answer, at least in part, is spending millions of dollars on usability tests, which are aimed at figuring out how easy—or difficult—a program is to use. In a typical usability test, Microsoft gathers a group of spreadsheet novices, watches them fumble around with the latest version of Excel, and then tweaks the program to make it more intuitive.

After producing Excel 2003, Microsoft finally decided that minor tune-ups couldn't fix Excel's overly complex, button-heavy toolbars. So they decided to undertake a radical redesign to create a user interface that actually makes sense. The centerpiece of this redesign is the super-toolbar called the *ribbon*.

The Ribbon

Everything you'll ever want to do in Excel—from picking a fancy background color to pulling information out of a database—is packed into the ribbon. To accommodate all these buttons without becoming an over-stuffed turkey, the ribbon uses *tabs*. Excel starts out with seven tabs in the ribbon. When you click one of these tabs, you see a whole new collection of buttons (Figure I-3).

Note: Wondering what each tab holds? You'll take a tab tour in Chapter 1 on page 29.

The ribbon is the best thing to hit the Excel scene in years. The ribbon makes it easier to find features and remember where they are, because each feature is grouped into a logically related tab. Even better, once you find the button you need, you can often find other, associated commands by looking at the section where the button is placed. In other words, the ribbon isn't just a convenient tool—it's also a great way to explore Excel.

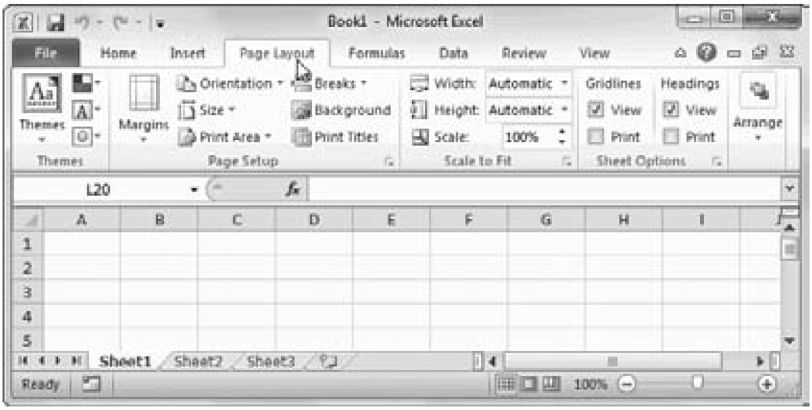

Figure I-3: *When you launch Excel, you start at the Home tab. But here's what happens when you click the Page Layout tab. Now, you have a slew of options for tasks like adjusting paper size and making a decent printout. The buttons in a tab are grouped into smaller boxes for clearer organization.*

The ribbon is full of craftsman-like detail. For example, when you hover over a button, you don't see a paltry two- or three-word description in a yellow box. Instead, you see a friendly pop-up box with a complete mini-description and a shortcut that lets you trigger this command from the keyboard. Another nice detail is the way you can jump through the tabs at high velocity by positioning the mouse pointer over the ribbon and rolling the scroll wheel (if your mouse has a scroll wheel). And you're sure to notice the way the ribbon rearranges itself when you change the size of the Excel window (see Figure I-4).

Using the Ribbon with the Keyboard

If you're an unredeemed keyboard lover, you'll be happy to hear that you can trigger ribbon commands with the keyboard. The trick is using *keyboard accelerators*, a series of keystrokes that starts with the Alt key (the same key you used to use to get to a menu). When using a keyboard accelerator, you *don't* hold down all the keys at the same time. (As you'll soon see, some of these keystrokes contain so many letters that you'd be playing Finger Twister if you tried holding them all down simultaneously.) Instead, you hit the keys one after the other.

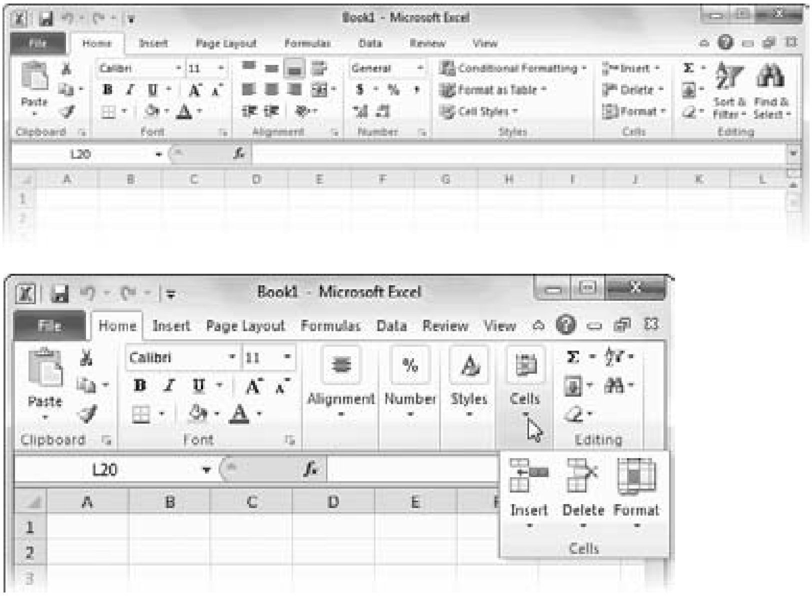

Figure I-4:
Top: A large Excel window gives you plenty of room to play. The ribbon uses the space effectively, making the most important buttons bigger.

Bottom: When you shrink the Excel window, the ribbon rearranges its buttons and makes some smaller (by shrinking the button's icon or leaving out the title). Shrink small enough, and you might run out of space for a section altogether. In that case, you get a single button (like the Number, Styles, and Cells sections in this example) for an entire section. Click this button and the missing commands appear in a drop-down panel.

The trick to using keyboard accelerators is to understand that once you hit the Alt key, there are two things you do, in this order:

1. **Pick the ribbon tab you want.**
2. **Choose a command in that tab.**

Before you can trigger a specific command, you *must* select the correct tab (even if it's already displayed). Every accelerator requires at least two key presses after you hit the Alt key. You need even more if you need to dig through a submenu.

By now, this whole process probably seems hopelessly impractical. Are you really expected to memorize dozens of different accelerator key combinations?

Fortunately, Excel is ready to help you out with a feature called KeyTips. Here's how it works. Once you press the Alt key, letters magically appear over every tab in the ribbon. Once you hit a key to pick a tab, letters appear over every button in that tab (Figure I-5). You can then press the corresponding key to trigger the command (Figure I-6).

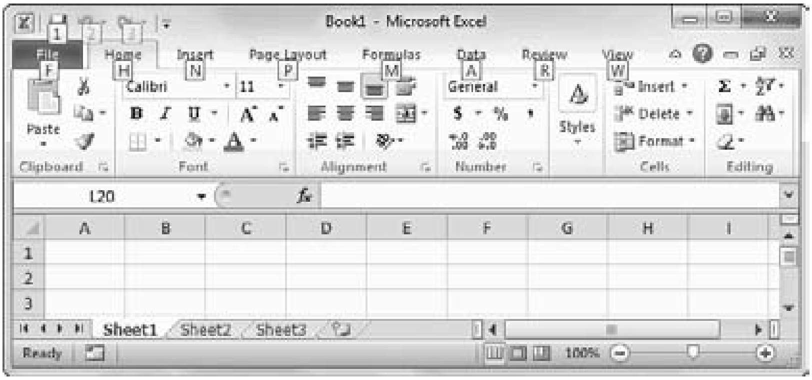

Figure I-5: *When you press Alt, Excel helps you out with KeyTips next to every tab, over the File menu, and over the buttons in the Quick Access toolbar. If you follow up with M (for the Formulas tab), you'll see letters next to every command in that tab, as shown in Figure I-6.*

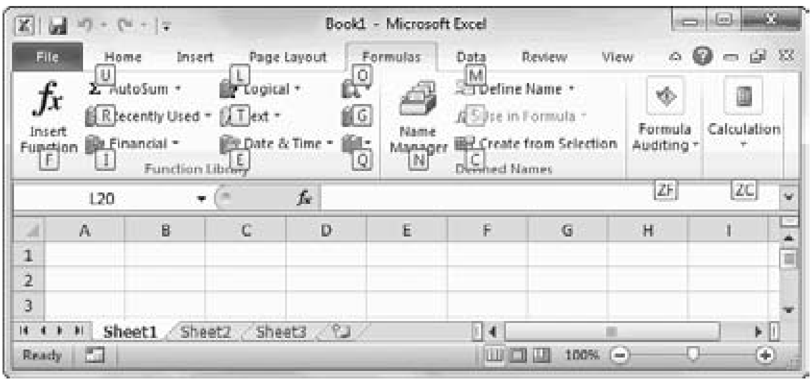

Figure I-6: *You can now follow up with F to trigger the Insert Function button, U to get to the AutoSum feature, and so on. Don't bother trying to match letters with tab or button names—there are so many features packed into the ribbon that in many cases the letters don't mean anything at all.*

In some cases, a command might have two letters, in which case you need to press both keys, one after the other. (For example, the Find & Select button on the Home tab has the letters FD. To trigger it, press Alt, then H, then F, and then D.)

Tip: You can back out of KeyTips mode without triggering a command at any time by pressing the Alt key again.

There are other shortcut keys that don't use the ribbon. These are key combinations that start with the Ctrl key. For example, Ctrl+C copies highlighted text and Ctrl+S saves your work. Usually, you find out about a shortcut key by hovering over a command with the mouse. For example, hover over the Paste button in the ribbon's Home tab, and you see a tooltip that tells you its timesaving shortcut key is Ctrl+V. And if you've worked with a previous version of Excel, you'll find that Excel 2010 keeps all of the same shortcut keys.

NOSTALGIA CORNER

Excel 2003 Menu Shortcuts

If you've worked with a previous version of Excel, you might have trained yourself to use menu shortcuts–key combinations that open a menu and pick out the command you want. For example, if you press Alt+E in Excel 2003, the Edit menu pops open. You can then press the S key to choose the Paste Special command.

At first glance, it doesn't look like these keyboard shortcuts will amount to much in Excel 2010. After all, Excel 2010 doesn't even have a corresponding series of menus! Fortunately, Microsoft went to a little extra trouble to make life easier for longtime Excel aficionados. The result is that you can still use your menu shortcuts, but they work in a slightly different way.

When you hit Alt+E in Excel 2010, you see a tooltip appear over the top of the ribbon (Figure I-7) that lets you know you've started to enter an Excel 2003 menu shortcut. If you go on to press S, you wind up at the familiar Paste Special dialog box, because Excel knows what you're trying to do. It's almost as though Excel has an invisible menu at work behind the scenes.

Of course, this feature can't help you out all the time. It doesn't work if you're trying to use one of the few commands that don't exist any longer. And if you need to see the menu to remember what key to press next, you're out of luck. All Excel gives you is the tooltip.

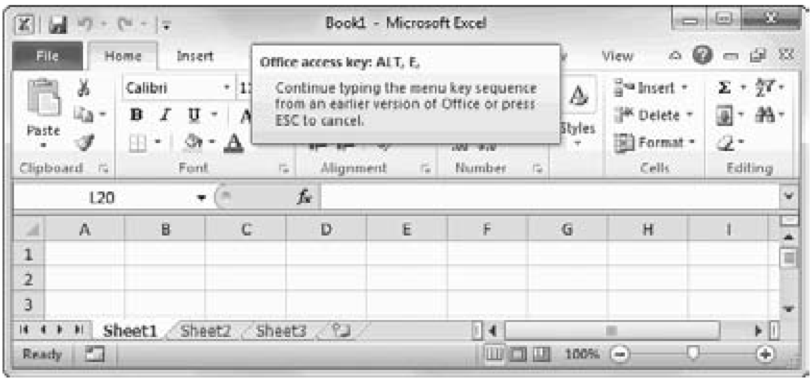

Figure I-7: *By pressing Alt+E, you've triggered the "imaginary" Edit menu from Excel 2003 and earlier versions. You can't actually see it (because in Excel 2010 this menu doesn't exist). However, the tooltip lets you know that Excel is paying attention. You can now complete your action by pressing the next key for the menu command you're nostalgic for.*

Backstage View

Excel 2010 doesn't introduce anything earth-shattering as the ribbon. However, it does have another not-so-small change to the way the program operates. Instead of sending you to an ordinary menu to open files, create them, and print your work, it devotes the entire window to these tasks—once you switch into a mode called *backstage view*.

To switch to backstage view, click the File button that appears just to the left of the Home tab in the ribbon. (The name of this button is a nod to Excel 2003 and other older, more traditional Windows programs, which group many of these tasks together in a menu named File.) To get out of backstage view, just click File again or press the Esc key.

Backstage view is split into two parts. On the left is a narrow strip of different commands. You click one of these to show a page for a different task. Depending on what you click, Excel may show additional options and information on the right, as shown in Figure I-8.

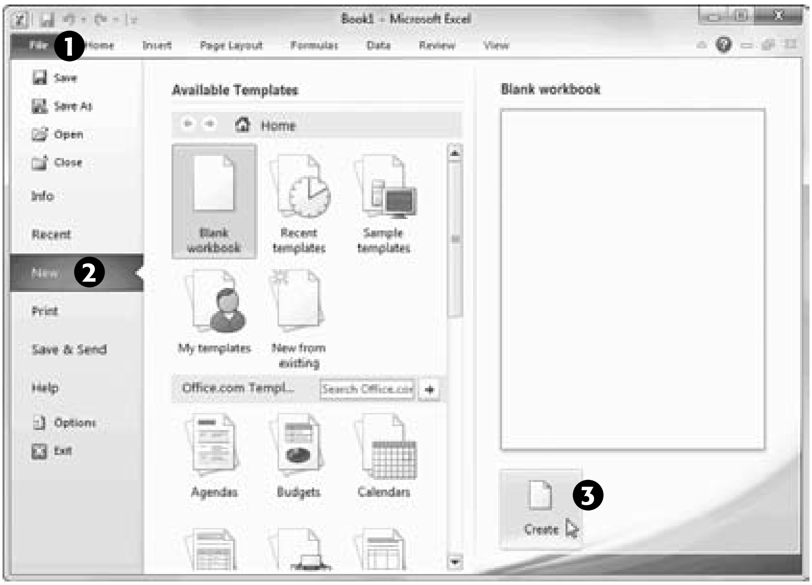

Figure I-8: *To create a new Excel workbook, start by clicking the File tab (1) and then the New command on the left (2). Excel uses its big backstage view to show you a slew of options. For a no-fuss blank workbook, just leave "Blank workbook" selected and click the Create button on the right (3). Or, if you want to get a head start with a premade template, choose one of the many other options underneath. (Chapter 16 tackles templates in detail.)*

Here are some of the things you'll do in Excel's backstage view:

- Work with files (creating, opening, closing, and saving them). You'll do plenty of this in Chapter 1.
- Print your work (Chapter 7) and send it off to other people by email (Chapter 25).
- Prepare a workbook to be shared with other people. For example, you can check its compatibility with old versions of Excel (Chapter 1) and use document protection to prevent other people from changing your numbers (Chapter 24).
- Configure how Excel behaves. Once you're in backstage view, just click the Options command to get to the Excel Options dialog box, an all-in-one place for configuring Excel (page 37).

The Quick Access Toolbar

Keen eyes will have noticed the tiny bit of screen real estate on the right side of the Office button, just above the ribbon. It holds a series of tiny icons, like the toolbars in older versions of Excel (Figure I-9). This is the Quick Access toolbar (or QAT, to Excel nerds).

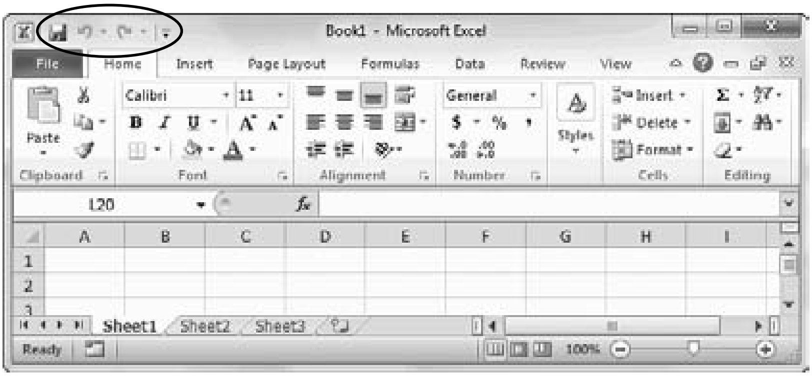

Figure I-9: *The Quick Access toolbar puts the Save, Undo, and Redo command right at your fingertips. These commands are singled out because most people use them more frequently than any other commands. But as you'll learn in the Appendix, you can add anything you want here.*

If the Quick Access toolbar were nothing but a specialized shortcut for three commands, it wouldn't be worth the bother. However, the nifty thing about the Quick Access toolbar is that you can customize it. In other words, you can remove commands you don't use and add your own favorites. The appendix of this book shows how.

Microsoft has deliberately kept the Quick Access toolbar very small. It's designed to provide a carefully controlled outlet for those customization urges. Even if you go wild stocking the Quick Access toolbar with your own commands, the rest of the ribbon remains unchanged. (And that means a co-worker or spouse can still use your computer without suffering a migraine.)

Excel's New Features

As you've already learned, Excel 2010 isn't nearly as radical as Excel 2007, which revamped the program's main window and introduced the now-infamous ribbon. However, Excel 2010 still has its share of enhancements. Here are the most important:

- **Backstage view.** Now Excel has a single go-to place for managing files. Whether you need to open an existing spreadsheet, create a new one, print your work, or tune up Excel options, Excel's backstage view gives you a bit more breathing room. You'll learn more about this mega-timesaver in Chapter 1.

- **Better AutoRecover.** Excel has always had an emergency failsafe feature that preserves your work in the event of computer disaster (for example, a power failure). But now you can use the same feature to rescue your work after a personal blunder (for example, forgetting to save your work or suddenly realizing you deleted a critical piece of information 40 minutes ago). That's because Excel 2010 stores a list of automatically backed-up files and lets you open them if you need to make an emergency recovery (page 82).
- **Paste preview.** Now you can see what the result of a copy-and-paste will be before you actually do it. It seems like a minor refinement, but when you're using advanced paste options—for example, trying to decide whether to copy your numbers and their formatting—it's a huge convenience. Page 106 has more.
- **Sparklines.** These miniature graphs are the hottest innovation in data display since the invention of the pie chart. They're small (so they can fit into a single cell), but still big enough for you to see trends and summaries at a glance. If a full-fledged chart seems like overkill in your spreadsheet, a sparkline just might be what you're looking for. Chapter 20 describes sparklines in detail.
- **Protected view.** Do you live in fear of marauding Excel viruses corrupting your computer? Well, even if you don't, you'll be happy to know that Excel opens potentially dangerous spreadsheets in a carefully restricted window, ensuring they can't get up to any trouble. Thanks to this change, you can open a spreadsheet straight from the Web or your email inbox, without worry. Page 57 shows how it works.
- **Slicers.** If you've ever run into Excel's pivot tables, you know that they're powerful but complicated. But life gets a bit easier with slicers (page 645), a new feature that gives you a visual way to slice and dice your pivot table data, filtering it down to the details that interest you the most.
- **Easier ribbon customization.** In Excel 2007, you couldn't change the ribbon unless you mastered an intimidating extensibility model based on XML. Now, you just need a leisurely trip to the Customize Ribbon section of the Excel Options dialog box, where you can add, remove, and reorder Excel's panoply of buttons to suit your personal preferences. The Appendix shows you how.
- **The Excel Web App.** Wouldn't it be cool to view your Excel spreadsheet in a web browser? And wouldn't it be even better if you could edit it and share with other people? And wouldn't it be just a little mind-blowing if you could use the Excel Web App to work on an Excel spreadsheet, even if you didn't have Excel installed on your computer? For the first time, Excel 2010 makes these scenarios possible. You'll get the scoop in Chapter 26.

Of course, this list is by no means complete. Excel 2010 is chock-full of refinements, tweaks, and tune-ups that make it easier to use than any previous version. You'll learn all the best tricks throughout this book.

About This Book

Despite the many improvements in software over the years, one feature hasn't improved a bit: Microsoft's documentation. In fact, with Office 2010, you get no printed user guide at all. To learn about the thousands of features included in this software collection, Microsoft expects you to read the online help.

Occasionally, the online help is actually helpful, like when you're looking for a quick description explaining a mysterious new function. On the other hand, if you're trying to learn how to, say, create an attractive chart, you're stuck with terse and occasionally cryptic instructions.

The purpose of this book, then, is to serve as the manual that should have accompanied Excel 2010. In these pages, you'll find step-by-step instructions and tips for using almost every Excel feature, including those you may not even know exist.

About the Outline

This book is divided into eight parts, each containing several chapters.

- **Part One: Worksheet Basics.** In this part, you'll get acquainted with Excel's interface and learn the basic techniques for creating spreadsheets and entering and organizing data. You'll also learn how to format your work to make it more presentable, and how to create sharp printouts.
- **Part Two: Formulas and Functions.** This part introduces you to Excel's most important feature—formulas. You'll learn how to perform calculations ranging from the simple to the complex, and you'll tackle specialized functions for dealing with all kinds of information, including scientific, statistical, business, and financial data.
- **Part Three: Organizing Your Information.** The third part covers how to organize and find what's in your spreadsheet. First, you'll learn to search, sort, and filter large amounts of information by using tables. Next, you'll see how to boil down complex tables with grouping and outlining. Finally, you'll turn your perfected spreadsheets into reusable templates.
- **Part Four: Charts and Graphics.** The fourth part introduces you to charting and graphics, two of Excel's most popular features. You'll learn about the wide range of different chart types available and when it makes sense to use each one. You'll also find out how you can use pictures to add a little pizazz to your spreadsheets.
- **Part Five: Advanced Data Analysis.** In this short part, you'll tackle some of the more advanced features that people often overlook or misunderstand. You'll see how to study different possibilities with scenarios, use goal seeking and the Solver add-in to calculate "backward" and fill in missing numbers, and create multi-layered summary reports with pivot tables.

- **Part Six: Sharing Data with the Rest of the World.** The sixth part explores ways that you can share your spreadsheets with other people. You'll learn how to collaborate with colleagues to revise a spreadsheet, without letting mistakes creep in or losing track of who did what. You'll also learn how to copy Excel tables and charts into other programs (like Word) and how to use the Excel Web App to share and edit spreadsheet on the Web.
- **Part Seven: Programming Excel.** This part presents a gentle introduction to the world of Excel programming, first by recording macros and then by using the full-featured VBA (Visual Basic for Applications) language, which lets you automate complex tasks.
- **Part Eight: Appendix.** The end of this book wraps up with an appendix that shows how to customize the ribbon to get easy access to your favorite commands.

About→These→Arrows

Throughout this book, you'll find sentences like this one: "Choose Insert→Illustrations→Picture." This a shorthand way of telling you how to find a feature in the Excel ribbon. It translates to the following instructions: "Click the **Insert** tab of the toolbar. On the tab, look for the **Illustrations** section. In the Illustrations box, click the **Picture** button." Figure I-10 shows the button you want.

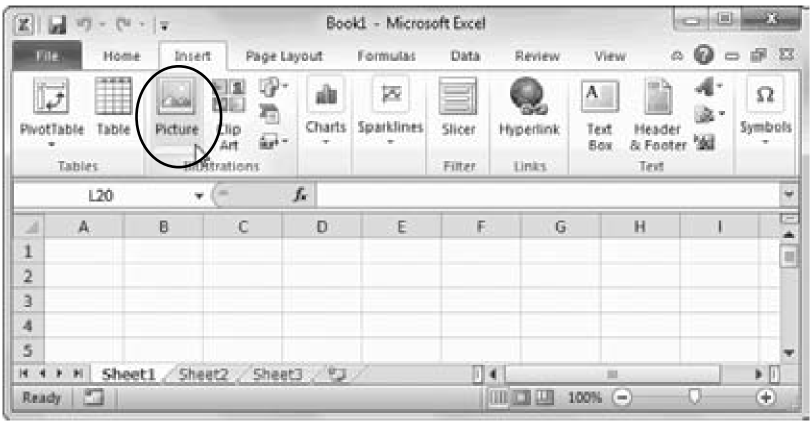

Figure I-10: *In this book, arrow notations help to simplify ribbon commands. For example, "Choose Insert→Illustrations→Picture" leads to the highlighted button shown here.*

Note: As you saw back in Figure I-4, the ribbon adapts itself to different screen sizes. Depending on the size of your Excel window, it's possible that the button you need to click won't include any text. Instead, it shows up as a small icon. In this situation, you can hover over the mystery button to see its name before deciding whether to click it.

Contextual tabs

There are some tabs that appear in the ribbon only when you're working on specific tasks. For example, when you create a chart, a Chart Tools section appears with three new tabs (see Figure I-11).

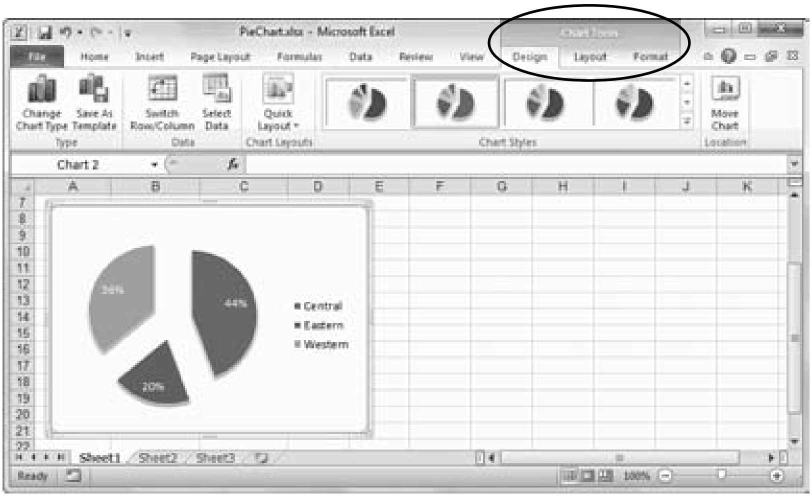

Figure I-11: *Excel doesn't bother to show these three tabs unless you're working on a chart, because it's frustrating to look at a bunch of buttons you can't use. This sort of tab, which appears only when needed, is called a contextual tab.*

When dealing with contextual tabs, the instructions in this book always include the title of the tab section (it's Chart Tools in Figure I-11). Here's an example: "Choose Chart Tools | Design→Type→Change Chart Type." Notice that the first part of this instruction includes the tab section title (Chart Tools) and the tab name (Design), separated by the | character. That way, you can't mistake the Chart Tools | Design tab for a Design tab in some other group of contextual tabs.

Buttons with menus

From time to time, you'll encounter buttons in the ribbon that have short menus attached to them. Depending on the button, this menu might appear as soon as you click the button, or it might appear only if you click the button's drop-down arrow, as shown in Figure I-12.

When dealing with this sort of button, the last step of the instructions in this book tells you what to choose from the drop-down menu. For example, say you're directed to "Home→Clipboard→Paste→Paste Special." That tells you to select the Home tab, look for the Clipboard section, click the drop-down part of the Paste button (to reveal the menu with extra options), and then choose Paste Special from the menu.

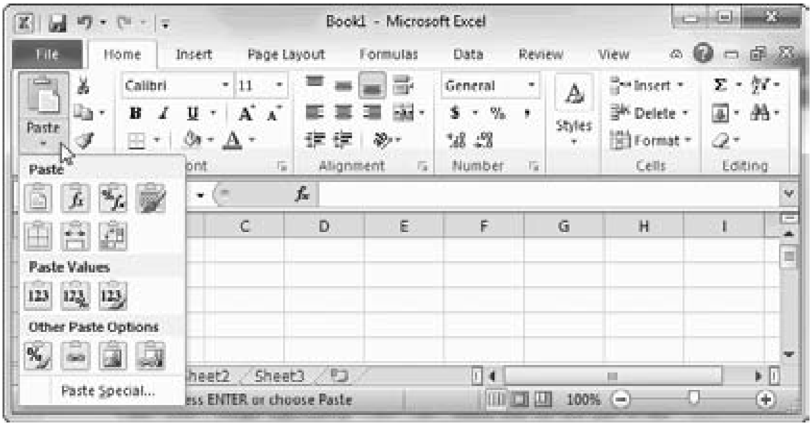

Figure I-12: *There are several options for pasting text from the clipboard. Click the top part of the Paste button to perform a plain-vanilla paste (with all the standard settings), or click the bottom part to see the menu of choices shown here.*

Tip: Be on the lookout for drop-down arrows in the ribbon—they're tricky at first. You need to click the ***arrow*** part of the button to see the full list of options. When you click the other part of the button, you don't see the list. Instead, Excel fires off the standard command (the one Excel thinks is the most common choice) or the command you used most recently.

Dialog box launchers

As powerful as the ribbon is, you can't do everything using the buttons it provides. Sometimes you need to use a good ol'-fashioned dialog box. (A *dialog box* is a term used in the Windows world to describe a small window with a limited number of options. Usually, dialog boxes are designed for one task and aren't resizable, although software companies like Microsoft break these rules all the time.)

There are two ways to get to a dialog box in Excel 2010. First, some ribbon buttons take you there straight away. For example, if you choose Home→Clipboard→Paste→Paste Special, you always get a dialog box. There's no way around it.

The second way to get to a dialog box is through something called a *dialog box launcher*, which is just a nerdified name for the tiny square-with-arrow icon that sometimes appears in the bottom-right corner of a section of the ribbon. The easiest way to learn how to spot a dialog box launcher is to look at Figure I-13.

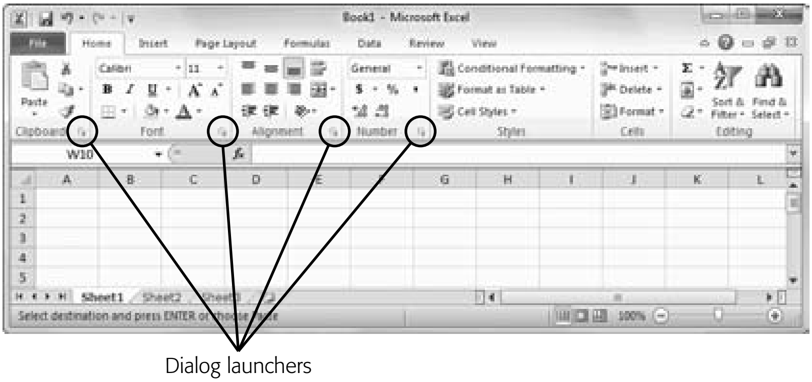

Figure I-13: *As you can see here, the Clipboard, Font, Alignment, and Number sections all have dialog box launchers. The Styles, Cells, and Editing sections don't.*

When you click a dialog box launcher, the related dialog box appears. For example, click the dialog box launcher for the Font section and you get a full Font dialog box that lets you scroll through all the typefaces on your computer, choose the size and color, and so on.

In this book, there's no special code word that tells you to use a dialog box launcher. Instead, you'll see an instruction like this: "To see more font options, look at the Home→Font section and click the dialog box launcher (the small icon in the bottom-right corner)." Now that you know what a dialog box launcher is, that makes perfect sense.

Backstage view

If you see an instruction that includes arrows but starts with the word File, it's telling you to go to Excel's backstage view. For example, the sentence "Choose File→New" means click the File button to switch to backstage view, then click the New command (which appears in the narrow list on the left). To take another look at backstage view and the list of commands it offers, jump back to Figure I-8 on page 9.

Ordinary menus

There are also a couple of cases where you'll still use the familiar Windows menu. One example is when you use the Visual Basic editor (in Chapter 29). In this case, the arrows refer to menu levels. For example the instruction "Choose File→Save" means "Click the File menu heading. Then, on the File menu, click the Save command."

About Shortcut Keys

Every time you take your hand off the keyboard to move the mouse, you lose a few microseconds. That's why many experienced computer fans use keystroke combinations instead of toolbars and menus wherever possible. Ctrl+S, for example, is a keyboard shortcut that saves your current work in Excel (and most other programs).

When you see a shortcut like Ctrl+S in this book, it's telling you to hold down the Ctrl key and, while it's down, press the letter S, and then release both keys. Similarly, the finger-tangling shortcut Ctrl+Alt+S means hold down Ctrl, and then press and hold Alt, and then press S (so that all three keys are held down at once).

Examples

As you read this book, you'll see a number of examples that demonstrate Excel features and techniques for building good spreadsheets. Most of these examples are available as Excel workbook files in a separate download. Just surf to *www.missingmanuals.com/cds* and click the link for this book to visit a page where you can download a ZIP file that includes the examples, organized by chapter.

About MissingManuals.com

At *http://www.missingmanuals.com*, you'll find news, articles, and updates to the books in the Missing Manual series.

But the website also offers corrections and updates to this book (to see them, click the book's title, and then click Errata). In fact, you're invited and encouraged to submit such corrections and updates yourself. In an effort to keep the book as up-to-date and accurate as possible, each time we print more copies of this book, we'll make any confirmed corrections you've suggested. We'll also note such changes on the website, so that you can mark important corrections in your own copy of the book.

In the meantime, we'd love to hear your own suggestions for new books in the Missing Manual series. There's a place for that on the website, too, as well as a place to sign up for free email notification of new titles in the series.

Safari® Books Online

Safari® Books Online is an on-demand digital library that lets you easily search over 7,500 technology and creative reference books and videos to find the answers you need quickly.

With a subscription, you can read any page and watch any video from our library online. Read books on your cellphone and mobile devices. Access new titles before they're available for print, and get exclusive access to manuscripts in development and post feedback for the authors. Copy and paste code samples, organize your favorites, download chapters, bookmark key sections, create notes, print out pages, and benefit from tons of other time-saving features.

O'Reilly Media has uploaded this book to the Safari Books Online service. To have full digital access to this book and others on similar topics from O'Reilly and other publishers, sign up for free at *http://my.safaribooksonline.com*.

CHAPTER
1

Creating Your First Spreadsheet

Every Excel grandmaster needs to start somewhere. In this chapter, you'll learn how to create a basic spreadsheet. First, you'll learn to move around Excel's grid of cells, typing in numbers and text as you go. Next, you'll take a quick tour of the Excel window, stopping to meet the different tabs in the ribbon and take a quick peek at the formula bar. Finally, you'll get a tour of Excel's innovative *backstage view*—the file-management hub where you can save your work for posterity, open recent files, and tweak Excel options.

Note: Even if you're an Excel old-timer, don't bypass this chapter. Although you already know how to fill in a simple spreadsheet, you haven't seen Excel's backstage view, which is a completely new feature in Excel 2010. It gives you a single, streamlined place to perform a whole variety of tasks, most of which have to do with managing your files. To skip the review and pick up your backstage pass, jump straight to page 36.

Creating a Basic Worksheet

When you first launch Excel, it starts you off with a new, blank *worksheet*, as shown in Figure 1-1. A worksheet is the grid of cells where you type your information and formulas. This grid takes up most of the Excel window. It's where you'll perform all your work, such as entering data, writing formulas, and reviewing the results.

Here are a few basics about Excel's grid:

- **The grid divides your worksheet into rows and columns**. Columns are identified with letters (A, B, C…), while rows are identified with numbers (1, 2, 3…).

- **The smallest unit in your worksheet is the** *cell*. Cells are identified by column and row. For example, C6 is the address of a cell in column C (the third column) and row 6 (the sixth row). Figure 1-2 shows this cell, which looks like a rectangular box. Incidentally, an Excel cell can hold up to 32,000 characters.
- **A worksheet can span an eye-popping 16,000 columns and 1 million rows**. In the unlikely case that you want to go beyond those limits—say, if you're tracking blades of grass on the White House lawn—you'll need to create a new worksheet. Every spreadsheet file can hold a virtually unlimited number of worksheets, as you'll learn in Chapter 4.
- **When you enter information, enter it one cell at a time**. However, you don't have to follow any set order. For example, you can start by typing information into cell A40 without worrying about filling any data in the cells that appear in the earlier rows.

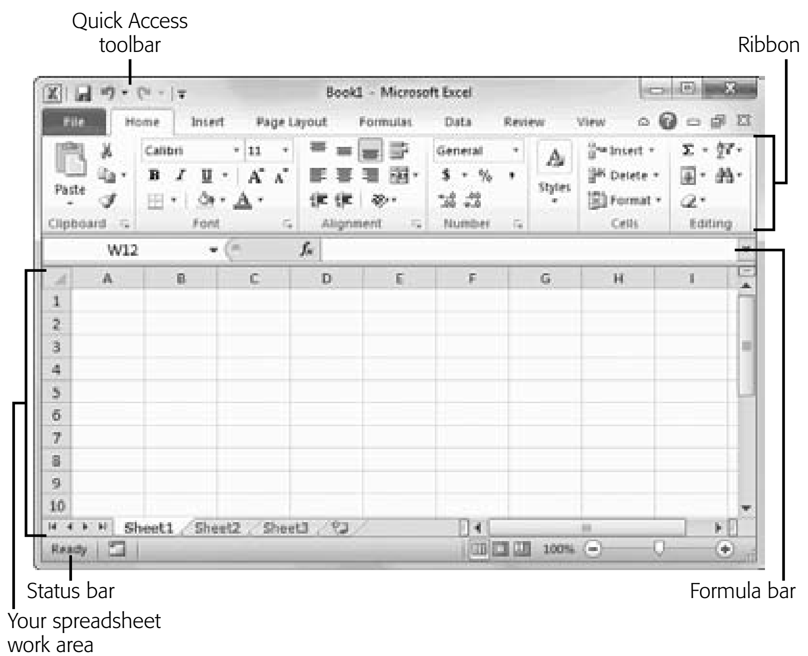

Figure 1-1: *The largest part of the Excel window is the worksheet grid, where you type in your information.*

Note: Obviously, once you go beyond 26 columns, you run out of letters. Excel handles this by doubling up (and then tripling up) letters. For example, column Z is followed by column AA, then AB, then AC, all the way to AZ and then BA, BB, BC—you get the picture. And if you create a ridiculously large worksheet, you'll find that column ZZ is followed by AAA, AAB, AAC, and so on.

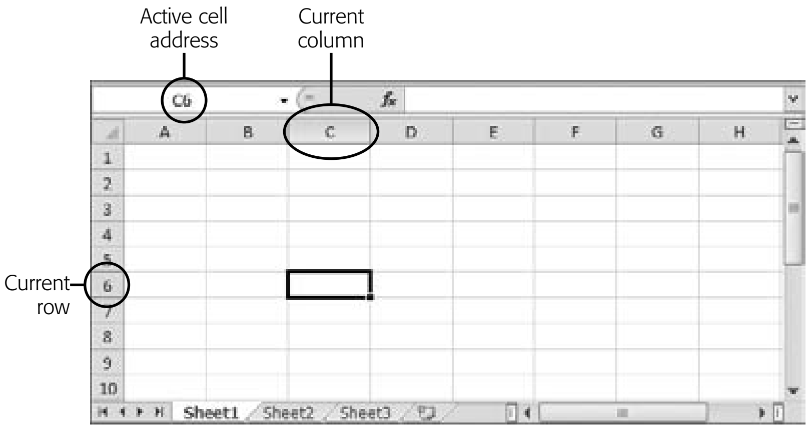

***Figure 1-2:** Here, the current cell is C6. You can recognize the current (or active) cell based on its heavy black border. You'll also notice that the corresponding column letter (C) and row number (6) are highlighted at the edges of the worksheet. Just above the worksheet, on the left side of the window, the formula bar tells you the active cell address.*

The best way to get a feel for Excel is to dive right in and start putting together a worksheet. The following sections cover each step that goes into assembling a simple worksheet. This one tracks household expenses, but you can use the same approach to create any basic worksheet.

Starting a New Workbook

When you fire up Excel, it opens a fresh workbook file. If you've already got Excel open and you want to create *another* workbook, just choose File→New (where File switches you into Excel's backstage view, and New is a command on the left side of the window). You'll see a variety of options for creating specialized types of spreadsheets. But to get started with a blank canvas, keep "Blank workbook" selected and click the Create button, as shown in Figure 1-3.

Note: A workbook is a collection of one or more worksheets. That distinction isn't terribly important now because you're using only a single worksheet in each workbook you create. But in Chapter 4, you'll learn how to use several worksheets in the same workbook to track related collections of data.

For now, all you need to know is that the worksheet is the grid of cells where you place your data, and the workbook is the spreadsheet file that you save on your computer.

You don't need to pick the file name for your workbook when you first create it. Instead, that decision happens later, when you *save* your workbook. For now, you start with a blank canvas that's ready to receive your numerical insights.

Note: Creating new workbooks doesn't disturb what you've already done. Whatever workbook you were using remains open in another window. You can use the taskbar to move from one workbook to the other, or you can use the Excel shortcuts explained on page 60.

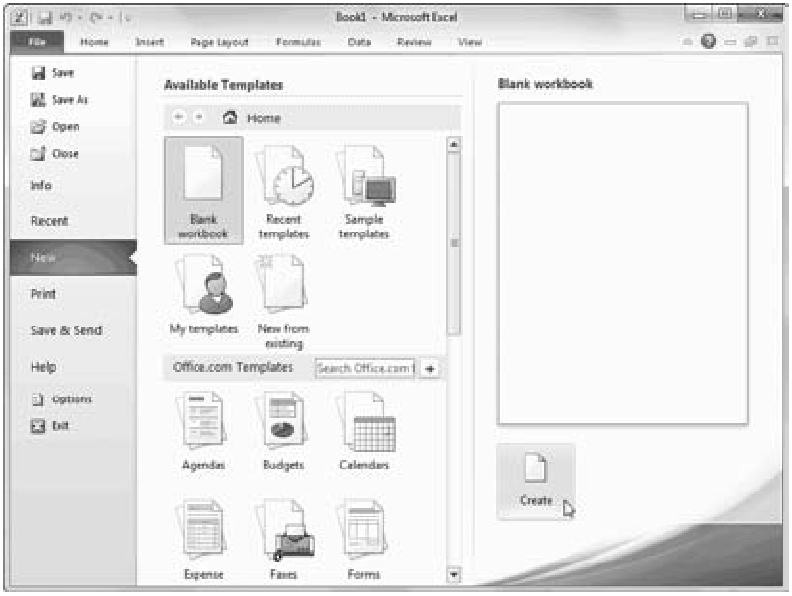

Figure 1-3:
Excel lets you create a new, blank workbook or a ready-made workbook from a template. For now, choose "Blank workbook" and then click the big Create button in the bottom-right corner. You'll learn about using (and making) templates in Chapter 16.

Adding the Column Titles

Excel allows you to arrange information in whatever way you like. There's nothing to stop you from scattering numbers left and right, across as many cells as you can. However, one of the most common (and most useful) ways to arrange your information is as a table, with headings for each column.

It's important to remember that even for the simplest worksheet, the decisions you make about what's going to go in each column can have a big effect on how easy it is to manipulate your information. For example, in a worksheet that stores a mailing list, you *could* have two columns: one for names and another for addresses. But if you create more than two columns, your life will probably be easier since you can separate first names from street addresses from ZIP codes, and so on. Figure 1-4 shows the difference.

You can, of course, always add or remove columns later. But you can avoid getting gray hairs by starting a worksheet with all the columns you think you'll need.

The first step in creating your worksheet is to add your headings in the row of cells at the top of the worksheet (row 1). Technically, you don't need to start right in the first row, but unless you want to add more information before your table—like a title for the chart or today's date—there's no point in wasting the space. Adding information is easy—just click the cell you want and start typing. When you're finished, hit Tab to complete your entry and move to the next cell to the right (or Enter to head to the cell just underneath).

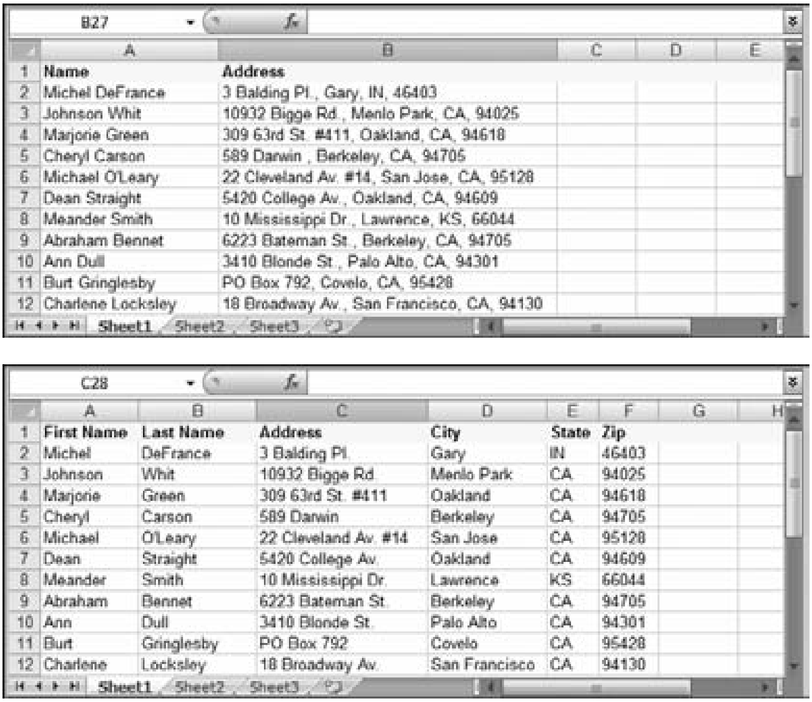

B27

	A	B	C	D	E
1	Name	Address			
2	Michel DeFrance	3 Balding Pl., Gary, IN, 46403			
3	Johnson Whit	10932 Bigge Rd., Menlo Park, CA, 94025			
4	Marjorie Green	309 63rd St. #411, Oakland, CA, 94618			
5	Cheryl Carson	589 Darwin , Berkeley, CA, 94705			
6	Michael O'Leary	22 Cleveland Av. #14, San Jose, CA, 95128			
7	Dean Straight	5420 College Av., Oakland, CA, 94609			
8	Meander Smith	10 Mississippi Dr., Lawrence, KS, 66044			
9	Abraham Bennet	6223 Bateman St., Berkeley, CA, 94705			
10	Ann Dull	3410 Blonde St., Palo Alto, CA, 94301			
11	Burt Gringlesby	PO Box 792, Covelo, CA, 95428			
12	Charlene Locksley	18 Broadway Av., San Francisco, CA, 94130			

Sheet1 Sheet2 Sheet3

C28

	A	B	C	D	E	F	G	H
1	First Name	Last Name	Address	City	State	Zip		
2	Michel	DeFrance	3 Balding Pl.	Gary	IN	46403		
3	Johnson	Whit	10932 Bigge Rd.	Menlo Park	CA	94025		
4	Marjorie	Green	309 63rd St. #411	Oakland	CA	94618		
5	Cheryl	Carson	589 Darwin	Berkeley	CA	94705		
6	Michael	O'Leary	22 Cleveland Av. #14	San Jose	CA	95128		
7	Dean	Straight	5420 College Av.	Oakland	CA	94609		
8	Meander	Smith	10 Mississippi Dr.	Lawrence	KS	66044		
9	Abraham	Bennet	6223 Bateman St.	Berkeley	CA	94705		
10	Ann	Dull	3410 Blonde St.	Palo Alto	CA	94301		
11	Burt	Gringlesby	PO Box 792	Covelo	CA	95428		
12	Charlene	Locksley	18 Broadway Av.	San Francisco	CA	94130		

Sheet1 Sheet2 Sheet3

Figure 1-4:
Top: If you enter the first and last names together in one column, Excel can sort only by the first names. And if you clump the addresses and ZIP codes together, you give Excel no way to count how many people live in a certain town or neighborhood, because Excel can't extract the ZIP codes.

Bottom: The benefit of a six-column table is significant: It lets you sort (reorganize) your list according to people's last names or where they live. It also allows you to filter out individual bits of information when you start using functions later in this book.

Note: The information you put in an Excel worksheet doesn't need to be in neat, ordered columns. Nothing stops you from scattering numbers and text in random cells. However, most Excel worksheets resemble some sort of table, because that's the easiest and most effective way to deal with large amounts of structured information.

For a simple expense worksheet designed to keep a record of your most prudent and extravagant purchases, try the following three headings:

- **Date Purchased** stores the date when you spent the money.
- **Item** stores the name of the product that you bought.
- **Price** records how much it cost.

Right away, you face your first glitch: awkwardly crowded text. Figure 1-5 shows how you can adjust column width for proper breathing room.

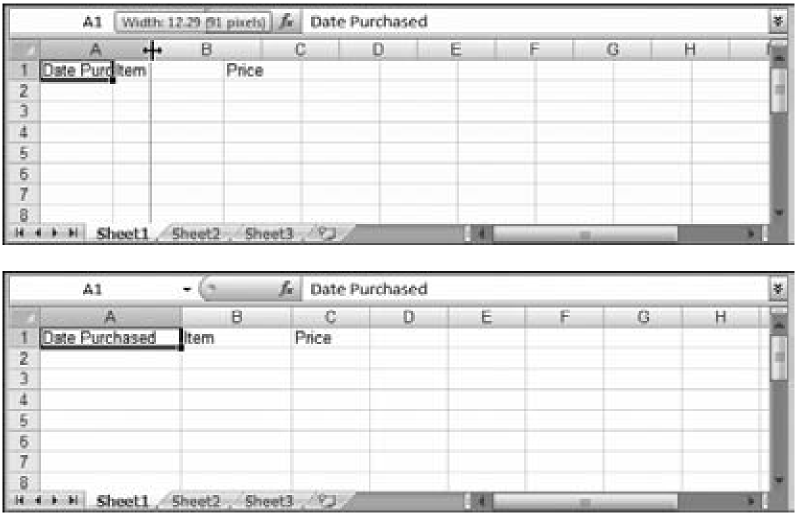

Figure 1-5: *Top: The standard width of an Excel column is 8.43 characters, which hardly allows you to get a word in edgewise. To solve this problem, position your mouse on the right border of the column header you want to expand so that the mouse pointer changes to the resize icon (it looks like a double-headed arrow). Now drag the column border to the right as far as you want. As you drag, a tooltip appears, telling you the character size and pixel width of the column. Both of these pieces of information play the same role–they tell you how wide the column is–only the unit of measurement changes.*

Bottom: When you release the mouse, the entire column of cells is resized to the new size.

Note: A column's character width doesn't really reflect how many characters (or letters) fit in a cell. Excel uses ***proportional*** fonts, in which different letters take up different amounts of room. For example, the letter W is typically much wider than the letter 1. All this means is that the character width Excel shows you isn't a real indication of how many letters can fit in the column, but it's still a useful measurement that you can use to compare different columns.

Adding Data

You can now begin adding your data: simply fill in the rows under the column titles. Each row in the expense worksheet represents a separate purchase you've made. (If you're familiar with databases, you can think of each row as a separate record.)

As Figure 1-6 shows, the first column is for dates, the second column is for text, and the third column holds numbers. Keep in mind that Excel doesn't impose any rules on what you type, so you're free to put text in the Price column. But if you don't keep a consistent kind of data in each column, you won't be able to easily analyze (or understand) your information later.

Figure 1-6: *This rudimentary expense list has three items (in rows 2, 3, and 4). The alignment of each column reflects the data type (by default, numbers and dates are right-aligned, while text is left-aligned), indicating that Excel understands your date and price information.*

That's it. You've now created a living, breathing worksheet. The next section explains how you can edit the data you've entered.

Editing Data

Every time you start typing in a cell, Excel erases any existing content in that cell. (You can also quickly remove the contents of a cell by just moving to it and pressing Delete, which clears its contents.)

If you want to edit cell data instead of replacing it, you need to put the cell in *edit mode*, like this:

1. **Move to the cell you want to edit.**

   Use the mouse or the arrow keys to get to the correct cell.

2. **Put the cell in edit mode by pressing F2.**

   Edit mode looks almost the same as ordinary text entry mode. The only difference is that you can use the arrow keys to move through the text you're typing and make changes. (When you aren't in edit mode, pressing these keys just moves you to another cell.)

   If you don't want to use F2, you can also get a cell into edit mode by double-clicking it.

3. **Complete your edit.**

   Once you've modified the cell content, press Enter to make your change or Esc to cancel your edit and leave the old value in the cell. Alternatively, you can click on another cell to accept the current value and go somewhere else. But while you're in edit mode, you can't use the arrow keys to move out of the cell.

Tip: If you start typing new information into a cell and you decide you want to move to an earlier position in your entry (to make an alteration, for instance), just press F2. The cell box still looks the same, but you're now in edit mode, which means that you can use the arrow keys to move within the cell (instead of moving from cell to cell). You can press F2 again to return to the normal data entry mode, which allows you to use the arrow keys to move to another cell.

As you enter data, you may discover the Bigtime Excel Display Problem (known to aficionados as BEDP): Cells in adjacent columns can overlap one another. Figure 1-7 shows the problem. One way to fix BEDP is to manually resize the column, as shown in Figure 1-5. Another option is to use wrapping to fit multiple lines of text in a single cell, as described on page 161.

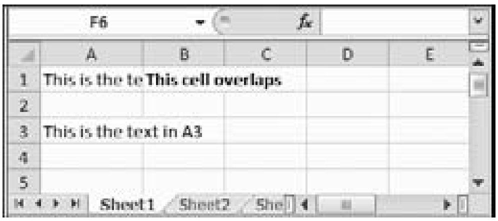

Figure 1-7:
Overlapping cells can create big headaches. For example, if you type a large amount of text into A1, and then you type some text into B1, you see only part of the data in A1 on your worksheet (as shown here). The rest is hidden from view. But if, say, A3 contains a large amount of text and B3 is empty, the content in A3 is displayed over both columns, and you don't have a problem.

Moving Around the Grid

Learning how to move around the Excel grid quickly and confidently is an indispensable skill. To move from cell to cell, you have two fairly obvious choices:

- **Use the arrow keys on the keyboard**. Keystrokes move you one cell at a time in any direction.
- **Click the cell with the mouse**. A mouse click jumps you directly to the cell you've clicked.

As you move from cell to cell, you see the black focus box move to highlight the currently active cell.

In some cases, you might want to cover ground a little quicker. One option is to use the scrollbars at the bottom and on the right side of the window to scroll off into the uncharted regions of your worksheet. Excel also provides two more powerful features—shortcut keys and the Go To feature—which are described in the following sections.

Shortcut Keys

Excel provides a number of handy key combinations that can transport you across your worksheet in great leaps and bounds (see Table 1-1). The most useful shortcut keys include the Home key combinations, which bring you back to the beginning of a row or the top of your worksheet.

POWER USERS' CLINIC

Using R1C1 Reference Style

Most people like to identify columns with letters and rows with numbers. This system makes it easy to tell the difference between the two, and it allows you to use short cell addresses like A10, B4, H99, and so on. When you first install Excel, it uses this style of cell addressing.

However, Excel also lets you use another cell addressing system called *R1C1*. In R1C1 style, columns and rows are both identified with numbers. That means the cell address A10 becomes R10C1 (read this as Row 10, Column 1). The letters R and C tell you which part of the address is the row and which part is the column. The order is reversed from conventional cell addressing.

R1C1 addressing isn't all that common. But it becomes useful if you need to deal with worksheets that have more than 26 columns. With normal cell addressing, Excel runs out of letters after column 26, and it starts using two letters (as in AA, AB, AC, and so on) to identify columns. This approach can become somewhat awkward. For example, if you want to find cell AX1, it isn't immediately obvious that this cell is in column 50. On the other hand, the R1C1 address for the same cell—R1C50—gives you a clearer idea of where to find the cell.

To use R1C1 for a spreadsheet, select File→Options. This shows the Excel Options window, where you can change a wide array of settings. In the list on the left, choose Formulas to hone in on the section you need. Then, look under the "Working with formulas" heading, and turn on the "R1C1 reference style" checkbox.

R1C1 is a file-specific setting, which means that if someone sends you a spreadsheet that he saved with R1C1, you'll see the R1C1 cell addresses when you open the file, regardless of what type of cell addressing you use in your own spreadsheets. Fortunately, you can change cell addressing at any time using the Excel Options window.

Note: Shortcut key combinations that use the + sign must be entered together. For example, "Ctrl+Home" means you hold down Ctrl and press Home at the same time. Key combinations with a comma work in sequence. For example, the key combination "End, Home" means press End first, release it, and then press Home.

Table 1-1. *Shortcut keys for moving around a worksheet*

Key combination	Result
→ (or Tab)	Moves one cell to the right.
← (or Shift+Tab)	Moves one cell to the left.
↑	Moves one cell up.
↓ (or Enter)	Moves one cell down.
Page Up	Moves up one screen. Thus, if the grid shows 10 cells at a time, this key moves to a cell in the same column, 10 rows up (unless you are already at the top of the worksheet).
Page Down	Moves down one screen. Thus, if the grid shows 10 cells at a time, this key moves to a cell in the same column, 10 rows down.
Home	Moves to the first cell (column A) of the current row.
Ctrl+Home	Moves to the first cell in the top row, which is A1.
Ctrl+End (or End, Home)	Moves to the last column of the last occupied row. This cell is at the bottom-right edge of your data.

Excel also lets you cross great distances in a single bound using a *Ctrl+arrow* key combination. These key combinations jump to the *edges* of your data. Edge cells include cells that are next to other blank cells. For example, if you press Ctrl+→ while you're inside a group of cells with information in them, you'll skip to the right, over all filled cells, and stop just before the next blank cell. If you press Ctrl+→ again, you'll skip over all the nearby blank cells and land in the next cell to the right that has information in it. If there aren't any more cells with data on the right, you'll wind up on the very edge of your worksheet.

The *Ctrl+arrow key* combinations are useful if you have more than one table of data in the same worksheet. For example, imagine you have two tables of data, one at the top of a worksheet and one at the bottom. If you are at the top of the first table, you can use Ctrl+↓ to jump to the bottom of the first table, skipping all the rows in between. Press Ctrl+↓ again, and you leap over all the blank rows, winding up at the beginning of the second table.

The Go To Feature

If you're fortunate enough to know exactly where you need to go, you can use the Go To feature to make the jump. Go To moves to the cell address you specify. It comes in useful in extremely large spreadsheets, where just scrolling through the worksheet takes half a day.

To bring up the Go To dialog box (shown in Figure 1-8), choose Home→Editing→Find & Select→Go To. Or you can do yourself a favor and just press Ctrl+G. Enter the cell address (such as C32), and then click OK.

Figure 1-8:
You'll notice that in the Go To list, cell addresses are written a little differently than the format you use when you type them in. Namely, dollar signs are added before the row number and column letter. Thus, C32 becomes $C$32, which is simply the convention that Excel uses for fixed cell references. (You'll learn much more about the different types of cell references in Chapter 8.)

The Go To feature becomes more useful the more you use it. That's because the Go To window maintains a list of the most recent cell addresses that you've entered. In addition, every time you open the Go To window, Excel automatically adds the current cell to the list. This feature makes it easy to jump to a far-off cell and quickly return to your starting location by selecting the last entry in the list.

The Go To window isn't your only option for leaping through a worksheet in a single bound. If you look at the Home→Editing→Find & Select menu, you'll find more specialized commands that let you jump straight to cells that contains formulas, comments, conditional formatting, and other advanced Excel ingredients that you haven't learned about yet. And if you want to hunt down cells that have specific text, you need the popular Find command (Home→Editing→Find & Select→Find), which is covered on page 127.

A Tour of the Excel Window

Finding your way around a worksheet is a fundamental part of mastering Excel. Knowing your way around the larger program window is no less important. The next few sections help you get oriented in the Excel window, pointing out the important stuff and letting you know what you can usually ignore.

The Tabs of the Ribbon

In the Introduction, you learned about the ribbon, the super-toolbar that offers one-stop shopping for all of Excel's features. All the most important Office applications—including Word, Access, PowerPoint, and Excel—use the ribbon. However, each program has a different set of tabs and buttons.

Throughout this book, you'll dig through the different tabs of the ribbon to find important features. But before you start your journey, it's nice to get a quick overview of what each tab provides. Here's the lowdown:

- **File** isn't really a toolbar tab, even though it appears first in the list. Instead, it's your gateway to Excel's backstage view, as described on page 36.
- **Home** includes some of the most commonly used buttons, like those for cutting and pasting information, formatting your data, and hunting down important bits of information with search tools. You've already used the Go To button on this tab (see the box "Getting Somewhere in a Hurry," on page 28).
- **Insert** lets you add special ingredients like tables, graphics, charts, and hyperlinks.
- **Page Layout** is all about getting your worksheet ready for the printer. You can tweak margins, paper orientation, and other page settings.
- **Formulas** are mathematical instructions that you use to perform calculations. This tab helps you build super-smart formulas and resolve mind-bending errors.
- **Data** lets you get information from an outside data source (like a heavy-duty database) so you can analyze it in Excel. It also includes tools for dealing with large amounts of information, like sorting, filtering, and subgrouping.
- **Review** includes the familiar Office proofing tools (like the spell checker). It also has buttons that let you add comments to a worksheet and manage revisions.

- **View** lets you switch on and off a variety of viewing options. It also lets you pull off a few fancy tricks if you want to view several separate Excel spreadsheet files at the same time.

GEM IN THE ROUGH

Collapsing the Ribbon

Most people are happy to have the ribbon sit at the top of the Excel window, with all its buttons on hand. However, serious number crunchers demand maximum space for their data. They'd rather look at another row of numbers than a pumped-up toolbar. If this describes you, then you'll be happy to find out you can *collapse* the ribbon, which shrinks it down to a single row of tab titles, as shown in Figure 1-9. To collapse it, just double-click the current tab title. (Or, you can click the tiny up-pointing icon that appears in the top-right corner of the ribbon, right next to the help icon.)

Even when the ribbon is collapsed, you can still use all its features. All you need to do is click a tab. For example, if you click Home, the Home tab pops up over your work-sheet. As soon as you click the button you want in the Home tab (or click a cell in your worksheet), the ribbon collapses itself again. The same trick works if you trigger a command in the ribbon using the keyboard, as described on page 5.

If you use the ribbon only occasionally, or if you prefer to use keyboard shortcuts, it makes sense to collapse the ribbon. Even when collapsed, the ribbon commands are available—it just takes an extra click to open the tab. On the other hand, if you make frequent trips to the ribbon or you're learning about Excel and like to browse the ribbon to see what features are available, don't bother collapsing it. The two or three rows that you'll lose are well worth it.

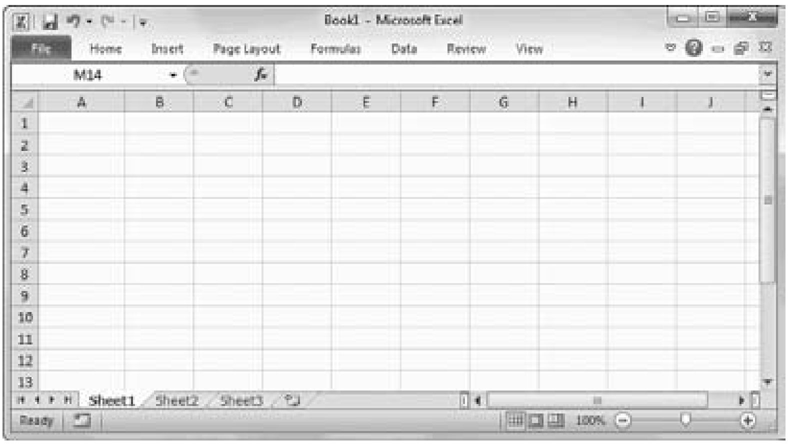

Figure 1-9: *Do you want to use every square inch of screen space for your cells? You can collapse the ribbon (as shown here) by double-clicking any tab. Click a tab to pop it open temporarily, or double-click a tab to bring the ribbon back for good. And if you want to perform the same trick without raising your fingers from the keyboard, you can use the shortcut key Ctrl+F1.*

Note: In some circumstances, you may see tabs that aren't listed here. Macro programmers and other highly technical types use the Developer tab. (You'll learn how to reveal this tab on page 817.) The Add-Ins tab appears when you're viewing workbooks that were created in previous versions of Excel and used custom toolbars. And finally, you can create a tab of your own if you're ambitious enough to customize the ribbon, as explained in the appendix.

The Formula Bar

The *formula bar* appears above the worksheet grid but below the ribbon (Figure 1-10). It displays the address of the active cell (like A1) on the left edge, and it also shows you the current cell's contents.

You can use the formula bar to enter and edit data, instead of editing directly in your worksheet. This approach is particularly useful when a cell contains a formula or a large amount of information. That's because the formula bar gives you more work room than a typical cell. Just as with in-cell edits, you press Enter to confirm your changes or Esc to cancel them. Or you can use the mouse: When you start typing in the formula bar, a checkmark and an "X" icon appear just to the left of the box where you're typing. Click the checkmark to confirm your entry or "X" to roll it back.

Note: You can hide (or show) the formula bar by choosing View→Show→Formula Bar. But the formula bar is such a basic part of Excel that you'd be unwise to get rid of it. Instead, keep it around until Chapter 8, when you'll learn how to build formulas.

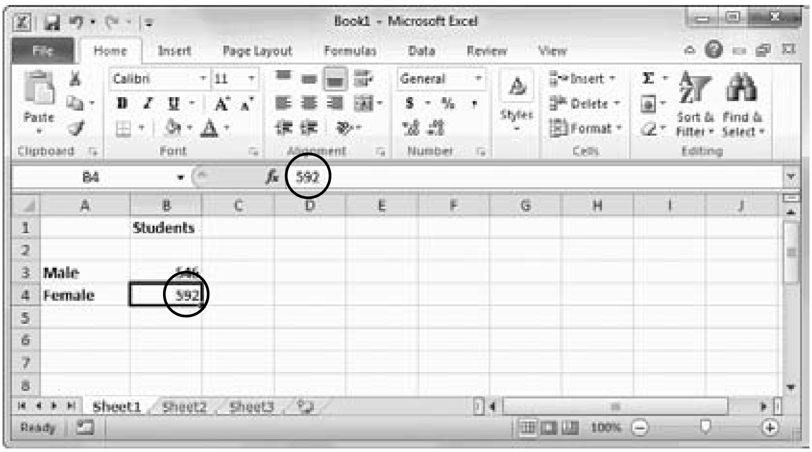

Figure 1-10: *The formula bar (just above the grid) shows information about the active cell. In this example, the formula bar shows that the current cell is B4 and that it contains the number 592. Instead of editing this value in the worksheet, you can click anywhere in the formula bar and make your changes there.*

Ordinarily, the formula bar is a single line. If you have a *really* long entry in a cell (like a paragraph's worth of text), you need to scroll from one side to the other. However, there's another option—you can resize the formula bar so it fits more information, as shown in Figure 1-11.

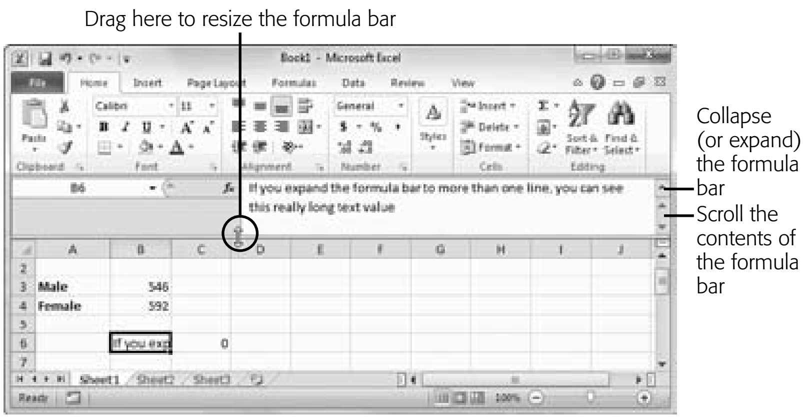

Figure 1-11: *To enlarge the formula bar, click the bottom edge and pull down. You can make it two, three, four, or many more lines large. Best of all, once you get the size you want, you can use the expand/collapse button on the right side of the formula bar to quickly expand it to your preferred size and collapse it back to the single-line view.*

The Status Bar

Though people often overlook it, the status bar (Figure 1-12) is a good way to keep on top of Excel's current state. For example, if you save or print a document, the status bar shows the progress of the printing process. If you're performing a quick action, the progress indicator may disappear before you have a chance to even notice it. But if you're performing a time-consuming operation—say, printing out an 87-page table of the airline silverware you happen to own—you can look to the status bar to see how things are coming along.

Figure 1-12: *In the status bar, you can see the basic status text (which just says "Ready" in this example), the view buttons (which are useful when you're preparing a spreadsheet for printing), and the zoom slider bar (which lets you enlarge or shrink the current worksheet view).*

The status bar combines several different types of information. The leftmost part of the status bar shows the Cell Mode, which displays one of three indicators:

- The word "Ready" means that Excel isn't doing anything much at the moment, other than waiting for you to take some action.

- The word "Enter" appears when you start typing a new value into a cell.
- The word "Edit" means the cell is currently in edit mode, and pressing the left and right arrow keys moves through the cell data, instead of moving from cell to cell. You can place a cell in edit mode or take it out of edit mode by pressing F2.

Just to the right of the Cell Mode information, you may see a small button that lets you start recording a macro. (A *macro* is a series of steps you can replay as often as you need to automate tiresome chores. You'll learn more about macros in Chapter 28.)

Farther to the right on the status bar are the view buttons, which let you switch to Page Layout View or Page Break Preview. These different views help you see what your worksheet will look like when you print it. They're covered in Chapter 7.

The zoom slider is next to the view buttons, at the far right edge of the status bar. You can slide it to the left to zoom out (which fits more information into your Excel window at once) or slide it to the right to zoom in (and take a closer look at fewer cells). You can learn more about zooming on page 196.

In addition, the status bar displays other miscellaneous indicators. For example, if you press the Scroll Lock key, a Scroll Lock indicator appears on the status bar (next to the "Ready" text). This indicator tells you that you're in *scroll mode*. In scroll mode, the arrow keys don't move you from one cell to another; instead, they scroll the entire worksheet up, down, or to the side. Scroll mode is a great way to check out another part of your spreadsheet without leaving your current position.

You can control what indicators appear in the status bar by configuring it. To see a full list of possibilities, right-click the status bar. A huge list of options appears, as shown in Figure 1-13. Table 1-2 describes the different status bar options.

Note: The Caps Lock indicator doesn't determine whether you can use the Caps Lock key—that feature always works. The Caps Lock indicator just lets you know when Caps Lock mode is on. That way you won't be surprised by an accidental keystroke that turns your next data entry INTO ALL CAPITALS.

Table 1-2. *Status bar indicators*

Indicator	Meaning
Cell Mode	Shows Ready, Edit, or Enter depending on the state of the current cell.
Signatures, Information Management Policy, and Permissions	Displays information about the rights and restrictions of the current spreadsheet. These features come into play only if you're using Office SharePoint Server to share spreadsheets among groups of people (usually in a corporate environment). SharePoint is introduced on page 746.
Caps Lock	Indicates whether Caps Lock mode is on. When Caps Lock is on, every letter you type is automatically capitalized. To turn Caps Lock mode on or off, hit Caps Lock.

Num Lock	Indicates whether Num Lock mode is on. When this mode is on, you can use the numeric keypad (typically at the right side of your keyboard) to type in numbers more quickly. When this sign's off, the numeric keypad controls cell navigation instead. To turn Num Lock on or off, press Num Lock.
Scroll Lock	Indicates whether Scroll Lock mode is on. When it's on, you can use the arrow keys to scroll through the worksheet without changing the active cell. (In other words, you can control your scrollbars by just using your keyboard.) This feature lets you look at all the information you have in your worksheet without losing track of the cell you're currently in. You can turn Scroll Lock mode on or off by pressing Scroll Lock.
Fixed Decimal	Indicates when Fixed Decimal mode is on. When this mode is on, Excel automatically adds a set number of decimal places to the values you enter in any cell. For example, if you set Excel to use two fixed decimal places and you type the number 5 into a cell, Excel actually enters 0.05. This seldom-used featured is handy for speed typists who need to enter reams of data in a fixed format. You can turn this feature on or off by selecting File®Options, choosing the Advanced section, and then looking under "Editing options" to find the "Automatically insert a decimal point" setting. Once you turn this checkbox on, you can choose the number of decimal places (the standard option is 2).
Overtype Mode	Indicates when Overwrite mode is turned on. Overwrite mode changes how cell edits work. When you edit a cell and Overwrite mode is on, the new characters that you type overwrite existing characters (rather than displacing them). You can turn Overwrite mode on or off by pressing Insert.
End Mode	Indicates that you've pressed End, which is the first key in many two-key combinations; the next key determines what happens. For example, hit End and then Home to move to the bottom-right cell in your worksheet. See Table 1-1 for a list of key combinations, some of which use End.
Macro Recording	Macros are automated routines that perform some task in an Excel spreadsheet. The Macro Recording indicator shows a record button (which looks like a red circle superimposed on a worksheet) that lets you start recording a new macro. You'll learn more about macros in Chapter 28.
Selection Mode	Indicates the current Selection mode. You have two options: normal mode and *extended selection.* When you press the arrows keys and extended selection is on, Excel automatically selects all the rows and columns you cross. Extended selection is a useful keyboard alternative to dragging your mouse to select swaths of the grid. To turn extended selection on or off, press F8. You'll learn more about selecting cells and moving them around in Chapter 3.
Page Number	Shows the current page and the total number of pages (as in "Page 1 of 4"). This indicator appears only in Page Layout view (as described on page 217).

Average, Count, Numerical Count, Minimum, Maximum, Sum	Show the result of a calculation on the selected cells. For example, the Sum indicator shows the total of all the numeric cells that are currently selected. You'll take a closer look at this handy trick on page 98.
Upload Status	Does nothing (that we know of). Excel does show a handy indicator in the status bar when you're uploading files to the Web, as you'll learn in Chapter 26. However, the upload status is always shown, and this setting doesn't seem to have any effect.
View Shortcuts	Shows the three view buttons that let you switch between Normal view, Page Layout View, and Page Break Preview.
Zoom	Shows the current zoom percentage (like 100 percent for a normal-sized spreadsheet, and 200 percent for a spreadsheet that's blown up to twice the magnification).
Zoom Slider	Shows a slider that lets you zoom in closer (by sliding it to the right) or out to see more information at once (by sliding it to the left).

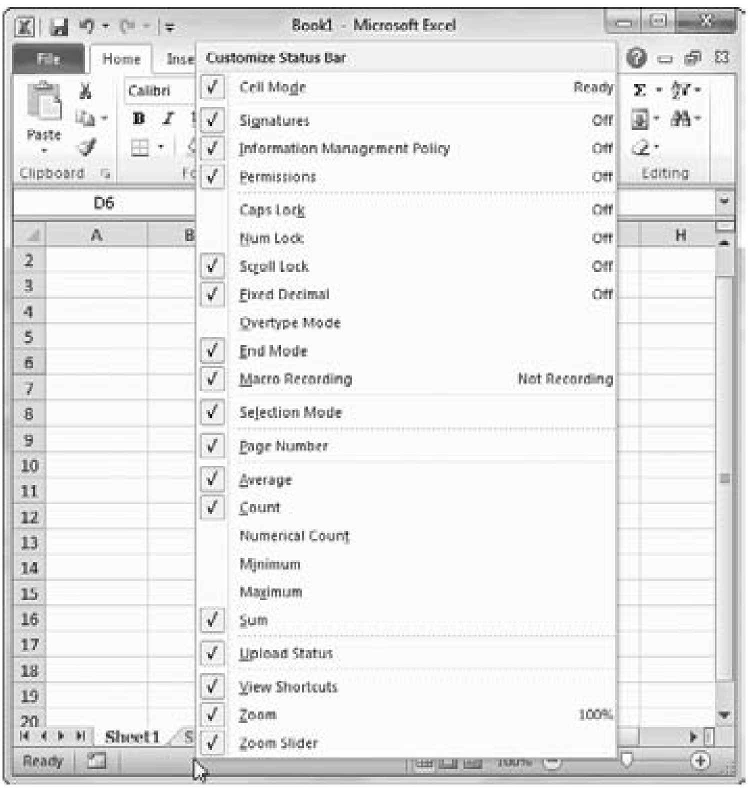

Figure 1-13:
Every item that has a checkmark appears in the status bar when you need it. For example, if you choose Caps Lock, the text "Caps Lock" appears in the status bar whenever you hit the Caps Lock key to switch to all-capital typing. The text that appears on the right side of the list tells you the current value of the indicator. In this example, Caps Lock mode is currently off and the Cell Mode text says "Ready."

Going Backstage

Your data is the star of the show. That's why the creators of Excel refer to your worksheet as being "on stage." The auditorium is the Excel main window, which—as you've just seen—includes the handy ribbon, formula bar, and status bar. Sure, it's a strange metaphor. But once you understand it, you'll realize the rationale for Excel's *backstage view*, which temporarily takes you away from your worksheet and lets you concentrate on other tasks that don't involve entering or editing data. These tasks include saving your spreadsheet, opening more spreadsheets, printing your work, and changing Excel settings.

To switch to backstage view, click the File button that's just to the left of the Home ribbon tab. Excel temporarily tucks your worksheet out of sight (although it's still open and waiting for you). This gives it space to show extra information related to the task you want to perform, as shown in Figure 1-14. For example, if you plan to print your spreadsheet, Excel's backstage view has room to show a preview of the printout. Or if you want to open an existing spreadsheet, Excel can show a detailed list of files you've recently worked on.

To get out of backstage view and return to your worksheet, just click the File button again, or press Esc.

Figure 1-14: *When you first switch to backstage view, Excel shows the Info page, which provides some basic information about your workbook file, its size, when it was last edited, who edited it, and so on (see the column on the far right). The Info page also provides the gateway to three important features: document protection (Chapter 24), compatibility checking (page 41), and AutoRecover backups (page 82). To go to another section, click a different command in the column on the far left.*

Excel Options

The Excel Options window provides a central hub where you can adjust how Excel looks, behaves, and calculates (see Figure 1-15). To get to this window, choose File→Options.

The various sections in the Excel Options window let you tweak a wide variety of different details. Some of these details are truly handy, like the options for opening and saving files (which are described at the end of this chapter). Others are seldom-used holdovers from the past, like the option that lets Excel act like Lotus—an ancient piece of spreadsheet software—when you hit the "/" key.

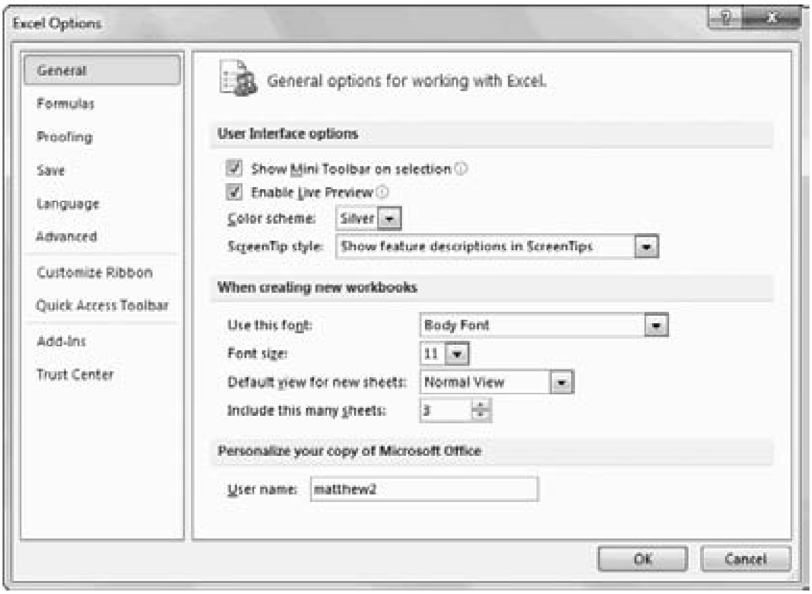

Figure 1-15: *The Excel Options window is divided into 10 sections. To pick which section to look at, choose an entry from the list on the left. In this example, you're looking at the General settings group. In each section, the settings are further subdivided into titled groups. You may need to scroll down to find the setting you want.*

Tip: Some important options have a small i-in-a-circle icon next to them, which stands for "information." Hover over this icon, and you see a tooltip that gives you a brief description about that setting.

While you're getting to know Excel, you can comfortably ignore most of what's in the Excel Options window. But you'll return here many times throughout this book to adjust settings and fine-tune the way Excel works.

Saving Files

As everyone who's been alive for at least three days knows, you should save your work early and often. Excel is no exception. You have three choices for saving a spreadsheet file:

- **Save**. This action updates the spreadsheet file with your most recent changes. If you use Save on a new file that hasn't been saved before Excel prompts you to choose a folder and file name. To use Save, select File→Save, or press Ctrl+S. Or look up at the top of the Excel window in the Quick Access toolbar for the tiny Save button, which looks like an old-style diskette.

Tip: Resaving a spreadsheet is an almost instantaneous operation, and you should get used to doing it regularly. After you've made any significant change, just hit Ctrl+S to make sure you've stored the latest version of your data.

- **Save As**. This choice allows you to save your spreadsheet file with a new name. You can use Save As the first time you save a new spreadsheet, or you can use it to save a copy of your current spreadsheet with a new name, in a new folder, or as a different file type. To use Save As, select File→Save As, or press F12.
- **Save & Send**. The File→Save & Send page in Excel's backstage view provides several of the same options you get from the Save As dialog box. The difference is that it makes them a bit more obvious and a bit more convenient. Figure 1-16 shows what to look for.

Figure 1-16:
Using File→Save & Send, you can email a copy of your workbook (page 723), transfer it to the Excel Web App (Chapter 26), or upload it to SharePoint (page 746). But ignore these options for now and focus on the "File Types" section underneath, which gives you shortcuts for saving your work in alternative file formats. If you click Change File Type, you get a list with the most common file formats (on the right). Double-click an entry to open the Save As dialog box with that selection. Or click Create PDF/XPS Document to create a print-ready PDF document, as described on page 44.

The Excel File Format

Since time immemorial, Excel fans have been saving their lovingly crafted spreadsheets in *.xls* files (as in *AirlineSilverware.xls*). But when Excel 2007 hit the streets, it introduced a completely new file format, with the extension *.xlsx* (as in *AirlineSilverware.xlsx*). Excel 2010 keeps the .xlsx revamped format, without introducing any more changes.

As an Excel user, you need to decide whether you want to save your spreadsheets in the latest and greatest *.xlsx* format or in the older *.xls* format. But before you make that decision, it helps to know a bit more about the advantages of the .xlsx format:

- **It's compact**. The .xlsx format uses ZIP file compression, so spreadsheet files are smaller—way smaller (as much as 75 percent smaller than their original sizes). And even though the average hard drive is already large enough to swallow thousands of old-fashioned Excel files, the new compact format is easier to email around.
- **It's less error-prone**. The .xlsx format carefully separates ordinary content, pictures, and macro code into separate sections. Microsoft claims that this change makes for tougher files. Now, if a part of your Excel file is damaged (for example, due to a faulty hard drive), there's a much better chance that you can still retrieve the rest of the information. (You'll learn about Excel disaster recovery on page 53.)
- **It's extensible**. The .xlsx format uses XML (the eXtensible Markup Language), which is a standardized way to store information. (You'll learn more about XML in Chapter 23.) XML storage doesn't benefit the average person, but it's sure to earn a lot of love from companies that plan to build custom software that uses Excel documents. As long as Excel documents are stored in XML, these companies can create automated programs that pull the information they need straight out of a spreadsheet, without going through Excel. These programs can also generate made-to-measure Excel documents all on their own.

For all these reasons, .xlsx is the format of choice for Excel 2010. However, Microsoft prefers to give people all the choices they could ever need (rather than make life really simple), and Excel file formats are no exception. In fact, the .xlsx file format actually has *two* additional flavors.

First, there's the closely related .xls*m* cousin, which adds the ability to store macro code. If you've added any macros to your spreadsheet, Excel prompts you to use this file type when you save your spreadsheet. (You'll learn about macros in Chapter 28.)

Second, there's the optimized .xls*b* format, which is a specialized option that just might be faster when you're opening and saving gargantuan spreadsheets. The .xlsb format has the same automatic compression and error-resistance as .xlsx, but it doesn't use XML. Instead, it stores information in raw *binary* form (good ol' ones and zeroes), which is speedier in some situations. To use the .xlsb format, choose File→Save As, and then, from the "Save as type" list, choose Excel Binary Workbook (.xlsb).

Most of the time, you don't need to think about Excel's file format. You can just create your spreadsheets, save them, and let Excel take care of the rest. The only time you need to stop and think twice is when you need to share your work with other, less fortunate people who have older versions of Excel, such as Excel 2003. You'll learn how to deal with this challenge in the following sections.

POWER USERS' CLINIC

Under the Hood with .xlsx Files

Here's a shocking secret: The new .xlsx file format is actually a ZIP file in disguise. It's composed of several files that are compressed and then packaged together as a single unit. With a little know-how, you can take a look at these hidden files-within-a-file, which makes for a great Excel party trick. Here's how:

1. Save your Excel spreadsheet in .xlsx format.
2. Browse to the file (using My Computer, Windows Explorer, or your favorite file management tool). If you're lazy, you can save the file on the desktop so you can manipulate it right there.
3. Right-click the file, and then choose Rename.
4. Change the file extension to .zip. So if you start with BlackMarketDinnerware.xlsx, change it to BlackMarketDinnerware.zip.
5. Now, open the ZIP file by double-clicking the file.
6. You can now see the files that are hidden inside your Excel file. They're organized into several folders (Figure 1-17). To find the actual content from your spreadsheet, head to xl→worksheets→sheet1.xml. Double-click it to open it up and take a look at what's inside.
7. When you're finished, rename the file using its .xlsx extension so you can open it in Excel.

To learn way more about the technical details of this type of file storage, you can read the Microsoft white paper at *http://msdn.microsoft.com/en-us/library/aa338205.aspx.*

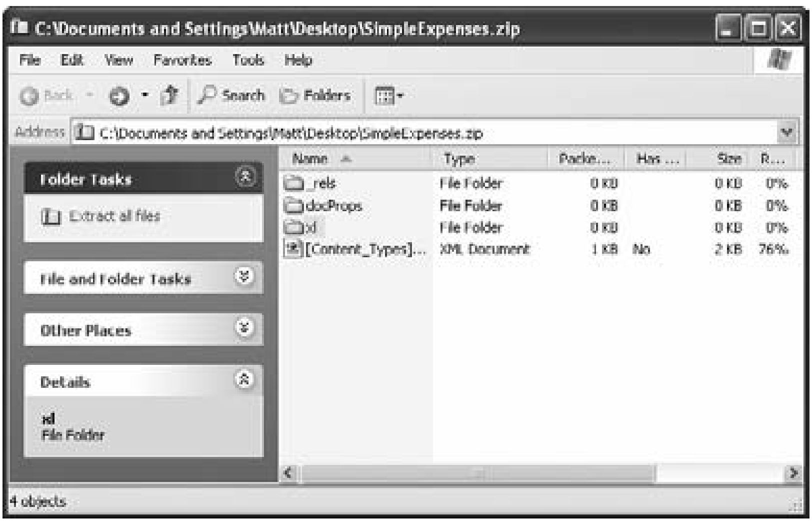

Figure 1-17: *Inside every .xlsx file lurks a number of compressed files, each with different information. For example, separate files store printer settings, styles, the name of the person who created the document, the composition of your workbook, and each individual worksheet.*

Tip: Don't use the .xlsb format unless you've tried it out and find it really does give better performance for one of your spreadsheets. Usually, .xlsx and .xlsb are just as fast. And remember, the only time you'll see any improvement is when you're loading or saving a file. Once your spreadsheet is open in Excel, everything else (like scrolling around and performing calculations) happens at the same speed.

Sharing Your Spreadsheet with Excel 2007

As you've just learned, Excel 2007 uses the same .xlsx file format as Excel 2010. That means that an Excel 2010 fan can exchange files with an Excel 2007 devotee, and there won't be any technical problems.

However, there are still a few issues that can trip you up when sharing spreadsheets between Excel 2010 and Excel 2007. For example, Excel 2010 introduces a few new formula functions, such as RANK.AVG (page 273). If you write a calculation that uses this function in Excel 2010, it won't work when someone else opens the spreadsheet in Excel 2007. Instead of seeing the numeric result that you saw, they'll see an error code mixed in with the rest of the spreadsheet data.

To avoid this sort of problem, you need the help of an Excel tool called the Compatibility Checker. The Compatibility Checker scans your spreadsheet to find features and formulas that will cause a problem in Excel 2007.

To use the Compatibility Checker, follow these steps:

1. **Choose File→Info.**

   Excel switches into backstage view.

2. **Click the Check for Issues button, and choose Check Compatibility.**

   The Compatibility Checker scans your spreadsheet, looking for signs of trouble. It then reports any problems back to you (Figure 1-18).

3. **Optionally, you can hide problems that don't affect Excel 2007.**

   The Compatibility Checker reports on two types of problems. The first type of problems affect anyone who doesn't have Excel 2010. These problems appear with the text "Excel 97-2003" and "Excel 2007" in the column on the right. The second problems affect Excel 2003 users but not Excel 2007 users. These problems appear with the text "Excel 97-2003" in the right column.

   If you plan to share your spreadsheet file with Excel 2007 only, there's no need to worry about Excel 2003 compatibility. You can hide these messages from the list by clicking the "Select versions to show" button and turning off Excel 97-2003.

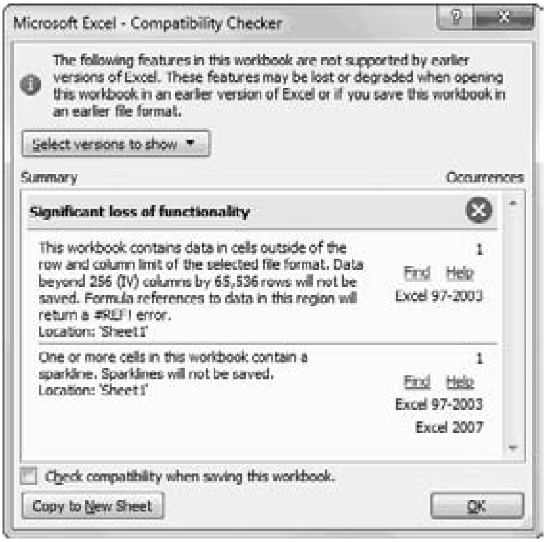

Figure 1-18:
In this example, the Compatibility Checker has found two potential problems. One will affect only Excel 2003 users, while the other will affect both Excel 2007 and Excel 2003 users.

4. **Review the problems.**

   You can choose to ignore the Compatibility Checker issues, click Find to hunt each one down, or click Help to figure out the exact problem. You can also click Copy to New Sheet to insert a full compatibility report into your spreadsheet as a separate worksheet. This way, you can print it up and review it in the comfort of your cubicle. (To get back to the worksheet with your data, click the Sheet1 tab at the bottom of the window. Chapter 4 has more about how to use and manage multiple worksheets.)

Note: The problems that the Compatibility Checker finds won't cause serious errors, like crashing your computer or corrupting your data. That's because Excel is designed to *degrade gracefully*. That means you can still open a spreadsheet that uses newer, unsupported features in an old version of Excel. However, you may receive a warning message and part of the spreadsheet may seem broken—that is, it won't work as you intended.

5. **Optionally, you can set the Compatibility Checker to run automatically for this workbook.**

   Turn on the "Check compatibility when saving this workbook" checkbox. Now, the Compatibility Checker will run each time you save your spreadsheet, just before the file is updated.

Saving Your Spreadsheet for Excel 2003

If you want to share your workbook with people who are using Excel 2003, the process is a bit more involved, because Excel 2003 uses the older .xls format instead of .xlsx.

When you find yourself in this situation, you have two choices:

- **Save your spreadsheet in the old format**. You can save a copy of your spreadsheet in the traditional .xls Excel standard that's been supported since Excel 97. To do so, choose File→Save As, and then, from the "Save as type" list, choose Excel 97-2003 Workbook.
- **Use a free add-in for older versions of Excel**. People who are stuck with Excel 2000, Excel 2002, or Excel 2003 *can* read your Excel 2010 files—they just need a free add-in that's provided by Microsoft. This is a good solution because it's doesn't require any work on your part. People with past-its-prime versions of Excel can find the add-in they need by surfing to *www.microsoft.com/downloads* and searching for "compatibility pack file formats" (or use the secret shortcut URL *http://tinyurl.com/y5w78r*). However, you should still run the Compatibility Checker to find out if your spreadsheet uses features that aren't supported in Excel 2003.

Tip: When you save your Excel spreadsheet in another format, make sure to keep a copy in the standard .xlsx format. Why bother? Because other formats aren't guaranteed to retain all your information, particularly if you choose a format that doesn't support some of Excel's newer features.

As you already know, each version of Excel introduces a small set of new features. Older versions of Excel don't support these features. The differences between Excel 2007 and Excel 2010 are small. But the differences between Excel 2003 and Excel 2010 are more significant.

Excel tries to help you out in two ways. First, whenever you save a file in .xls format, Excel automatically runs the Compatibility Checker to review your spreadsheet and detect compatibility issues. Second, whenever you open a spreadsheet that's in the old .xls file format, Excel switches into *compatibility mode*. While the Compatibility Checker points out potential problems after the fact, compatibility mode is designed to prevent you from using unsupported features in the first place. For example, in compatibility mode you'll face these restrictions:

- Excel limits you to a smaller grid of cells (65,536 rows instead of 1,048,576).
- Excel prevents you from using really long or deeply nested formulas.
- Excel doesn't let you use some pivot table features.

In compatibility mode, these missing features aren't anywhere to be found. In fact, compatibility mode is so seamless that you might not even notice you're being limited. The only clear indication is the title bar at the top of the Excel window. Instead of seeing something like CateringList.xlsx, you'll see CateringList.xls [Compatibility Mode].

Note: When you save an Excel workbook in .xls format, Excel won't switch into compatibility mode right away. Instead, you need to close the workbook and reopen it.

If you decide at some point that you're ready to move into the modern world and convert your file to the .xlsx format favored by Excel 2010, you can use the trusty File→Save As command. However, there's an even quicker shortcut. Just choose File→Info, and click the Convert button. This saves an Excel 2010 version of your file with the same name but with the extension .xlsx, and reloads the file so you get out of compatibility mode. It's up to you to delete your old .xls original if you don't need it anymore.

Saving a Spreadsheet in Other Formats

Some eccentric individuals have even older or stranger spreadsheet software on their computers. If you want to save a copy of your spreadsheet in a more exotic file type, you can choose File→Save As, and then find the desired format in the "Save as type" drop-down list (Figure 1-19). Alternatively, you can get quick access to the most popular file formats by choosing File→Save & Send and then clicking Change File Type. Double-click one of the file types that's shown here to open the Save As dialog box with your choice already filled in the "Save as type" box.

Excel lets you save your spreadsheet using a variety of different formats, including the classic Excel 95 format from more than a decade ago. If you're looking to view your spreadsheet using a mystery program, use the CSV file type, which produces a comma-delimited text file that almost all spreadsheet applications on any operating system can read (comma-delimited means the information has commas separating each cell).

Saving Your Spreadsheet As a PDF

Sometimes you want to save a copy of your spreadsheet so that people can read it even if they don't have Excel (and even if they're running a different operating system, like Linux or Apple's OS X). In this situation, you have several choices:

- **Use the Excel Viewer**. Even if you don't have Excel, you can install a separate tool called the Excel Viewer, which is available from Microsoft's website (search for "Excel Viewer" at *http://www.microsoft.com/downloads*). However, few people have the viewer, and even though it's free, few want to bother installing it. And it doesn't work on non-Windows computers.
- **Save your workbook as an HTML web page**. That way, all you need to view the workbook is a web browser (and who doesn't have one of those?). The only disadvantage is that you could lose complex formatting. Some worksheets may make the transition to HTML gracefully, while others don't look very good when they're squashed into a browser window. And if you're planning to let other people print the exported worksheet, the results might be unsatisfactory. The next section has more about saving your worksheet as a web page.

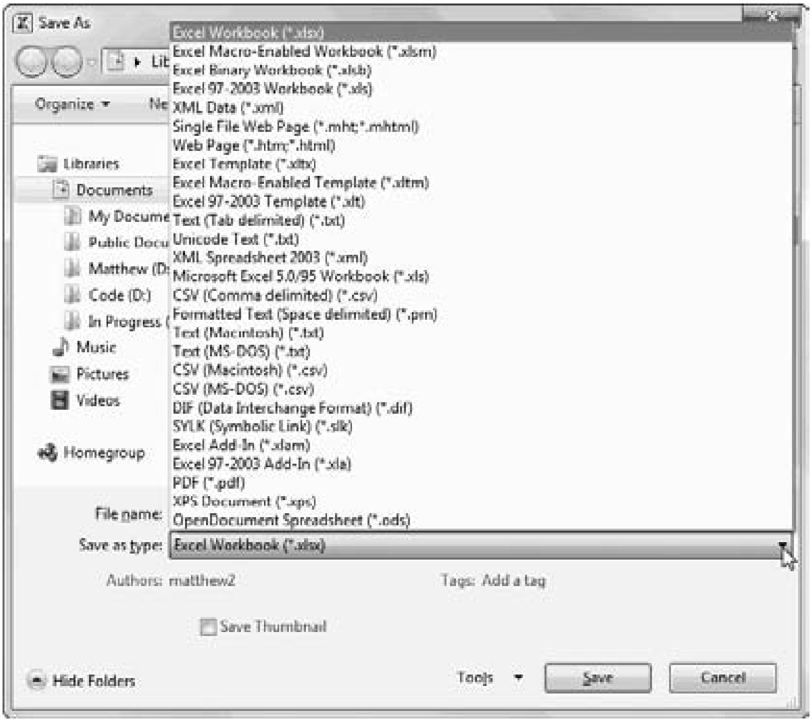

Figure 1-19: *Excel offers a few useful file type options in the "Save as type" list. CSV format is the best choice for compatibility with truly old software (or when nothing else seems to work). The four most commonly used formats–regular workbooks, macro-enabled workbooks, binary workbooks, and old-style Excel 2003 workbooks–sit at the top of the list.*

- **Save your workbook as a PDF file**. This gets you the best of both worlds—you keep all the rich formatting (so your workbook can be printed), and you let people who don't have Excel (and possibly don't even have Windows) view your work. Excel 2007 introduced the Save As PDF feature, but it forced Excel fans to download an add-in to get it. Excel 2010 has no such limitation. To save your spreadsheet as a PDF, you simply select File→Save As, and then pick PDF from the "Save as type" list. Or choose File→Save & Send, click Create PDF/XPS Document (in the "File Types" section), and then click the Create PDF/XPS button.

It doesn't matter whether you use the File→Save As or the File→Save & Send approach. Either way, you end up at a modified version of the Save As dialog box that has a few additional options (Figure 1-20).

The "Publish as PDF" dialog box gives you some control over the quality settings with the "Optimize for" options. If you're just saving a PDF copy so other people can *view* the information in your workbook, choose "Minimum size (publishing online)" to save some space. On the other hand, if there's a possibility that the people reading your PDF might want to print it out, choose "Standard (publishing online and printing)" to save a slightly larger PDF that makes for a better printout.

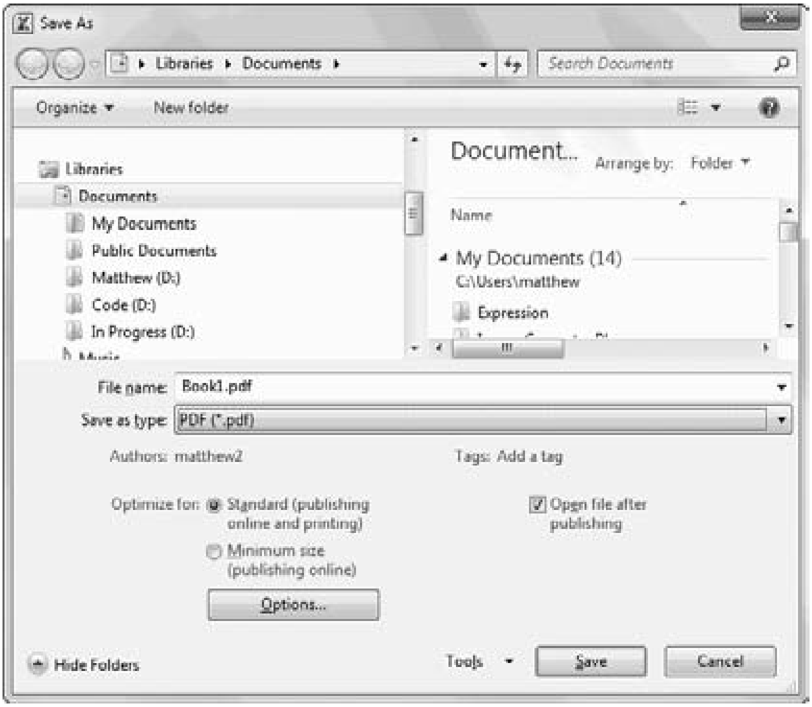

Figure 1-20: *PDF files can be saved with different resolution and quality settings (which mostly affect any graphical objects like pictures and charts that you've placed in your workbook). Normally, you use higher quality settings if you're planning to print your PDF file, because printers use higher resolutions than computer monitors.*

UP TO SPEED

Learning to Love PDFs

You've probably heard about PDFs, Adobe's popular format for sharing formatted, print-ready documents. People use PDFs to pass around product manuals, brochures, and all sorts of electronic documents. Unlike a document format like .xlsx, PDF files are designed to be viewed and printed, but not edited.

The best part about PDFs is that you can view them on just about any type of computer and operating system using the free Adobe Reader. You can download Adobe Reader at *http://get.adobe.com/reader*, but you probably don't need to. Most computers already have Adobe Reader installed, because it comes bundled with so many different programs (usually so you can view their electronic documentation). It's also widely used on the Web.

Incidentally, PDF isn't the only kid on the block. Microsoft's newer operating systems, Windows Vista and Windows 7, include another electronic paper format called XPS (XML Paper Specification). In time, as XPS is integrated into more and more products, it might become a true PDF competitor. But for now, PDF is dramatically more popular and widespread, so it's the one to stick with. (If you're interested in saving an Excel document as a XPS file, you can do that too—just choose XPS from the "Save as type" list.)

You can switch on the "Open file after publishing" setting to tell Excel to open the PDF file in Adobe Reader (assuming you have it installed) after the publishing process is complete, so you can check the result.

Finally, if you want to publish only a portion of your spreadsheet as a PDF file, click the Options button to open a dialog box with even more settings. You can choose to publish just a fixed number of pages, just the selected cells, and so on. These options mirror the choices you get when sending a spreadsheet to the printer (page 211). You also see a few more cryptic options, most of which you can safely ignore. (They're intended for PDF nerds.) One exception is the "Document properties" option—turn this off if you don't want the PDF to keep track of certain information that identifies you, like your name. (Excel document properties are discussed in more detail on page 719.)

Saving Your Spreadsheet As an HTML File

HTML (short for *Hypertext Markup Language*) is the language of the Web. Web authors use it to craft pages with text, links, and graphics.

In the early days of the Web, most programs had export-to-HTML features that weren't worth a second glance. They distorted formatting, mangled text, and generated HTML so ugly that professional web developers fainted at the sight of it. Fortunately, the situation has improved. Though Excel's HTML exporting might never match the graphical flair of the most talented web artists, it's still downright impressive. Best of all, you can have your data ready for web surfers in a matter of minutes. Here's what to do:

1. **Choose File→Save As.**

   This action opens the Save As dialog box.

2. **From the "Save as type" list, choose Web Page.**

   When you do, the Save As dialog box changes a little bit, as shown in Figure 1-21.

3. **Choose which portion of your workbook you want to export to HTML.**

   If you want to export every worksheet, select Entire Workbook. If you just want to export the current worksheet, then select Selection: Sheet. (Every Excel workbook begins with three worksheets, and so far you've only used the first one. You'll learn how to use the others—and add more—in Chapter 4.)

Tip: If you export the entire workbook, then Excel creates a web page that includes worksheet tab buttons, which you can use to switch from one worksheet to another. Generally, it's simpler to just export a single worksheet–that makes more sense to web surfers.

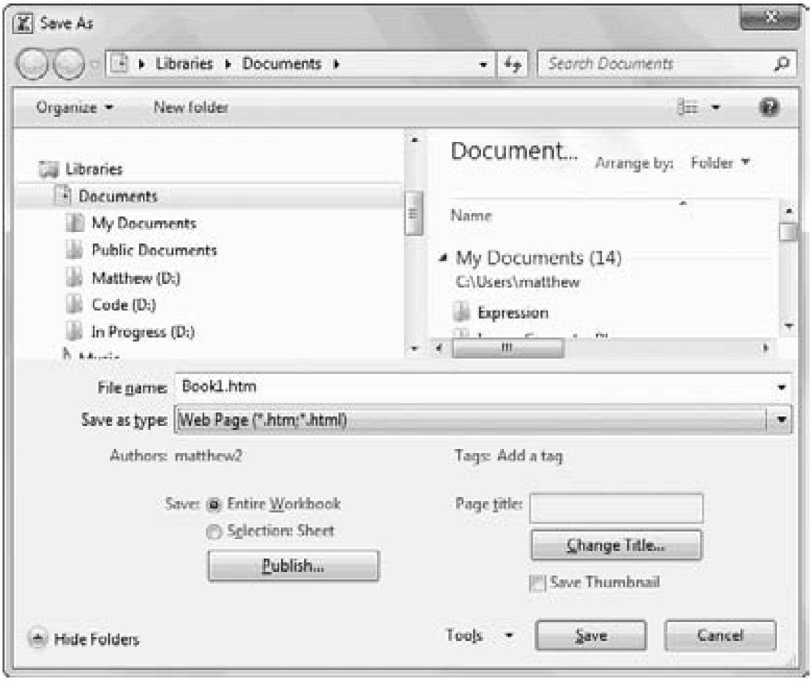

Figure 1-21: *When you're saving a spreadsheet as a web page, the Save As dialog box gets a couple of new buttons. Use Change Title to set the web page's title. You can also choose to save either the entire workbook or just the current selection.*

If you select a group of cells before you choose File→Save, then you won't see the Selection: Sheet option. Instead, you'll see an option for the range of cells you chose, like "Selection: $A$2:$B$5." (This means "the box of cells stretching from cell A2 in the top-left to cell B5 in the bottom-right.") Choose this to create a web page that includes just that section of content.

4. **If you want to add a title, click the Change Title button. When the Set Page Title dialog box appears, type in a title for your web page, and then click OK.**

   If you add a descriptive title, it appears in large bold font centered over the rest of your content. Titles don't have any restrictions, so feel free to use something clear and descriptive like "Blue Skies Budget Report" or "Bankruptcy Projections for 2011." This title also appears in the title bar of the web browser window. (Without the title information, most web browsers simply show the web page file name in their title bar.)

5. **Browse to the location where you want to save the web page.**

   This location can be any of the places where you normally save files.

6. **In the "File name" box, enter the name of the HTML file you want to create.**

   Depending on the content in your worksheet, Excel may create more than one file. If your worksheet contains embedded graphics or charts, or if you're printing the entire workbook, Excel creates additional files. Excel puts these files in a newly created folder that has the same name as your file, plus the text "_files."

   If you save the web page BudgetReport.htm, Excel creates a folder named BudgetReport_files to hold the extra files. You need to keep the HTML file and this folder together at all times, because the folder contains some information that the HTML file uses. (The box on page 51 has more about these folders.)

7. **You now have a choice to either save or publish your web page.**

   If you want to perform a direct save of your file, effectively converting your current workbook into the HTML format, click the Save button. The original copy of your workbook remains in an .xlsx file, but Excel won't update it again unless you choose File Menu→Save As and explicitly select it.

   If you want to publish your file, which creates a *copy* of your data in the HTML format, click Publish. This launches the "Publish as Web Page" dialog box, which gives you a last-minute chance to select the portion of the workbook you want to publish and change the file name or web page title. You've already set all the options you need, so just click Publish to save your HTML file. Your workbook remains in the .xlsx format, but Excel makes an HTML copy suitable for viewing in your browser.

Note: When you use the "Publish as Web Page" dialog box, you can select "AutoRepublish every time this workbook is saved" to tell Excel to save the HTML copy of your workbook every time you save the .xlsx workbook file.

The exported copy of your worksheet is amazingly faithful. Excel preserves the formatting, layout, and content of your original worksheet. If your worksheet contains pictures or charts, Excel saves a separate graphic file for each object and displays it in the same web page using the linking power of HTML. Figure 1-22 shows an exported worksheet that includes a chart.

Note: It goes without saying that while other people can view an Excel-created web page, they can't edit it. However, Excel 2010 blows the lid off this limitation with the Excel Web App—a simplified version of Excel that runs right in the browser. This tool is entirely free to use, even for people who don't have the desktop version of Excel. Chapter 26 covers the Excel Web App in detail.

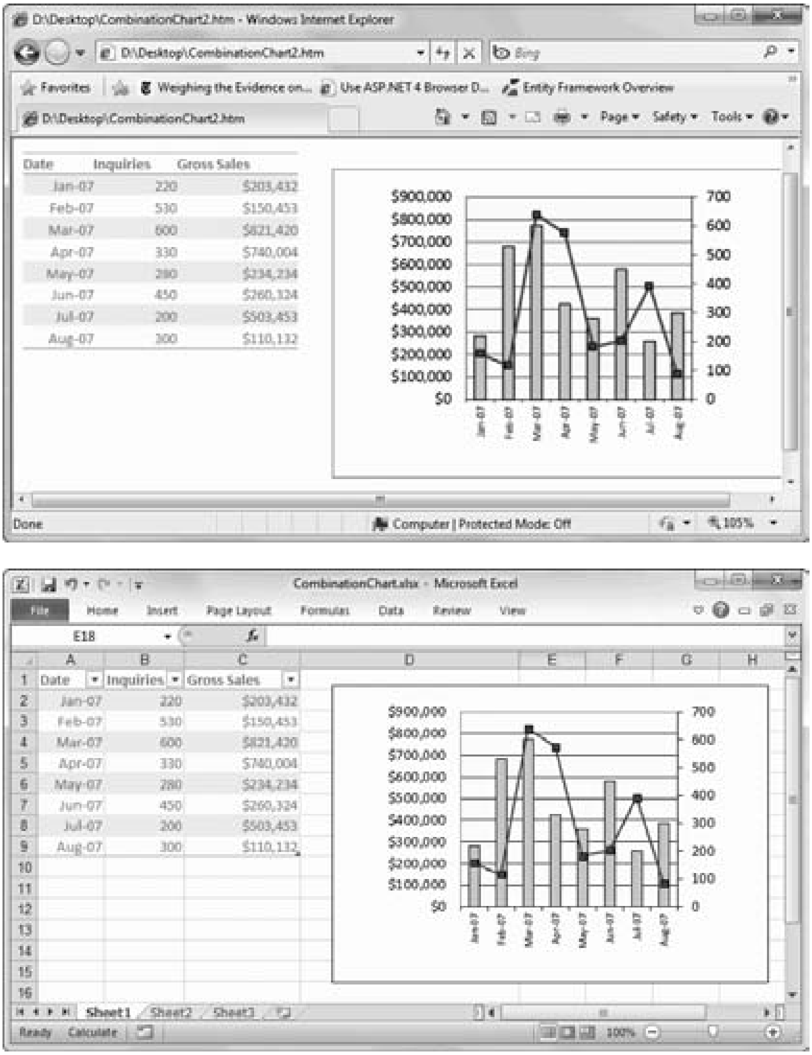

Figure 1-22: *The HTML version of a spreadsheet, as displayed in Internet Explorer (top), mimics its appearance in Excel (bottom) with surprising accuracy. You need to get used to a few minor changes–like the fact that the HTML version doesn't show any gridlines and can't display nonstandard fonts. The solution? Stick to commonly supported Web fonts, like Arial, Courier New, Times New Roman, and Verdana (to name the most popular).*

Saving Your Spreadsheet with a Password

Occasionally, you might want to add confidential information to a spreadsheet—for example, a list of the airlines from which you've stolen spoons. If your computer is on a network, the solution may be as simple as storing your file in the correct, protected location. But if you're afraid that you might inadvertently email the spreadsheet to the wrong people (say, executives at American Airlines), or if you're about to expose systematic accounting irregularities in your company's year-end statements, you'll be happy to know that Excel provides a tighter degree of security. It allows you to *password-protect* your spreadsheets, which means anyone who wants to open them has to know the password you've set.

Excel actually has two layers of password protection that you can apply to a spreadsheet:

- You can prevent others from *opening* your spreadsheet unless they know the correct password. This level of security, which scrambles your data for anyone without the password (a process known as *encryption*), is the strongest.
- You can let others read a spreadsheet, but you can prevent them from *modifying* it unless they know the correct password.

You can apply one or both of these restrictions to a spreadsheet. Applying them is easy. Just follow these steps:

1. **Choose File→Save As.**

   The Save As dialog box appears.

GEM IN THE ROUGH

A Convenient Web Page Package

When you save your Excel data in an HTML page, Excel usually creates several files–one each for pictures and charts. Web browsers don't know anything about Excel charts, so Excel turns those into .gif files that any browser can display. Excel also uses additional files if you save an entire workbook, in which case it places each worksheet in a separate file.

In order to keep everything organized, Excel stores all these extra files in a linked folder. If you save your Excel file as the web page BudgetReport.htm, Excel creates a folder called BudgetReport_files to hold extra files. Usually, you don't need to worry about which files are in this folder. Instead, you just load the main page (the one called BudgetReport.htm) in your browser.

You do, however, need to make sure the folder is always available, because the main page links to the files it contains. That means if you want to move your HTML page to another location or another computer, then you need to make sure that you copy the linked folder with it.

Keeping track of all these files can be a bit of a headache. What if you want to email your Excel web page to another person? You'll need to not only email all the files, but also explain how to create the required directory to hold all these files.

Fortunately, you can use an easier solution. The trick is to save everything using a special web archive format called .mht. When you save a .mht file, everything is combined into one compound file, including graphics. You can't easily edit .mht files, and they're supported only in Internet Explorer (other browsers, like Firefox or Opera, need not apply). However, .mht files are a great way to package everything up into one file when you need to send it in an email.

To save a .mht file, select File→Save As. In the "Save as type" list, choose Single File Web Page. Choose any other options you want, and click Save or Publish to finish the job.

2. **Click the Tools button, and then, from the pop-up menu, choose General Options.**

   If you're using a Windows XP computer, you'll find the Tools button in the bottom-left corner of the Save As dialog box. But if you're running Windows Vista or Windows 7, it's at the bottom right, just next to the Save button.

   The General Options dialog box appears.

3. **Type a password next to the security level you want to turn on (as shown in Figure 1-23). Then click OK.**

   The General Options dialog box also gives you a couple of other unrelated options:

   - Turn on the "Always create backup" checkbox if you want an extra copy of your file, just in case something goes wrong. (Think of it as insurance.) Excel creates a backup that has the file extension *.xlk*. For example, if you're saving a workbook named SimpleExpenses.xlsx and you use the "Always create backup" option, Excel creates a file named "Backup of *SimpleExpenses.xlk*" every time you save your spreadsheet. You can open the .xlk file in Excel just like an ordinary Excel file. When you do, you see that it has an exact copy of your work.
   - Turn on the "Read-only recommended" checkbox to prevent other people from accidentally making changes to your spreadsheet. When you use this option, Excel shows a message every time you (or anyone else) opens the file. This message politely suggests that you open the spreadsheet in *read-only mode*, in which case Excel won't allow any changes. Of course, it's entirely up to the person opening the file whether to accept this recommendation.

Figure 1-23:
You can use any sequence of letters and numbers as a password. Passwords are case-sensitive (which means that PanAm is different from panam), and they are masked (which means that all that appears in the window as you type is a series of asterisks).

4. **Click Save to store the file.**

   If you use a password to restrict people from opening the spreadsheet, Excel prompts you to supply the "password to open" the next time you open the file (Figure 1-24, top).

If you use a password to restrict people from modifying the spreadsheet, the next time you open this file you'll be given the choice—shown in Figure 1-24, bottom—to open it in *read-only mode* (which requires no password) or to open it in full edit mode (in which case you'll need to supply the "password to modify").

Figure 1-24:
Top: You can give a spreadsheet two layers of protection. Assign a "password to open," and you'll see this window when you open the file

Bottom: If you assign a "password to modify," you'll see the choices in this window. If you use both passwords, you'll see both windows, one after the other.

Disaster Recovery

The corollary to the edict "Save your data early and often" is the truism "Sometimes things fall apart quickly…before you've even had a chance to back up." Fortunately, Excel includes an invaluable safety net called AutoRecover.

AutoRecover periodically saves backup copies of your spreadsheet while you work. If you suffer a system crash, you can retrieve the last AutoRecover backup even if you never managed to save the file yourself. Of course, even the AutoRecover backup won't necessarily have *all* the information you entered in your spreadsheet before the problem occurred. But if AutoRecover saves a backup every 10 minutes (the standard), at most you'll lose 10 minutes of work.

If your computer does crash, when you get it running again, you can easily retrieve your last AutoRecover backup. In fact, the next time you launch Excel, it automatically checks the backup folder and, if it finds a backup, it opens a Document Recovery panel on the left of the Excel window.

If your computer crashes mid-edit, the next time you open Excel you'll probably see the same file listed twice in the Document Recovery window, as shown in Figure 1-25. The difference is the status. The status [AutoSaved] indicates the most recent backup created by Excel. The status [Original] indicates the last version of the file that *you* saved (which is safely stored on your hard drive, right where you expect it).

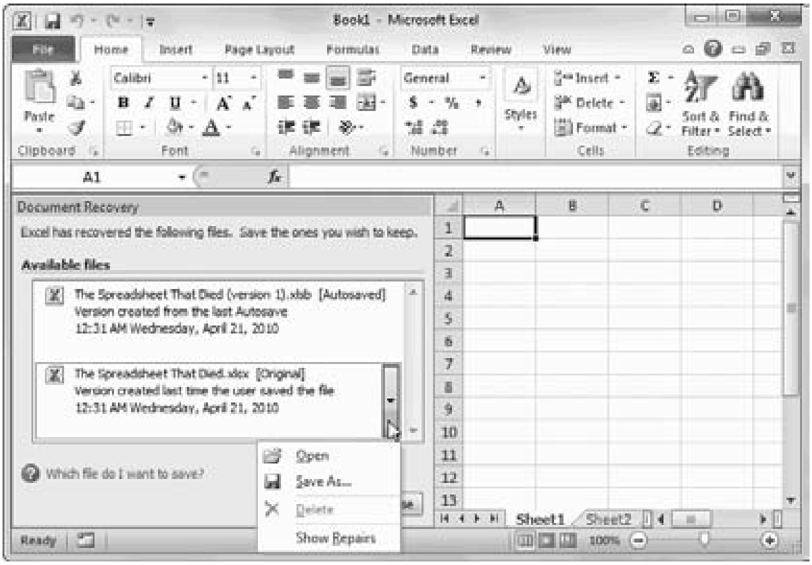

Figure 1-25: *You can save or open an AutoRecover backup just as you would an ordinary Excel file; simply click the item in the list. Once you've dealt with all the backup files, close the Document Recovery window by clicking the Close button. If you haven't saved your backup, Excel asks you at this point whether you want to save it permanently or delete the backup.*

To open a file that's in the Document Recovery window, just click it. You can also use a drop-down menu with additional options (Figure 1-24). Make sure to save the file before you leave Excel. After all, it's just a temporary backup.

If you attempt to open a backup file that's somehow been scrambled (technically known as *corrupted*), Excel automatically attempts to repair it. You can choose Show Repairs to display a list of any changes Excel had to make to recover the file.

AutoRecover settings

AutoRecover comes switched on when you install Excel, but you can tweak its settings. Choose File→Options, and then choose the Save section. Under the "Save workbooks" section, make sure that "Save AutoRecover information" is turned on.

You can also make a few other changes to AutoRecover settings:

- You can adjust the backup frequency in minutes. (See Figure 1-26 for tips on timing.)
- You can control whether Excel keeps a backup if you create a new spreadsheet, work on it for at least 10 minutes, and then close it without saving your work. This sort of AutoRecover backup is called a *draft*, and it's discussed in more detail on page 83. Ordinarily, the setting "Keep the last Auto Recovered file if I exit without saving" is switched on, and Excel keeps drafts. (To find all the drafts that Excel has saved for you, choose File→Recent and click the Recover Unsaved Workbooks link at the bottom of the window.)

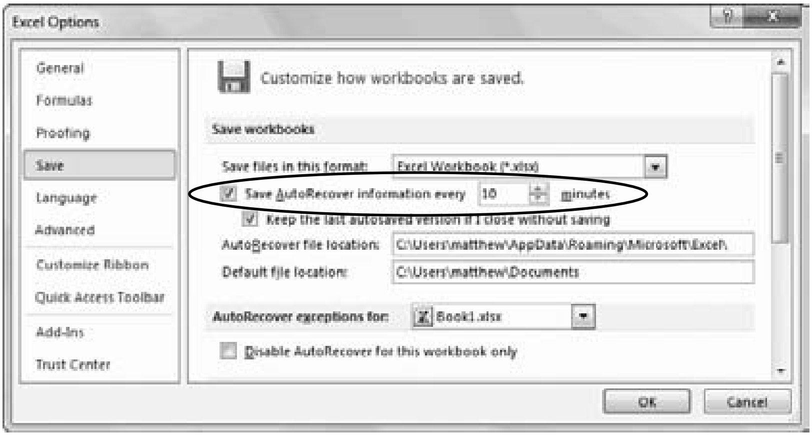

Figure 1-26: *You can configure how often AutoRecover saves backups. There's really no danger in being too frequent. Unless you work with extremely complex or large spreadsheets—which might suck up a lot of computing power and take a long time to save—you can set Excel to save the document every five minutes with no appreciable slowdown.*

- You can choose the folder where you'd like Excel to save backup files. The standard folder works fine for most people, but feel free to pick some other place. Unfortunately, there's no handy Browse button to help you find the folder, so you need to find the folder you want in advance (using a tool like Windows Explorer), write it down somewhere, and then copy the full folder path into this dialog box.
- Under the "AutoRecover exceptions" heading, you can tell Excel not to bother saving a backup of a specific spreadsheet. Pick the spreadsheet name from the list (which shows all the currently open spreadsheet files), and then turn on the "Disable AutoRecover for this workbook only" setting. This setting is exceedingly uncommon, but you might use it if you have a gargantuan spreadsheet full of data that doesn't need to be backed up. For example, this spreadsheet might hold records that you've pulled out of a central database so you can take a closer look. In this case, there's no need to create a backup because your spreadsheet just has a copy of the data that's in the database. (If you're interested in learning more about this scenario, check out Chapter 23.)

Opening Files

Opening existing files in Excel works much the same as it does in any Windows program. To get to the standard Open dialog box, choose File→Open (or use the keyboard shortcut Ctrl+O). Using the Open dialog box, you can browse to find the spreadsheet file you want, and then click Open to load it into Excel.

Excel can open many file types other than its native .xlsx format. To learn the other formats it supports, launch the Open dialog box and, at the bottom, open the "Files of type" menu, which shows you the whole list. If you want to open a file but you don't know what format it's in, try using the first option on the menu, "All Files."

Once you choose a file, Excel scans the beginning of the file and informs you about the type of conversion it will attempt to perform (based on what type of file Excel thinks it is).

Note: Depending on your computer settings, Windows might hide file extensions. That means that instead of seeing the Excel spreadsheet file MyCoalMiningFortune.xlsx, you'll just see the name MyCoalMiningFortune (without the .xlsx part on the end). In this case, you can still tell what the file type is by looking at the icon. If you see a small Excel icon next to the file name, that means Windows recognizes that the file is an Excel spreadsheet. If you see something else (like a tiny paint palette, for example), you need to make a logical guess about what type of file it is.

The Open dialog box also lets you open several spreadsheets in one step, as long as they're all in the same folder. To use this trick, hold down the Ctrl key and click to select each file. Whe you click Open, Excel puts each one in a separate window, just as if you'd opened them one after the other.

Opening Recent Documents

Plan to take another crack at a recent spreadsheet? You can find the most recently opened documents in Excel's Recent Documents list. To see this list, just choose File→Recent. The list holds 20 files. The best part about the Recent Documents list is the way you can *pin* a document there so it stays forever, as shown in Figure 1-27.

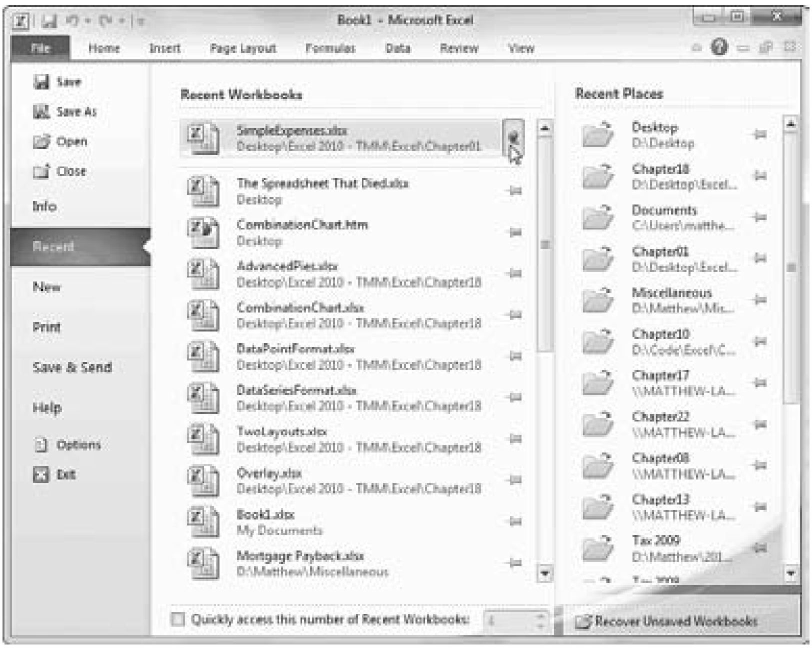

Figure 1-27: *To keep a spreadsheet around on the Recent Documents list, click the thumbtack on the right. Excel moves your workbook to the top of the list and pins it in place. That means it won't ever leave the list, no matter how many documents you open. If you decide to stop working with it later on, just click the thumbtack again to release it. Pinning is a great trick for keeping your most important files at your fingertips.*

You can put recent documents even closer at hand by switching on the "Quickly access this number of Recent Workbooks" checkbox, which you'll find underneath the list of recent documents. You can then enter a number in the adjacent box. For example, if you use 4 (the standard value), the four most recently used workbooks will appear directly in the column of commands on the left side of Excel's backstage view. That means that if one of them is named MyFavoriteFile.xlsx, you can open it by choosing File→MyFavoriteFile.xlsx, rather than choosing File→Recent, and then clicking MyFavoriteFile.xlsx in the recent files list.

Tip: Do you want to hide your recent editing work? You can remove any file from the recent document list by right-clicking it and choosing "Remove from list." And if the clutter is keeping you from finding the workbooks you want, just pin the important files, then right-click any file and choose "Clear unpinned workbooks." This action removes every file that isn't pinned down.

Just to the right of the Recent Documents list is the Recent Places list. It keeps track of the folders where you store your Excel files. So if you like to place Excel spreadsheets on your desktop and in your documents folder, you'll see both folders in the list (as in in Figure 1-27). When you click a location, Excel heads to the appropriate location and shows the standard Open dialog box.

If you store your work in several different locations, far and wide, this feature can save you from digging through the directory tree. It's particularly handy if you store documents on different drives or on a network. And just like the Recent Documents list, you can pin folders to the Recent Places list.

Protected View

Even something that seems as innocent as an Excel file can't always be trusted. Protected view is a new security feature in Excel 2010. It opens potentially risky Excel files in a specially limited Excel window. You'll know that you're in protected view because Excel doesn't allow you to edit any of the data in the workbook, and it displays a message bar at the top of the window (Figure 1-28).

Excel automatically uses protected view when you download a spreadsheet from the Web or open it from your email inbox. This is actually a huge convenience, because Excel doesn't need to hassle you with questions when you try to view the file (such as "are you sure you want to open this file?"). Because Excel's protected view has bulletproof security, it's a safe way to view the most suspicious spreadsheet.

At this point, you're probably wondering about the risks of rogue spreadsheets. Truthfully, they're quite small. The most obvious danger is *macro code*: miniature programs that are stored in a spreadsheet file and can perform Excel tasks. Poorly written or malicious macro code can tamper with your Excel settings, lock up the program, and even scramble your data. But before you panic, consider this: Excel macro viruses are very rare, and the .xlsx file format doesn't even allow macro code. Instead, macro-containing files must be saved as .xlsm or .xlsb files.

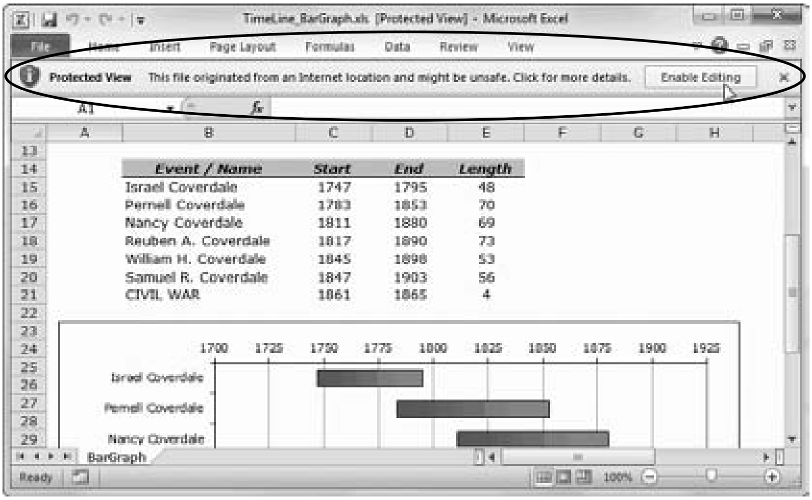

Figure 1-28: *Currently, this file is in protected view. If you decide that it's safe and you need to edit its content, click the Enable Editing button to open in the normal Excel window with no security safeguards.*

The more subtle danger is that crafty hackers could create corrupted Excel files that might exploit tiny security holes in the program. One of these files could scramble Excel's brains in a dangerous way, possibly causing it to execute a scrap of malicious computer code that could do almost anything. Once again, this sort of attack is extremely rare. It might not even be possible with the up-to-date .xlsx file format. But protected view completely removes any chance of an attack, which helps corporate bigwigs sleep at night.

Opening Files—with a Twist

The Open dialog box harbors a few tricks. To see these hidden secrets, first select the file you want to use (by clicking it once, not twice), and then click the drop-down arrow on the right-side of the Open button. A menu with several additional options appears, as shown in Figure 1-29.

Here's what these different choices do:

- **Open** opens the file in the normal way.
- **Open Read-Only** opens the file, but won't let you *save* changes. This option is great if you want to make sure you don't accidentally overwrite an existing file. (For example, if you're using last month's sales invoice as a starting point for this month's sales invoice, you might use Open Read-Only to make sure you can't accidentally wipe out the existing file.) If you open a document in read-only mode, you can still make changes—you just have to save the file with a new file name (choose File→Save As).

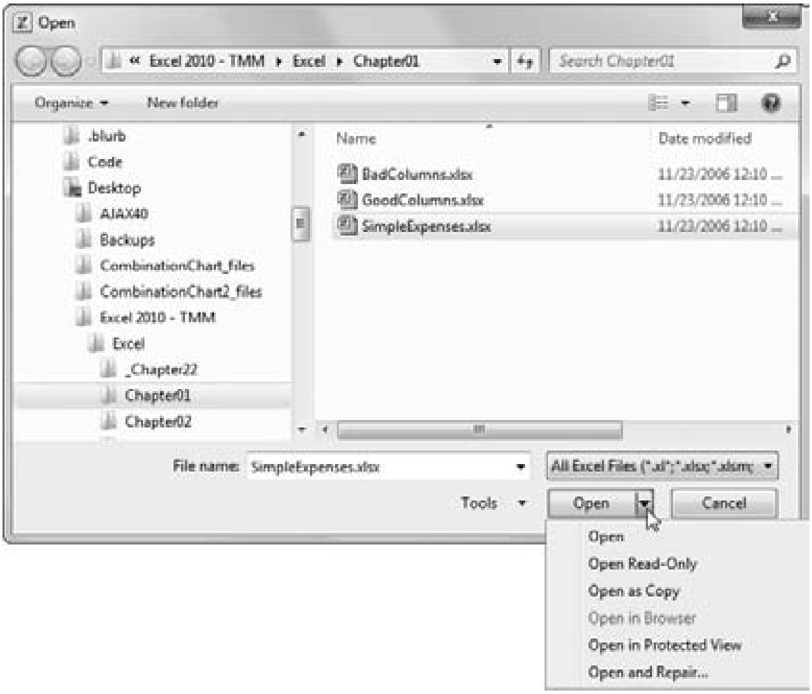

***Figure 1-29:** Why settle for the plain-vanilla Open command when you have all these other choices?*

- **Open as Copy** creates a copy of the spreadsheet file in the same folder. If your file is named Book1.xlsx, the copy will be named "Copy of Book1.xlsx". This feature comes in handy if you're about to start editing a spreadsheet and want to be able to look at the last version you saved. Excel won't let you open the same file twice. However, you can load the previous version by selecting the same file and using "Open as Copy". (Of course, this technique works only when you have changes you haven't saved yet. Once you save the current version of a file, the older version is overwritten and lost forever.)
- **Open in Browser** is only available when you select an HTML file. This option allows you to open the HTML file in your computer's web browser, which is something you'll only want to attempt when trying to convert Excel files to web pages (page 47).
- **Open in Protected View** prevents a potentially dangerous file from running any code. However, you'll also be restrained from editing the file, as explained on page 57.

- **Open and Repair** is useful if you need to open a file that's corrupted. If you try to open a corrupted file by just clicking Open, Excel warns you that the file has problems and refuses to open it. To get around this problem, you can open the file using the "Open and Repair" option, which prompts Excel to make the necessary corrections, display them for you in a list, and then open the document. Depending on the type of problem, you might not lose any information at all.

Working with Multiple Open Spreadsheets

As you open multiple spreadsheets, Excel creates a new window for each one. Although this helps keep your work separated, it can clause a bit of clutter and make it harder to track down the window you really want. Fortunately, Excel provides a couple of shortcuts that are indispensable when dealing with several spreadsheets at a time:

- To jump from one spreadsheet to another, find the window in the View→Window→Switch Windows list, which includes the file name of all the currently open spreadsheets (Figure 1-30).

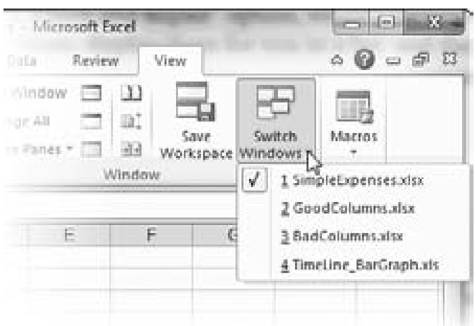

Figure 1-30:
When you have multiple spreadsheets open at the same time, you can easily move from one to the other using the Switch Windows list.

- To move to the next spreadsheet, use the keyboard shortcut Ctrl+Tab or Ctrl+F6.
- To move to the previous spreadsheet, use the shortcut key Ctrl+Shift+Tab or Ctrl+Shift+F6.

When you have multiple spreadsheets open at the same time, you need to take a little more care when closing a window so you don't accidentally close the entire Excel application—unless you want to. Here are your choices:

- **You can close all the spreadsheets at once**. To do so, you need to close the Excel window. Select File→Exit from any active spreadsheet, or just hold down Shift while you click the close icon (the infamous X button) in the top-right corner of the Excel window.

- **You can close a single spreadsheet**. To do so, right-click the spreadsheet on the taskbar and click Close. Or switch to the spreadsheet you want to close (by clicking the matching taskbar button) and then choose File→Close.

Note: One of the weirdest limitations in Excel occurs if you try to open more than one file with the same name. No matter what steps you take, you can't coax Excel to open both of them at once. It doesn't matter if the files have different content or if they're in different folders or even different drives. When you try to open a file that has the same name as a file that's already open, Excel displays an error message and refuses to go any further. Sadly, the only solution is to open the files one at a time, or rename one of them.

POWER USERS' CLINIC

Handy Options for Opening and Saving Files

If you're in the habit of configuring your software programs to get the most out of them, you'll be happy to hear that Excel has several useful details to tweak. To see them, choose File→Options.

Here are the most useful things you can do:

- **Adjust your starting point**. When you open a file or save it for the first time, Excel starts you off in your personal documents folder. This folder is a Windows-specific folder that many programs assume you use for all your files. If you don't use this folder, you can tell Excel to look elsewhere when saving and opening files. Choose the Save section, and then look under the "Save workbooks" heading for the "Default file location" text box. You can modify it so that it points to the folder where you usually store files (as in *C:\John Smith\MyExcel Files*). Sadly, you can't browse and pick the path from a dialog box–instead, you need to type it in by hand.
- **Keep track of more recent documents**. Why stick with 20 recent documents when you can show dozens? If you want to keep track of more recent work, and you aren't deterred by a long Recent Documents list, choose the Advanced section, scroll down to the Display group of settings, and then change the "Show this number of Recent Documents." You can pick any number from 0 to 50.
- **Change the standard file type**. Most Excel fans prefer the new .xlsx file format, which Excel uses every time you save a new file (unless you explicitly choose another option in the "Save as type" list). However, if you've decided that something else suits you better, like the binary .xlsb format (page 39) or the legacy .xls format, you can tell Excel to use this by default instead. Choose the Save section, look under the "Save workbooks" heading, and then change the "Save files in this format" setting by choosing another file type from the list.
- **Get started with a bang.** You can tell Excel to automatically open a whole group of spreadsheet files every time it starts up. To find this setting, choose the Advanced section, and then scroll to the General group of settings. You can use the "At startup, open all files in" text box to specify a folder where you put all the Excel files on which you're currently working. Then, the next time you start Excel, it automatically opens every Excel file it finds in a separate window. Of course, if you decide to use this option, make sure you don't clutter your in-progress folder with too many files, or Excel opens a dizzying number of windows when it starts.

CHAPTER
2

Adding Information to Worksheets

Now that you've created a basic worksheet, and you're acquainted with Excel's spiffy interface, it's time to get down and dirty adding data. Whether you want to plan your household budget, build a sales invoice, or graph your soaring (or plunging) net worth, you first need to understand how Excel interprets the information you put in your worksheet.

Depending on what kind of data you type into a cell, Excel classifies it as a date, a number, or a piece of text. In this chapter, you'll learn how Excel makes up its mind and how you can make sure it makes the right decision. You'll also learn how to use Excel's best timesavers, including the indispensable Undo feature and the spell checker. Finally, you'll consider inserting something completely different into your worksheets: a Web-style hyperlink.

Adding Different Types of Data

One of Excel's most important features is its ability to distinguish between different types of information. A typical worksheet contains both text and numbers. There isn't a lot you can do in Excel with ordinary text (other than alphabetize a list, perform a simple spell check, and apply some basic formatting). On the other hand, Excel gives you a wide range of options for numeric data. For example, you can string your numbers together into complex calculations and formulas, or you can graph them on a chart. Programs that don't distinguish between text and numbers—like Microsoft Word, for example—can't provide these features.

Most of the time, when you enter information in Excel, you don't explicitly indicate the type of data. Instead, Excel examines the information you've typed in and, based on your formatting and other clues, classifies it automatically. Excel distinguishes between four core data types:

- **Ordinary text**. This data type includes column headings, descriptions, and any content that Excel can't identify as one of the other data types.
- **Numbers**. This data type includes currency values, integers, fractions, percentages, and every other type of numeric data. Numbers are the basic ingredient of most Excel worksheets.
- **Dates and times**. This data type includes dates (like Oct 3, 2010), times (like 4:30 p.m.), and combined date and time information (like Oct 3, 2010, 4:30 p.m.). You can enter date and time information in a variety of formats.
- **True or false values**. This data type (known in geekdom as a *Boolean* value) can contain one of two things: TRUE or FALSE (displayed in all capitals). You don't need Boolean data types in most worksheets, but they're useful in worksheets that include Visual Basic macro code (see Chapter 29) or use complex formulas that evaluate conditions (see Chapter 13).

One useful way to tell how Excel is interpreting your data is to look at cell alignment, as explained in Figure 2-1.

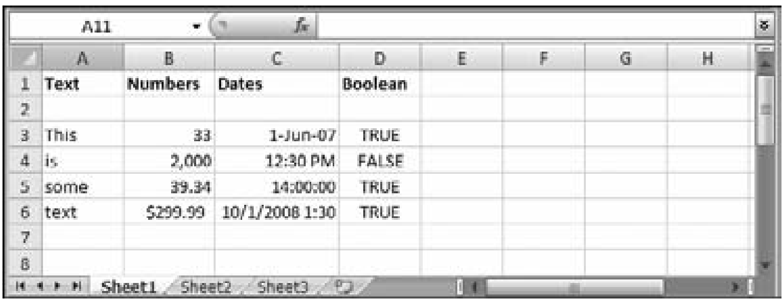

Figure 2-1: *Unless you explicitly change the alignment, Excel always left-aligns text (that is, it lines it up against the left edge of a cell), as in column A. On the other hand, it always right-aligns numbers and dates, as in columns B and C. And it centers Boolean values, as in column D.*

Note: The standard alignment of text and numbers doesn't just represent the whims of Excel—it also matches the behavior you want most of the time. For example, when you type in text, you usually want to start at the left edge so that subsequent entries in a column line up. But when entering numbers, you usually want them to line up on the *decimal point* so that it's easier to scan a list of numbers and quickly spot small and large values. Of course, if you don't like Excel's standard formatting, you're free to change it, as you'll see in Chapter 5.

As Figure 2-1 shows, Excel can display numbers and dates in several different ways. For example, some of the numbers include decimal places, one uses a comma, and one has a currency symbol. Similarly, one of the time values uses the 12-hour clock, while another uses the 24-hour clock. Other entries include only date information or both date and time information. You assume that when you type in a number, it will appear in the cell exactly the way you typed it. For example, when you type 3-comma-0-0-0 you expect to see 3,000. However, that doesn't always happen. To see the problem in action, try typing *3,000* in a cell. It shows up exactly the way you entered it. Then, type over that value with *2000*. The new number appears as 2,000. In this example, Excel remembers your first entry, and assumes that you want to use thousand separators in this cell *all the time*.

These differences may seem like a spreadsheet free-for-all, but don't despair—you can easily set the formatting of numbers and dates. (In fact, that's the subject of Chapter 5.) At this point, all you need to know is that the values Excel *stores* in each cell don't need to match exactly the values that it *displays* in each cell. For example, the number 4300 could be formatted as plain old 4300 or as the dollar amount $4,300. Excel lets you format your numbers so you have exactly the representation you want. At the same time, Excel treats all numbers equivalently, no matter how they're formatted, which lets you combine them together in calculations. Figure 2-2 shows you how to find the underlying stored value of a cell.

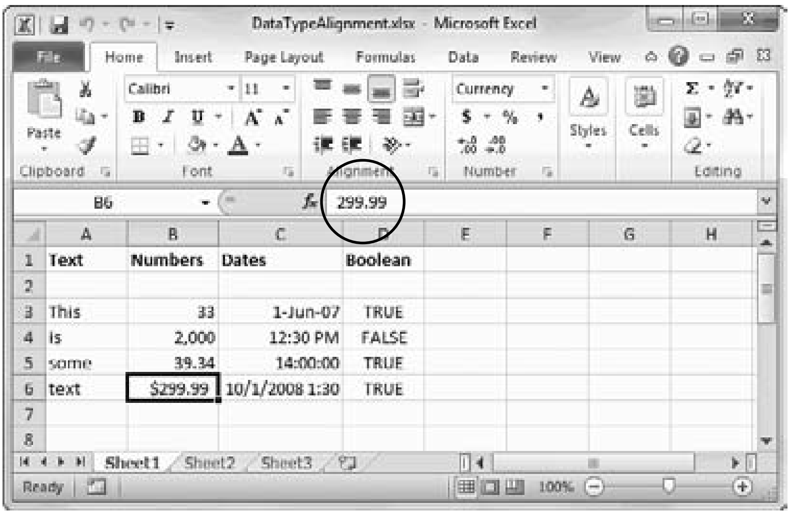

Figure 2-2: *You can see the underlying value that Excel is storing for a cell by selecting the cell and then glancing at the formula bar. In this sheet, you can see that the value $299.99 is actually stored without the dollar currency symbol, which Excel applied only as part of the display format. Similarly, Excel stores the number 2,000 without the comma; it stores the date 1-Jun-10 as 6/1/2010; the time 12:30 p.m. as 12:30:00 PM, and the time 14:00:00 as 2:00:00 PM.*

Note: Excel assigns data types to each cell in your worksheet, and you can't mix more than one data type in the same cell. For example, when you type in *44 fat cats*, Excel interprets the whole thing as text because it contains letters. If you want to treat 44 as a number (so that you can perform calculations with it, say), then you need to split this content into two cells–one that contains the number 44 and one that contains the remaining text.

By looking at cell alignment, you can easily tell how Excel is interpreting your data. That's helpful. But what happens when Excel's interpretation is at odds with your wishes? For example, what if you type in something you consider a *number* but Excel freakishly treats it as *text* or vice versa? The first step to solving this problem is grasping the logic behind Excel's automatic decision-making process.

How Excel Identifies Text

If your cell meets any of the following criteria, Excel automatically treats the content as ordinary *text*:

- **It contains any letters**. Thus, C123 is text, not a number.
- **It contains any punctuation that Excel can't interpret numerically**. Punctuation allowed in numbers and dates includes the comma (,), the decimal point (.), and the forward slash (/) or dash (-) for dates. When you type in any other punctuation, Excel treats the cell as text. Thus, 14! is text, not a number.

Occasionally, Excel reads your data the wrong way. For example, you may have a value—like a Social Security number or a credit card number—that's made up entirely of numeric characters but that you want to treat like text because you don't ever want to perform calculations with it. In this case, Excel doesn't know what you're up to, and so it automatically treats the value as a number. You can also run into problems when you precede text with the equal sign (which tells Excel that you have a formula in progress), or when you use a series of numbers and dashes that you don't intend to be part of a date (for example, you want to enter 1-2-3 but you don't want Excel to read it as January 2, 2003—which is what it wants to do).

In all these cases, the solution's simple. Before you type the cell value, start by typing an apostrophe ('). The apostrophe tells Excel to treat the cell content as text. Figure 2-3 shows you how it works.

Figure 2-3:
To have Excel treat any number, date, or time as text, just precede the value with an apostrophe (you can see the apostrophe in the formula bar but not in the cell). This worksheet shows the result of typing 1-2-3, both with and without the initial apostrophe. When you store 1-2-3 as text, Excel left-aligns it, as if it were an all-text cell (and puts a tiny green triangle in the corner of the cell to let you know you may have made a mistake). The date, on the other hand, is right-aligned.

When you precede a numeric value with an apostrophe, Excel checks out the cell to see what's going on. When Excel determines that it can represent the content as a number, it places a green triangle in the top left corner of the cell and gives you a few options for dealing with the cell, as shown in Figure 2-4.

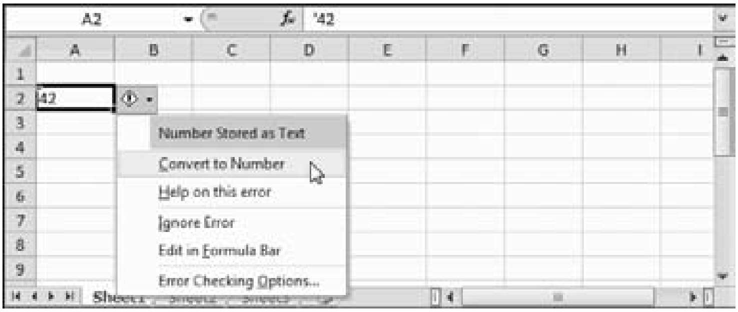

Figure 2-4: *In this worksheet, the number 42 is stored as text, thanks to the apostrophe that precedes it. Excel notices the apostrophe, wonders if it's an unintentional error, and flags the cell by putting a tiny green triangle in the top-left corner. If you move to the cell, an exclamation mark icon appears, and, if you click that, a menu appears, letting you choose to convert the number or ignore the issue for this cell. Excel provides a similar menu if you enter a text date that has a two-digit year, as in '1-1-07. In this case, the menu allows you to convert the two-digit date to a four-digit date that has a year starting with 19 or 20.*

Tip: When you type in either *false* or *true* (using any capitalization you like), Excel automatically recognizes the data type as Boolean value instead of text, converts it to the uppercase word FALSE or TRUE, and centers it in the cell. If you want to make a cell that contains *false* or *true* treat it as text and *not* as Boolean data, start by typing an apostrophe (') at the beginning of the cell.

How Excel Identifies Numbers

Excel automatically interprets any cell that contains only numeric characters as a number. In addition, you can add the following nonnumeric characters to a number without causing a problem:

- One decimal point (but not two). For example, 42.1 is a number, but 42.1.1 is text.
- One or more commas, provided you use them to separate groups of three numbers (like thousands, millions, and so on). Thus 1,200,200 is a valid number, but 1,200,20 is text.
- A currency sign ($ for U.S. dollars), provided it's at the beginning of the number.

- A percent symbol at the beginning or end of the number (but not both).
- A plus (+) or minus (–) sign before the number. You can also create a negative number by putting it in parentheses. In other words, entering (33) is the same as entering –33.
- An equal sign at the start of the cell. This tells Excel that you're starting a formula (page 227).

The most important thing to understand about entering numbers is that when you choose to add other details like commas or the dollar sign, you're actually doing two things at once: You're entering a value for the cell *and* you're setting the format for the cell, which affects how Excel displays the cell. Chapter 5 provides more information about number styles and shows how you can completely control cell formatting.

How Excel Identifies Dates and Times

When typing in a date, you have a choice of formats. You can type in a full date (like *July 4, 2010*) or you can type in an abbreviated date using dashes or slashes (like *7-4-2010* or *7/4/2010*), which is generally easier. If you enter some numbers formatted as a date, but the date you entered doesn't exist (like the 30th day in February or the 13th month), then Excel interprets it as text. Figure 2-5 shows you the options.

Figure 2-5: *Whichever way you type the date in a cell, it always appears the same on the formula bar (the specific formula bar display depends on the regional settings on your computer, explained next). To fine-tune the way the date appears in the worksheet, you can use the formatting features discussed on page 138.*

Because you can represent dates a few different ways, working with them can be tricky, and you're likely to encounter some unexpected behavior from Excel. Here are some tips for using dates, trouble-free:

- **Instead of using a number for the month, you can use a three-letter month abbreviation, but you must put the month in the middle**. In other words, you can use *7/4/2010* and *4/Jul/2010* interchangeably.

- **When you use a two-digit year as part of a date, Excel tries to guess whether the first two digits of the year should be 20 or 19**. When the two-digit year is from 00 to 29, Excel assumes it belongs to the 21st century. If the year is from 30 to 99, Excel plants it in the 1900s. In other words, Excel translates 7/4/29 into 7/4/2029, while 7/4/30 becomes 7/4/1930.

Tip: If you're a mere mortal and forget where the cutoff point is, then enter the year as a four-digit number, which prevents any confusion.

- **If you don't type in any year at all, Excel automatically assumes you mean the current year**. For example, when you enter 7/4, Excel inserts the date 7/4/2010 (assuming it's currently 2010 on your computer's internal clock). When you enter a date this way, the year component doesn't show up in the cell, but it's still stored in the worksheet (and visible on the formula bar).
- **Excel understands and displays dates differently depending on the regional settings on your computer**. Windows has a setting that determines how your computer interprets dates (see the next section, page 70.) On a computer that's configured with U.S. settings, Month-Day-Year is the standard progression. But on a UK-configured computer, Day-Month-Year is the deal. For example, in the U.S., either 11-7-09 or 11/7/09 is shorthand for November 7, 2009. In the UK or in Canada, the same notations refer to July 11, 2009.

  Thus, if your computer has U.S. regional settings turned on, and you type in *11/7/09*, then Excel understands it as November 7, 2009, and the formula bar displays 11/7/2009.

Note: The way Excel recognizes and displays dates varies according to the regional settings on your computer, but the way Excel *stores* dates does not. This feature comes in handy when you save a worksheet on one computer and then open it on another computer with different regional settings. Because Excel stores every date the same way, the date information remains accurate on the new computer, and Excel can display it according to the current regional settings.

Typing in times is more straightforward than typing in dates. You simply use numbers, separated by a colon (:). You need to include an hour and minute component at minimum (as in 7:30), but you can also add seconds (as in 7:30:10). You can use values from 1 to 24 for the hour part, though if your system's set to use a 12-hour clock, Excel converts the time accordingly (in other words, 19:30 becomes 7:30 PM). If you want to use the 12-hour clock when you type in a time, follow your time with a space and the letters P or PM (or A or AM).

Finally, you can create cells that have both date and time information. To do so, just type the date portion first, followed by a space, and then the time portion. For example, Excel happily accepts this combo: 7/4/2010 1:30 PM.

Behind the scenes, Excel stores dates as *serial numbers*. It considers the date January 1, 1900 to be day 1. January 2, 1900 is day 2, and so on, up through the year 9999. This system is quite nifty, because if you use Excel to subtract one date from another, then you actually end up calculating the difference in days, which is exactly what you want. On the other hand, it means you can't enter a date in Excel that's earlier than January 1, 1900 (if you do, Excel treats your date like text).

Similarly, Excel stores times as fractional numbers from 0 to 1. The number 0 represents 12:00 a.m. (the start of the day) and 0.99999 represents 11:59:59 p.m. (the end of the day). As with dates, this system allows you to subtract one time value from another. See Chapter 11 for more information on how to perform calculations that use dates and times.

Regional Dating

Windows has regional settings for your computer, which affect the way Microsoft programs understand things like dates and currency. You can change the settings, and they don't have to correspond to where you live—you can set them for your company headquarters on another continent, for instance. But keep in mind that these affect all the programs on your computer.

Every version of Windows uses the same system for regional settings. (Note for tech nerds: You can find them in the Region and Language section of the Control Panel.) But if you're not used to digging into your computer's Control Panel, you'll face a small problem, because every version of Windows puts them in a slightly different place. Here's the easiest way to find them:

- **If you're using Windows 7 or Windows Vista**: Click the Start button. Then, in the search box at the bottom of the Start menu, type *region*. When Region and Language appears in the list of matches, click it.
- **If you're using Windows XP:** Click the Start button and choose Run. In the Run window, type *intl.cpl* and click OK.

Either way, the Region and Language window will appear (see Figure 2-6). The most important setting is in the first box, which has a drop-down list you can use to pick the region you want, like English (United States) or Swedish (Finland).

If you want to fine-tune the settings in your region, click the "Additional settings" button (in Windows 7), the "Customize this format" button (in Windows Vista), or the Customize button (in Windows XP). Although the name has changed over the years, it's still the same button.

You might decide to customize your region's settings if you have a particular preference that doesn't match the standard options. For example, you might decide that you want U.K.-formatted dates on a computer that's set to use U.S.-region settings for everything else.

When you customize a region, a new window appears (Figure 2-7) with a long list of settings.

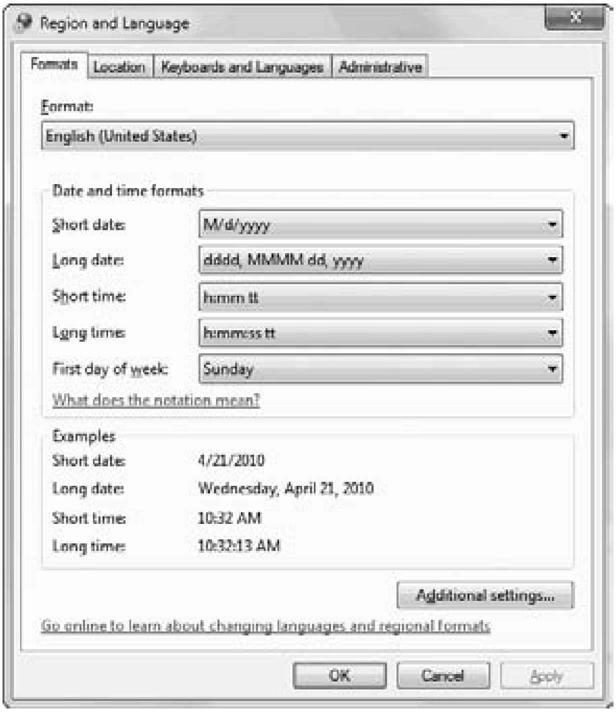

Figure 2-6:
In the Region and Language window, you choose a geographical region and your computer stores a set of preferences about number and date display. Excel heeds these settings.

Tip: No matter what the regional settings are, you can always use the international date standard when typing dates into Excel, which is Year/Month/Day, though you must supply a four-digit year (as in 2008/7/4). If you use a two-digit year, Excel assumes you're trying to use the Month-Day-Year or Day-Month-Year pattern.

Handy Timesavers

Some of Excel's frills can make your life easier when you're entering data in a worksheet. This section covers four such timesaving features: AutoComplete, AutoCorrect, AutoFill, and AutoFit.

Note: Excel really has two types of automatic features. First off, there are features that do things to your spreadsheets *automatically*, namely AutoComplete and AutoCorrect. Sometimes that's cool and convenient, but other times it can send you running for the old manual typewriter. Fortunately, you can turn off both. Excel also has "auto" features that really aren't that automatic. These include AutoFill and AutoFit, which never run on their own.

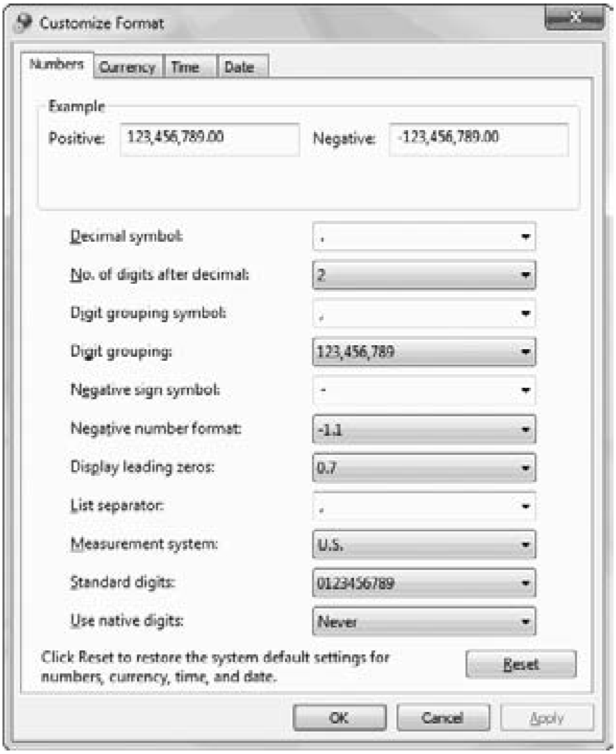

Figure 2-7: *Tweaking the regional settings on your computer gives you complete control over how Excel recognizes dates. Use the drop-down lists to specify the date separator, order of month, day, and year components in a date, and how Excel should interpret two-digit years. You can mix and match these settings freely, although you could wind up with a computer that's completely counterintuitive to other people.*

AutoComplete

Some worksheets require that you type in the same information row after row. For example, if you're creating a table to track the value of all your *Sesame Street* collectibles, you can type in *Kermit* only so many times before you start turning green. Excel tries to help you out with its AutoComplete feature, which examines what you type, compares it against previous entries in the same column, and, if it recognizes the beginning of an existing word, fills it in.

For instance, in your *Sesame Street* worksheet, if you already have Kermit in the Characters column, when you start typing a new entry in that column beginning with the letter K, Excel automatically fills in the whole word Kermit. Excel then selects the letters that it's added (in this case, *ermit*). You now have two options:

- **If you want to accept the AutoComplete text, move to another cell**. For example, when you hit the right arrow key or press Enter to move down, Excel leaves the word Kermit behind.

- **If you want to blow off Excel's suggestion, just keep typing**. Because Excel automatically selects the AutoComplete portion of the word (*ermit*), your next keystrokes type over that text. Or, if you find the AutoComplete text is distracting, press Delete to remove it right away.

Tip: When you want to use the AutoComplete text but change it slightly, turn on edit mode for the cell by pressing F2. Once you enter edit mode, you can use the arrow keys to move through the cell and make modifications.

AutoComplete has a few limitations. It works only with text entries, ignoring numbers and dates. It doesn't pay any attention to the entries you've placed in other columns. And if there's a blank row between the cell you're working on and the values above it, Excel ignores them, assuming that they're part of a different list that just happens to use the same column.

Finally, Excel won't give you a suggestion unless the text you've typed in matches the value in another cell *unambiguously*. This means that when your column contains two words that start with K, like Kermit and kerplop, Excel doesn't make any suggestion when you type K into a new cell, because it can't tell which option is the most similar. But when you type *Kerm*, Excel realizes that kerplop isn't a candidate, and it supplies the AutoComplete suggestion Kermit.

If you find AutoComplete annoying, you can get it out of your face with a mere click of the mouse. Just choose File→Options, choose the Advanced section, and look under the "Editing options" heading for the "Enable Auto-Complete for cell values" setting. Turn this setting off to banish AutoComplete from your life.

AutoCorrect

As you type text in a cell, AutoCorrect cleans up behind you—correcting things like wrongly capitalized letters and common misspellings. AutoCorrect is subtle enough that you may not even realize it's monitoring your every move. To get a taste of its magic, look for behaviors like these:

- If you type *HEllo*, AutoCorrect changes it to *Hello*.
- If you type *friday*, AutoCorrect changes it to *Friday*.
- If you start a sentence with a lowercase letter, AutoCorrect uppercases it.
- If you scramble the letters of a common word (for example, typing *thsi* instead of *this*, or *teh* instead of *the*), AutoCorrect replaces the word with the proper spelling.
- If you accidentally hit Caps Lock key, and then type *jOHN sMITH* when you really wanted to type *John Smith*, Excel not only fixes the mistake, it also switches off the Caps Lock key.

Note: AutoCorrect doesn't correct most misspelled words, just common typos. To correct other mistakes, use the spell checker described on page 85.

For the most part, AutoCorrect is harmless and even occasionally useful, as it can spare you from delivering minor typos in a major report. But if you need to type irregularly capitalized words, or if you have a garden-variety desire to rebel against standard English, then you can turn off some or all of the AutoCorrect actions.

To reach the AutoCorrect settings, choose File→Options. Choose the Proofing section, and then click the AutoCorrect Options button. (All Auto-Correct options are language-specific, and the title of the dialog box that opens indicates the language you're currently using.) Most of the actions are self-explanatory, and you can turn them off by turning off their checkboxes. Figure 2-8 explains the "Replace text as you type" option, which isn't just for errors.

Figure 2-8:
Under "Replace text as you type" is a long list of symbols and commonly misspelled words (the column on the left) that Excel automatically replaces with something else (the column on the right). But what if you want the copyright symbol to appear as a C in parentheses? You can remove individual corrections (select one, and then click Delete), or you can change the replacement text. And you can add your own rules. For example, you might want to be able to type PESDS and have Excel insert Patented Electronic Seltzer Delivery System. Simply type in the "Replace" and "With" text, as shown here, and then click OK.

Tip: For really advanced AutoCorrect settings, you can use the Exceptions button to define cases where Excel *won't* use AutoCorrect. When you click this button, the AutoCorrect Exceptions dialog box appears with a list of exceptions. For example, this list includes abbreviations that include the period but shouldn't be capitalized (like pp.) and words where mixed capitalization is allowed (like WordPerfect).

AutoFill

AutoFill is a quirky yet useful feature that lets you create a whole column or row of values based on just one or two cells that Excel can extrapolate into a series. Put another way, AutoFill looks at the cells you've already filled in a column or row, and then makes a reasonable guess about the additional cells you'll want to add. People commonly use AutoFill for sequential numbers, months, or days.

Here are a few examples of lists that AutoFill can and can't work with:

- The series 1, 2, 3, 4 is easy for Excel to interpret—it's a list of steadily increasing numbers. The series 5, 10, 15 (numbers increasing by five) is just as easy. Both of these are great AutoFill candidates.
- The series of part numbers CMP-40-0001, CMP-40-0002, CMP-40-0003 may seem more complicated because it mingles text and numbers. But clever Excel can spot the pattern, as long as the numbers are at the end of the value (so CMP-40-A won't work).
- Excel readily recognizes series of months (*January, February, March*) and days (*Sun, Mon, Tue*), either written or in three-letter abbreviations.
- A list of numbers like 47, 345, 6 doesn't seem to follow a regular pattern. But by doing some analysis, Excel can guess at a relationship and generate more numbers that fit the pattern. There's a good chance, however, that these won't be the numbers you want, so take a close look at whatever Excel adds in cases like these.

Bottom line: AutoFill is a great tool for generating simple lists. When you're working with a complex sequence of values, it's no help—unless you're willing to create a custom list (page 76) that spells it out for Excel.

Tip: AutoFill doubles as a quick way to copy a cell value multiple times. For example, if you select a cell in which you've typed *Cookie Monster*, you can use the AutoFill technique described below to fill every cell in that row or column with the same text.

To use AutoFill, follow these steps:

1. **Fill in a couple of cells in a row or column to start off the series.**

   Technically, you can use AutoFill if you fill in only one cell, although this approach gives Excel more room to make a mistake if you're trying to generate a series. Of course, when you want to copy only a single cell several times, one cell is a sufficient start.

2. **Select the cells you've entered so far. Then click (and hold) the small black square at the bottom-right corner of the selected box.**

   You can tell that your mouse is in the correct place when the mouse pointer changes to a plus symbol (+).

3. **Drag the border down (if you're filling a column of items) or to the right (if you're filling a row of items).**

   As you drag, a tooltip appears, showing the text that Excel is generating for each cell.

   While you're dragging, you can hold down Ctrl to affect the way that Excel fills a list. When you've already filled in at least *two* cells, Ctrl tells Excel to just copy the list multiple times, rather than look for a pattern. When you want to expand a range based on just *one* cell, Ctrl does the opposite: It tells Excel to try to predict a pattern, rather than just copy it.

   When you release the mouse, Excel automatically fills in the additional cells, and a special AutoFill icon appears next to the last cell in the series, as shown in Figure 2-9.

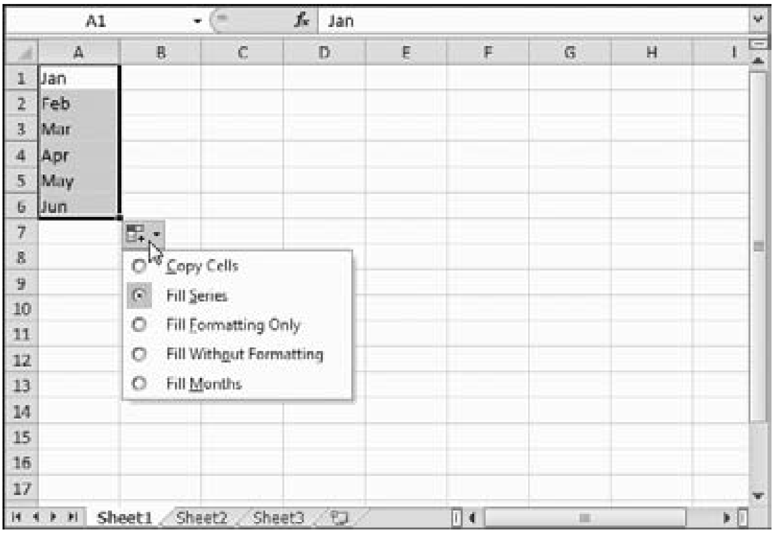

Figure 2-9: *After AutoFill does its magic, Excel displays a menu that lets you fill the series without copying the formatting, or copy the formatting without filling the series. You can also choose to copy values instead of generating a list. For example, if you choose to copy values–or Copy Cells, as Excel calls it–then in the two-item series Jan, Feb, you end up with Jan, Feb, Jan, Feb, rather than Jan, Feb, Mar, Apr.*

Custom AutoFill lists

Excel stores a collection of AutoFill lists that it refers to every time you use the feature. You can add your own lists to the collection, which extends the series AutoFill recognizes. For example, Excel doesn't come set to understand Kermit, Cookie Monster, Grover, Big Bird, Oscar, and Snuffleupagus as a series, but you can add it to the mix.

But why bother to add custom lists to Excel's collection? After all, if you need to type in the whole list before you use it, is AutoFill really saving you any work? The benefit occurs when you need to create the same list in *multiple* worksheets, in which case you can type it in just once and then use AutoFill to recreate it as often as you'd like.

To create a custom list, follow these steps:

1. **Choose File→Options.**

   The familiar Excel Options window appears.

2. **Choose the Advanced section, scroll down to the "General" heading and, at the bottom of that section, click the Edit Custom Lists button.**

   Here, you can take a gander at Excel's predefined lists and add your own (Figure 2-10).

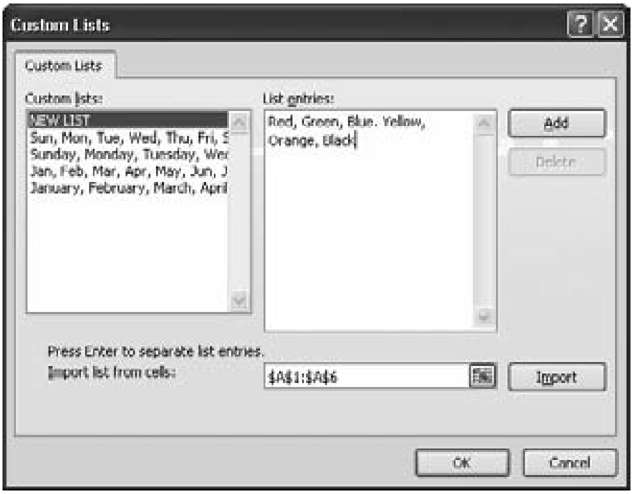

Figure 2-10:
Here, a new custom list of colors is being added.

3. **In the "Custom lists" box on the left side of the dialog box, select NEW LIST.**

   This action tells Excel that you're ready to create a new list.

4. **In the "List entries" box on the right side of the dialog box, type in your list.**

   Separate each item with a comma or by pressing Enter. The list in Figure 2-10 shows a series of color names separated by commas.

   If you've already typed your list into your worksheet, you can save some work. Instead of retyping the list, click inside the text box labeled "Import list from cells." Then, click the worksheet and drag to select the cells that contain the list. (Each item in the list must be in a separate cell, and the whole list should be in a series of adjacent cells in a single column or a single row.) When you're finished, click Import, and Excel copies the cell entries into the new list you're creating.

5. **Click Add to store your list.**

   At any later point in time, you can return to this dialog box, select the saved list, and modify it in the window on the right. Just click Add to commit your changes after making a change, or click Delete to remove the list entirely.

6. **Click OK to close the Custom Lists dialog box and OK again to close the Excel Options window.**

   You can now start using the list with the current worksheet or in a new worksheet. Just type the first item in your list and then follow the AutoFill steps outlined in the previous section.

AutoFit

Page 24 (Figure 1-5) explained how you can drag the edge of a column to resize it. For greater convenience, Excel also provides an AutoFit feature that automatically enlarges or shrinks a column to fit its content perfectly.

The AutoFit feature springs into action in three situations:

- When you type a number or date that's too wide to fit into a cell, Excel automatically widens the column to accommodate the new content. (Excel doesn't automatically expand columns when you type in text, however.)
- If you double-click the right edge of a column header, Excel automatically sizes the column to fit the widest entry it contains. This trick works for all types of data, including dates, numbers, and text.
- If you choose Home→Cells→Format→AutoFit Column Width, Excel automatically sizes the column to fit the content in the current cell. Or you can select a group of cells (see page 95) and use this command to size the column to fit the widest value in the group. This feature is helpful if you have a column that's made up of relatively narrow entries but also has a long column title. In this situation, you may not want to expand the column to the full width of the title. Instead, you may wish to size the column to fit the contents and allow the title to spill over to the next column.

While AutoFit automatically widens columns when you type in a number or date in a cell, you can still shrink a column after you've entered your information. Keep in mind, however, that when your columns are too narrow, Excel displays the cell data differently, depending on the type of information. When your cells contain *text*, it's entirely possible for one cell to overlap (and thereby obscure) another, a problem first described in Chapter 1. However, if Excel allowed truncated *numbers*, it could be deceiving. For example, if you squash a cell with the price of espresso makers so that they appear to cost $2 (instead of $200), you might wind up ordering a costly gift for all your coworkers. To prevent this problem, Excel never truncates a number or date. Instead, if you've shrunk a cell's width so that the number can't fit, you'll see a series of number signs (like #####) filling in the whole cell. This warning is just Excel's way of telling you that you're out of space. Once you enlarge the column by hand (or by using AutoFit), the original number reappears. (Until then, you can still see the number stored in the cell by moving to the cell and looking in the formula bar.)

GEM IN THE ROUGH

A Few More Ways to Adjust Column Width

Excel gives you the ability to precisely control column widths. To change the width of a column, right-click the column header at the top of the column, and then choose Column Width. The standard unadjusted column size is a compact 8.43 characters, but you can change that to any number of characters. (Remember that because different fonts use different size letters, the number of characters you specify here may not correspond directly to the number of characters in your column.)

You can also adjust multiple column widths at the same time. Just select multiple columns (click the first column header, and then drag to the left or to the right to select more columns). Now, when you apply a new width, Excel uses it for all the selected columns.

Finally, you can customize the standard width for columns, which is the width that Excel assigns to columns in every new worksheet that you create. To set the standard width, choose Home→Cells→Format→Default Width from the menu, and then change the number.

Dealing with Change: Undo, Redo, and AutoRecover

While editing a worksheet, an Excel guru can make as many (or more) mistakes as a novice. These mistakes include copying cells to the wrong place, deleting something important, or just making a mess of the cell formatting. Excel masters can recover much more quickly, however, because they rely on Undo and Redo. Get in the habit of calling on these features, and you'll be well on your way to Excel gurudom.

Undo and Redo

As you create your worksheet, Excel records every change you make. Because the modern computer has vast resources of extra memory and computing power (that is, when it's not running the latest three-dimensional real-time action game), Excel can keep this log without slowing your computer down one bit.

If you make a change to your worksheet that you don't like (say you inadvertently delete your company's entire payroll plan), you can use Excel's Undo history to reverse the change. In the Quick Access toolbar, simply click the Undo button (Figure 2-11), or press the super-useful keyboard shortcut Ctrl+Z. Excel immediately restores your worksheet to its state just before the last change. If you change your mind again, you can revert to the changed state (known to experts as "undoing your undo") by choosing Edit→Redo, or pressing Ctrl+Y.

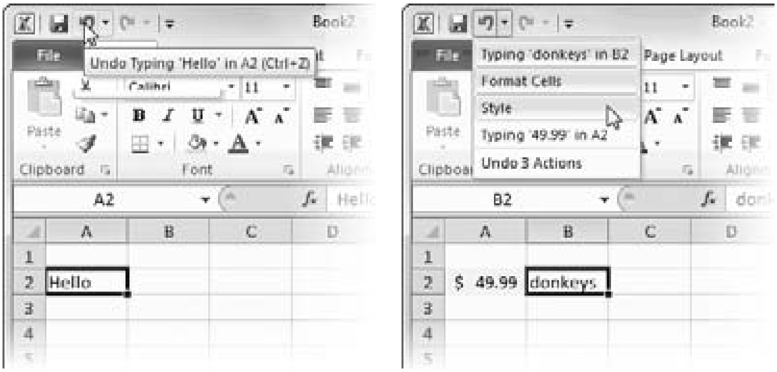

Figure 2-11: *Left: When you hover over the Undo button, you see a text description for the most recent action, which is what you'll undo if you click away. Here, the text Hello has just been typed into a cell, as Excel explains.*

Right: Click the down-pointing arrow on the edge of the Undo button to see a history of all your recent actions, from most recent (top) to oldest (bottom). If you click an item that's down the list, you'll perform a mega-undo operation that undoes all the selected actions. In this example, three actions are about to be rolled back–the text entry in cell B2, and a style and format operation that changed the number format and background fill of cell A2.

Things get interesting when you want to go farther back than just one previous change, because Excel doesn't just store one change in memory. Instead, it tracks the last *100* actions you made. And it tracks just about anything you do to a worksheet, including cell edits, cell formatting, cut and paste operations, and much more. As a result, if you make a series of changes you don't like, or if you discover a mistake a little later down the road, you can step back through the entire series of changes, one at a time. Every time you press Ctrl+Z, you go back one change in the history. This ability to reverse multiple changes makes Undo one of the most valuable features ever added to any software package.

Note: The Undo feature means you don't need to be afraid of performing a change that may not be what you want. Excel experts often try out new actions, and then simply reverse them if the actions don't have the desired effect.

The Undo feature raises an interesting dilemma. When you can go back 100 levels into the history of your document, how do you know exactly what changes you're reversing? Most people don't remember the previous 100 changes they made to a worksheet, which makes it all too easy to reverse a change you actually *want* to keep. Excel provides the solution by not only keeping track of old worksheet versions, but also by keeping a simple description of each change. You don't see this description if you use the Ctrl+Z and Ctrl+Y shortcuts. However, when you hover over the button in the Quick Access toolbar, you'll see the action you're undoing listed there.

For example, consider what happens if you type *hello* into cell A1 and then clear the cell by pressing the Delete key. Now, when you hover over the Undo button in the Quick Access toolbar, it says "Undo Clear (Ctrl+Z)". If you click Undo, the word *hello* returns. And if you hover over the Undo button again, it now says, "Undo Typing 'hello' in A2 (Ctrl+Z)", as shown in Figure 2-11, left.

Incidentally, Excel doesn't clear the Undo history when you save your spreadsheet. Instead, the Undo history remains until you close your workbook.

Note: Very rarely, when you perform an advanced analysis task with an extremely complex worksheet, Excel may decide it can't afford to keep an old version of your worksheet in memory. When Excel hits this point, it warns you before you make the change, and gives you the chance to either cancel the edit or continue (without the possibility of undoing the change). In this rare situation, you may want to cancel the change, save your worksheet as a backup, and then continue.

GEM IN THE ROUGH

Using Repeat to Automate Repetitive Tasks

Redo is commonly used to reverse an Undo. In other words, if you cancel an action and then change your mind, you can use Redo to quickly reapply the change. But Redo also has a much more interesting relative called Repeat, which lets you repeat any action multiple times. The neat thing is that you can repeat this action *on other cells.*

For example, imagine you hit Ctrl+B to change a cell to bold. If you move to another cell and hit Ctrl+Y, Excel *repeats* your operation and applies the bold formatting to the new cell.

In this case, you're not saving much effort, because it's just as easy to use Ctrl+B or Ctrl+Y. However, imagine you finish an operation that applies a set of sophisticated formatting changes to a cell. For example, say you go to the Home→Font section of the ribbon, click the dialog launcher to get to the Format Cells dialog box, and then increase the font size, bold the text, and apply a border around the cell. (Chapter 5 tells you how to do these things.) Now, when you move to another cell and press Ctrl+Y, Excel applies all the changes at once—which is much easier than calling up the Format Cells dialog box again and selecting the same options.

The trick when using Repeat is to make sure you don't perform another action until you've finished repeating your changes. For example, if you make some formatting changes and then stop to delete an incorrect cell value, then you can no longer use Repeat to apply your formatting because Excel applies the last change that you made—in this case, clearing the cell. (Of course, when you mistakenly apply Repeat, you can always call on Undo to get out of the mess.)

If you're ever in doubt about what will happen when you use Repeat, just hover over the Undo button in the Quick Access toolbar. You'll see a text description, like Undo Font or Undo Format Cells, which describes your last action.

AutoRecover

Undo and Redo are undeniably useful. However, there's another data-saving trick that just might save your sanity (and your job) if you need it in a pinch: AutoRecover.

In Chapter 1, you learned how AutoRecover can help you out in the event of catastrophic failure—for example, if Excel crashes or your computer loses power. In this situation, Excel will present you with your automatically backed-up work the next time you start the program. However, you can also use AutoRecover when the problem is your fault—for example, if you've just realized that you wiped out a critical column of numbers 20 minutes ago while editing your spreadsheet.

In this situation, you can open one of the automatically saved versions of your spreadsheet in a separate Excel window. You can then find the missing or modified data and copy it back to the current version of your spreadsheet. Or, if you've really made a mess of things, you can revert to the older version with a single click.

Tip: AutoRecover is perfect in situations where Undo isn't enough–for example, if the change was made long ago and isn't in the Undo history anymore, or if you don't want to reverse every change you've made since that mistake in order to fix the problem.

To see the AutoRecover files for your current workbook, choose File→Info. Excel switches into backstage view and shows three sections and three big buttons (Figure 2-12). The bottom section is named Versions, and it lists all the automatically saved copies of your work.

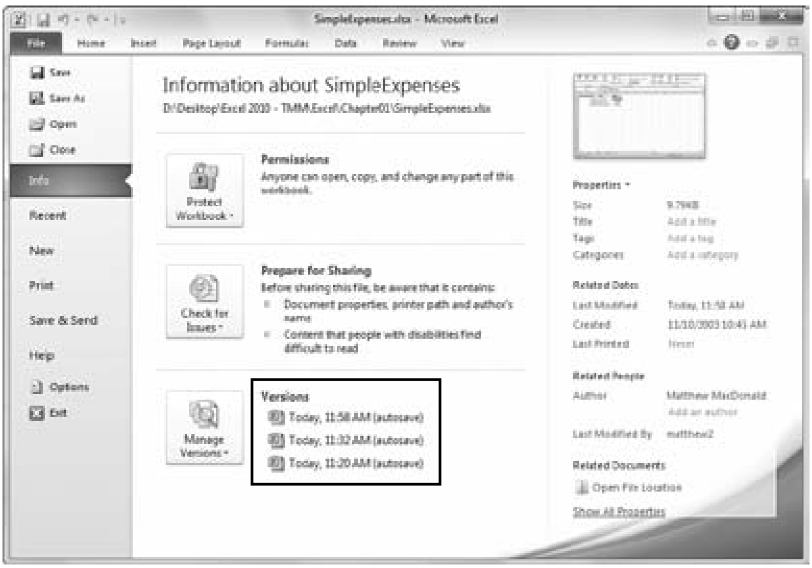

Figure 2-12: *As long as you're making changes, Excel will continue making automatic backups according to the AutoRecover settings you've configured (page 54). If you're using the standard AutoRecover settings, Excel will create a new backup as often as every 10 minutes. Here, there are three previous versions of the current workbook. Backups are listed from most recent to oldest and identified by the time they were made.*

To take a look at an AutoRecover backup, just click the file you want in the Versions list. Excel opens the backup in a new window, but it adds a helpful warning to make sure you don't confuse it with the current version of the file (Figure 2-13).

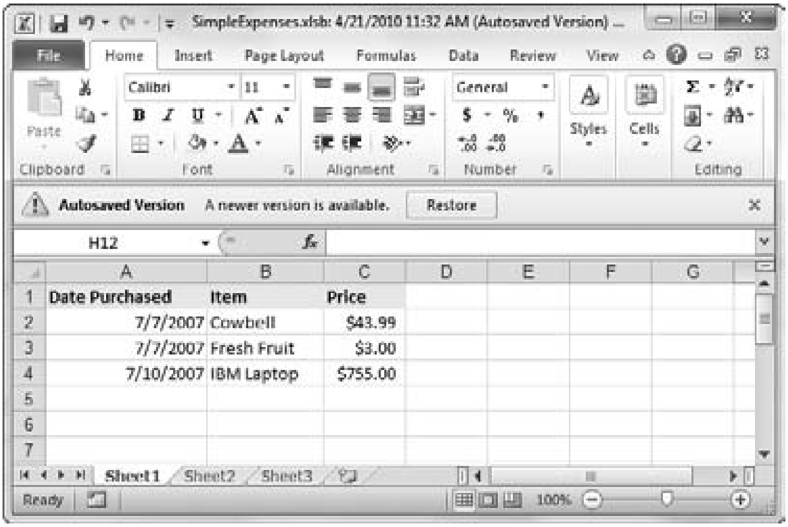

Figure 2-13: *The yellow bar at the top of the worksheet indicates that this is an autosaved version of the file. You can copy data from this window to the window with your current version, or you can click Restore to abandon your current version and use this in its place.*

Ordinarily, Excel deletes AutoRecover files when you close the workbook. However, there's one significant exception: If you close a workbook without saving it, the AutoRecover version will remain stored on your computer. To get it back, open the file again and look at the list of old Versions (choose File→Info, as shown in Figure 2-12). There will be just one backed-up copy, with the note "(when I closed without saving)" next to its name. This is the version of your spreadsheet that you didn't save.

Interestingly, if you start a new workbook, work on it long enough for Excel to create an AutoRecover backup, and then exit without saving it, the AutoRecover file will *still* remain, even though you haven't saved your work even once. A similar thing happens if you open a file, edit it for a while, and close it without saving your changes. Excel calls this sort of file a *draft*, and you can dig it up later out of the recent files list (choose File→Recent). To dig it up later, choose File→Info, click the Manage Versions button, and choose Recover Unsaved Workbooks. An Open dialog box appears with a list of drafts (Figure 2-14).

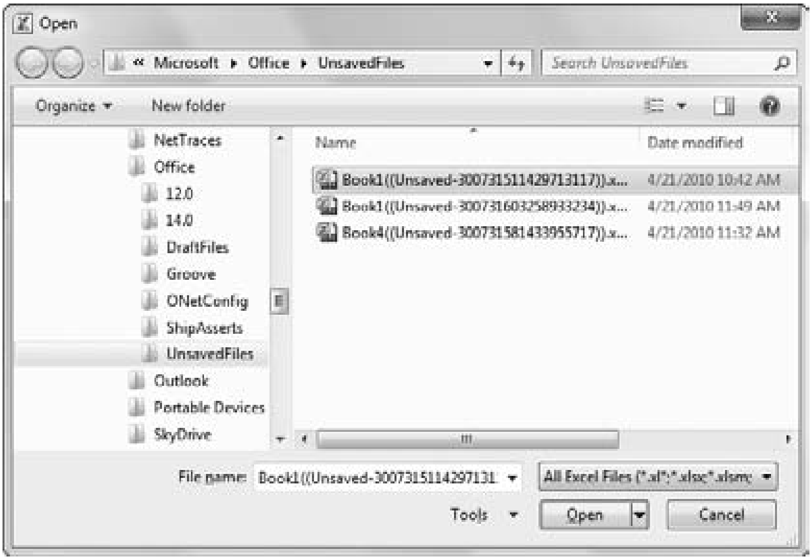

Figure 2-14: *Here, Excel is currently storing three unsaved workbooks. The file name is the name Excel originally gave the new document (like Book1), followed by a unique string of numbers. To open one, select it and click Open.*

You shouldn't depend on Excel to keep copies of unsaved work. After all, the AutoRecover file is a temporary backup, and Excel can delete it at any time. Instead, the AutoRecover is an emergency measure that's there to bail you out if you make a serious blunder. For example, almost every Excel guru has a horror story that usually runs like this: "It was 3:00 a.m., and I had just finished a long bout of number crunching. I had eight Excel windows opened at once, so I could scratch down different calculations as I worked. Relieved that my job was done, I saved my file, closed the program, and hit Don't Save a few times to get rid of my temporary work—only to realize there was another new spreadsheet with the company financials, and I'd just banished it into oblivion." Fortunately, you won't face this nightmare. If you act fast, you'll find that your abandoned file is still available in the recent files list.

Note: Excel files usually aren't that big, so you don't need to worry about drafts cluttering up your computer's hard drive. However, if you want to clear them out anyway, choose File→Info, click the big Versions button, and choose Delete All Draft Versions.

Incidentally, Excel stores its AutoRecover backups in a specific location, based on the person currently logged on to the computer. For example, if you're logged in as cindy_k in Windows Vista or Windows 7, your AutoRecover backups are stored in a folder like *C:\Users\cindy_k\AppData\Roaming\Microsoft\Excel*. (Excel creates a separate folder there for each workbook, so each workbook can keep its own carefully separated set of backups.) Drafts—the AutoRecover files for workbooks that have never been saved—are stored in a folder like *C:\Users\cindy_k\AppData\Roaming\Microsoft\DraftFiles*. You can change this location and other AutoRecover settings by choosing File→Options, and then choosing the Save category. Page 54 has a detailed description of AutoRecover settings.

Spell Check

A spell checker in Excel? Is that supposed to be for people who can't spell 138 correctly? The fact is that more and more people are cramming text—column headers, boxes of commentary, lists of favorite cereal combinations—into their spreadsheets. And Excel's designers have graciously responded by providing the very same spell checker that you've probably used with Microsoft Word. As you might expect, Excel's spell checker examines only text as it sniffs its way through a spreadsheet.

To start the spell checker, follow these simple steps:

1. **Move to where you want to start the spell check.**

   If you want to check the entire worksheet from start to finish, move to the first cell. Otherwise, move to the location where you want to start checking. Or, if you want to check a portion of the worksheet, select the cells you want to check.

   Excel's spell check can check only one worksheet at a time.

2. **Choose Review→Proofing→Spelling, or press F7.**

   The Excel spell checker starts working immediately, starting with the current cell and moving to the right, going from column to column. After it finishes the last column of the current row, checking continues with the first column of the next row.

   If you don't start at the first cell (A1) in your worksheet, Excel asks you when it reaches the end of the worksheet whether it should continue checking from the beginning of the sheet. If you say yes, it checks the remaining cells and stops when it reaches your starting point (having made a complete pass through all of your cells).

When the spell check finishes, a dialog box informs you that all cells have been checked. If your cells pass the spell check, this dialog box is the only feedback you receive. On the other hand, if Excel discovers any potential spelling errors during its check, it displays a Spelling window, as shown in Figure 2-15, showing the offending word and a list of suggestions.

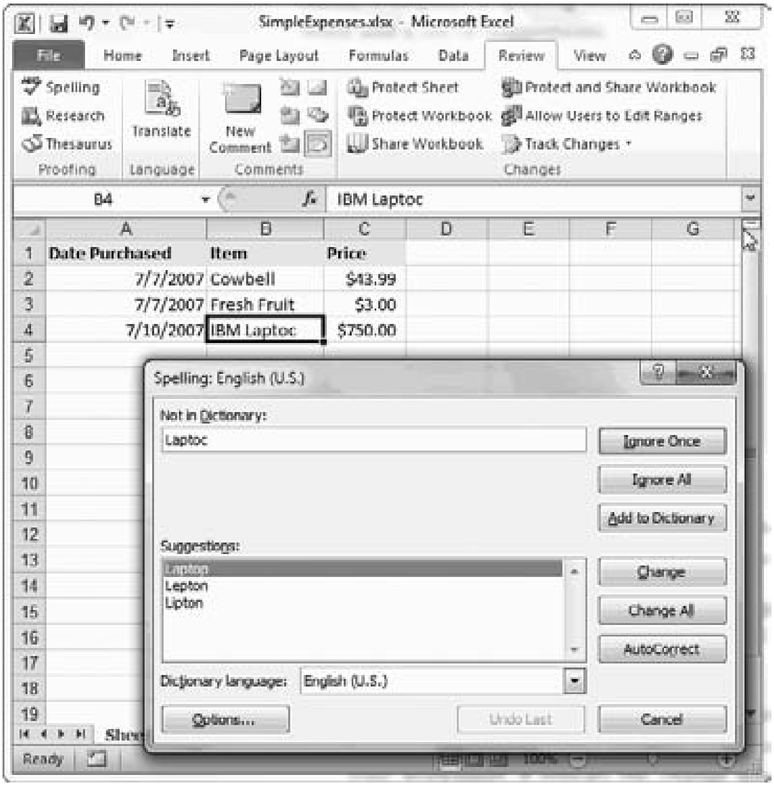

Figure 2-15: *When Excel encounters a word it thinks is misspelled, it displays the Spelling window. The cell containing the word—but not the actual word itself—gets highlighted with a black border. Excel doesn't let you edit your spreadsheet while the Spelling window is active. However, if you don't see the word you want in the list of suggestions, you can type a replacement in the "Not in Dictionary" box and click Change.*

The Spelling window offers a wide range of choices. If you want to use the list of suggestions to perform a correction, you have three options:

- Click one of the words in the list of suggestions, and then click Change to replace your text with the proper spelling. Double-clicking the word has the same effect.
- Click one of the words in the list of suggestions, and click Change All to replace your text with the proper spelling. If Excel finds the same mistake elsewhere in your worksheet, it repeats the change automatically.
- Click one of the words in the list of suggestions, and click AutoCorrect. Excel makes the change for this cell and for any other similarly misspelled words. In addition, Excel adds the correction to its AutoCorrect list (described on page 73). That means if you type the same unrecognized word into another cell (or even another workbook), Excel automatically corrects your entry. This option is useful if you've discovered a mistake that you frequently make.

Tip: If Excel spots an error but it doesn't give you the correct spelling in its list of suggestions, just type the correction into the "Not in Dictionary" box and hit Enter. Excel then inserts your correction into the corresponding cell.

On the other hand, if Excel is warning you about a word that doesn't represent a mistake (like your company name or some specialized term), you can click one of the following buttons:

- **Ignore Once** skips the word and continues the spell check. If the same word appears elsewhere in your spreadsheet, Excel prompts you again to make a correction.
- **Ignore All** skips the current word and all other instances of that word throughout your spreadsheet. You might use Ignore All to force Excel to disregard something you don't want to correct, like a person's name. The nice thing about Ignore All is that Excel doesn't prompt you again if it finds the same name, but it does prompt you again if it finds a different spelling (for example, if you misspelled the name).
- **Add to Dictionary** adds the word to Excel's custom dictionary. Adding a word is great if you plan to keep using a word that's not in Excel's dictionary. (For example, a company name makes a good addition to the custom dictionary.) Not only does Excel ignore any occurrences of this word, but if it finds a similar but slightly different variation of that word, it provides the custom word in its list of suggestions. Even better, Excel uses the custom dictionary in every workbook you spell check.
- **Cancel** stops the operation altogether. You can then correct the cell manually (or do nothing) and resume the spell check later.

Spell Checking Options

Excel lets you tweak how the spell checker works by letting you change a few basic options that control things like the language used and which, if any, custom dictionaries Excel examines. To set these options (or just to take a look at them), choose File→ Options, and then select the Proofing section (Figure 2-16). You can also reach these options by clicking the Spelling window's Options button while a spell check is underway.

The most important spell check setting is the language (at the bottom of the window), which determines what dictionary Excel uses. Depending on the version of Excel that you're using and the choices you made while installing the software, you might be using one or more languages during a spell check operation.

GEM IN THE ROUGH

Other Proofing Tools

Spreadsheet spell checking is a useful proofing tool. But Excel doesn't stop there. It piles on a few more questionable extras to help you enhance your workbooks.

Along with the spellchecker, Excel offers the following goodies. All of them require an Internet connection.

- **Research**. Click Review→Proofing→Research button to open a Research window, which appears on the right side of the Excel window, and lets you retrieve all kinds of information from the Web. The Research window provides a small set of Internet-driven services, including the ability to search a dictionary for a detailed definition, look in the Encarta encyclopedia, or get a delayed stock market quote from MSN Money. Page 690 has more.
- **Thesaurus**. Itching to promulgate your prodigious prolixity? (Translation: wanna use big words?) The thesaurus can help you take ordinary language and transform it into clear-as-mud jargon. Or, it can help you track down a synonym that's on the edge of your tongue. Either way, use this tool with care. To get started, click Review→Proofing→Thesaurus.
- **Translate**. Click this button to translate words or short phrases from one language to another. Behind the scenes, a free web service (like *www.worldlingo.com*) does the actual translation work. To try it out, click Review→Language→Translate.

Figure 2-16: *The spell checker options allow you to specify the language and a few other miscellaneous settings. This figure shows the standard settings that Excel uses when you first install it.*

Some of the other spelling options you can set include:

- **Ignore words in UPPERCASE**. If you choose this option, Excel won't bother checking any word written in all capitals (which is helpful when your text contains lots of acronyms).

- **Ignore words that contain numbers**. If you choose this option, Excel won't check words that contain numeric characters, like *Sales43* or *H3ll0*. If you don't choose this option, then Excel flags these entries as errors unless you've specifically added them to the custom dictionary.
- **Ignore Internet and file addresses**. If you choose this option, Excel ignores words that appear to be file paths (like *C:\Documents and Settings*) or website addresses (like *http://FreeSweatSocks.com*).
- **Flag repeated words**. If you choose this option, Excel treats the same words appearing consecutively ("the the") as an error.
- **Suggest from main dictionary only**. If you choose this option, the spell checker doesn't suggest words from the custom dictionary. However, it still *accepts* a word that matches one of the custom dictionary entries.

You can also choose the file Excel uses to store custom words—the unrecognized words that you add to the dictionary while a spell check is underway. Excel automatically creates a file named *custom.dic* for you to use, but you might want to use another file if you're sharing someone else's custom dictionary. (You can use more than one custom dictionary at a time. If you do, Excel combines them all to get one list of custom words.) Or, you might want to edit the list of words if you've mistakenly added something that shouldn't be there.

To perform any of these tasks, click the Custom Dictionaries button, which opens the Custom Dictionaries dialog box (Figure 2-17). From this dialog box, you can remove your custom dictionary, change it, or add a new one.

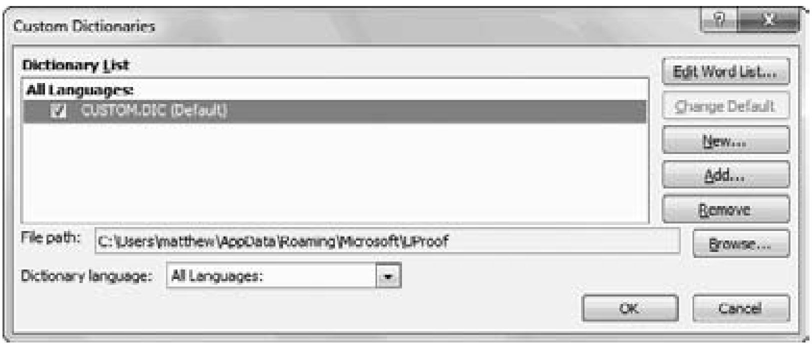

***Figure 2-17:** Excel starts you off with a custom dictionary named custom.dic (shown here). To add an existing custom dictionary, click Add and browse to the file. Or, click New to create a new, blank custom dictionary. You can also edit the list of words a dictionary contains (select it and click Edit Word List). Figure 2-18 shows an example of dictionary editing.*

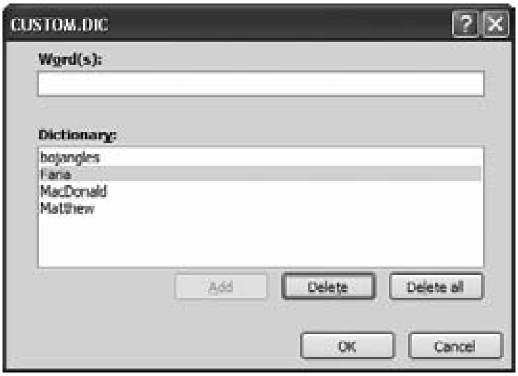

Figure 2-18:
This custom dictionary is fairly modest. It contains three names and an unusual word. Excel lists the words in alphabetical order. You can add a new word directly from this window (type in the text and click Add), remove one (select it and click Delete), or go nuclear and remove them all (click Delete All).

Note: All custom dictionaries are ordinary text files with the extension *.dic*. Unless you tell it otherwise, Excel assumes that custom dictionaries are located in a folder named *AppData\Roaming\UProof* inside the folder Windows uses for user-specific settings. For example, if you're using Windows Vista or Windows 7 and you're logged in under the user account Brad_Pitt, you'd find the custom dictionary in the *C:\Users\Brad_Pitt\AppData\Roaming\UProof* folder.

Adding Hyperlinks

Web browsers aren't the only programs that use *hyperlinks*—the underlined pieces of text that let you easily travel around the Web. In fact browsers weren't even the first. (Most believe that honor goes to the ambitious, 1960s-era cataloging project called Xanadu, which was never finished.) You may be surprised to find out that hyperlinks are quite useful in Excel, letting you link together different types of content and even navigate large spreadsheets. Here are three common examples:

- **You can create a hyperlink to a web page**. In this case, Excel opens your web browser in a new window and points it to the appropriate page.
- **You can create a hyperlink to a different type of file**. You can link to a Word document or a PowerPoint presentation, among other things. In this case, Excel opens whatever program is registered on your computer to handle this type of file. If you have a link to a .doc or .docx file and have Word installed, Excel opens a new Word window to display the document.
- **You can create a hyperlink to another worksheet or another part of the current spreadsheet**. This technique is helpful if you have a large amount of data and you want the people using your workbook to be able to quickly jump to the important places.

In Excel, you can place a maximum of one hyperlink in each cell.

Tip: Excel can create web page hyperlinks automatically. If you type some text that clearly corresponds to a web address (like text that starts with "http://" or "www."), Excel converts it to a hyperlink. When you're done typing, a smart tag appears, which you can click to undo this automatic adjustment and convert the cell back to ordinary text.

Creating a Link to a Web Page or Document

To insert a hyperlink into a cell, follow these steps:

1. **Move to the cell where you want to place the hyperlink.**
2. **Choose Insert→Links→Hyperlink (or press the shortcut key Ctrl+K).**

   The Insert Hyperlink dialog box appears, as shown in Figure 2-19.

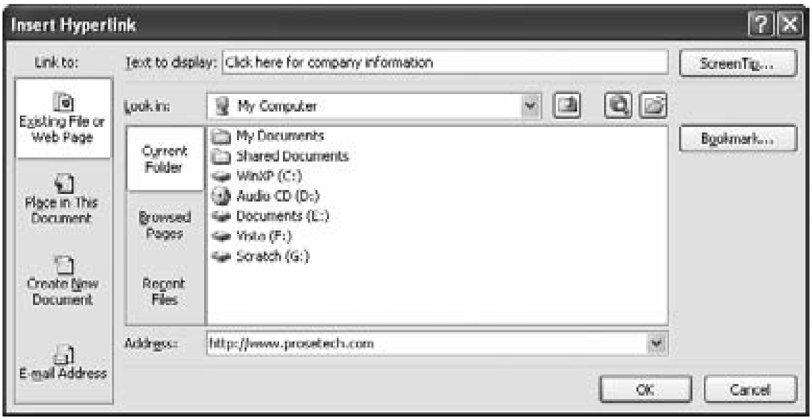

Figure 2-19: *In this example, someone's about to create a new hyperlink. It'll appear in the worksheet with the text "Click here for company information" (which, of course, you can edit to say anything you want) and will take the clicker to the website www.prosetech.com.*

Tip: You can also create a hyperlink on a picture object, so that the web page opens when you click the image. To do so, right-click the picture box, choose Hyperlink, and then continue with step 3. (For help inserting the picture in the first place, see Chapter 19.)

3. **Click the Existing File or Web Page option on the left side of the dialog box.**

   You can also use the Create New Document option to create and link to a new file in one step. The trick is that you need to remember to add the correct file extension. If you want to create a new Word document, you need to make sure you add .docx to the end of the file name so Windows knows what program to use for viewing and editing the file. If you want, select the "Edit this document now" option to open the file immediately in the appropriate program.

4. **At the top of the dialog box, in the "Text to display" text box, enter whatever you would like the link to say.**

   Common choices for the text include the actual web address (like *www.my-company.com*) or a descriptive message (like "Click here to go to my company's Website"). If the current cell already contains text, that text appears in the "Text to display" text box. If you change it, the new text replaces the current cell contents.

5. **If you want to set a custom tooltip for this hyperlink, click the ScreenTip button. Type in your message and click OK (see Figure 2-20).**

   A custom tooltip is a little message-bearing window that appears above a hyperlink when your mouse pointer hovers over the link. If you don't specify a custom tooltip, Excel shows the full path or URL.

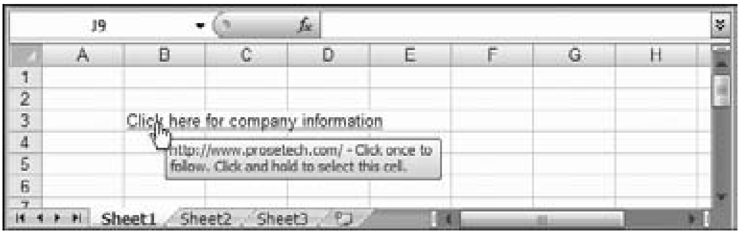

Figure 2-20: *A tooltip appears telling you the target web address, file path, or worksheet location. Optionally, you can replace this tooltip with a custom message when you create or edit the hyperlink.*

6. **If you want to add a link to a document, then browse to the appropriate file and select it. If you want to add a link to a web page, type the URL address into the Address text box.**

   If you're adding a link to a document, Excel sets the address to the full file path, as in *C:\MyDocuments\Resume.doc*. You can type this path in manually, and if your network supports it, you can use UNC (Universal Naming Convention) paths that point to a file on another computer using the name of the computer, as in *\\SalesComputer\Documents\CompanyPolicy.doc*.

Note: You're free to use files on your computer or those that are stored on network drives. Just remember that when you click the link, Excel looks in the exact location you've specified. That means if you move the target file to a new location, or you open the worksheet on another computer, Excel won't be able to find it.

7. **Click OK to insert the hyperlink.**

   When you insert a hyperlink, Excel formats the cell with blue lettering and adds an underline, so it looks like a web browser hyperlink. However, you can reformat the cell to change its appearance by selecting Home→Font section of the ribbon.

To use a hyperlink, just click it. You'll notice that the mouse pointer changes to a pointing hand as soon as you move over the hyperlink. If you want to move to a cell that contains a hyperlink without activating it, you can use the arrow keys, or you can click *and hold* the cell for about one second.

Creating a Link to a Worksheet Location

Hyperlinks also make for helpful navigational aids. If you have a worksheet with multiple tables of data, you can use a hyperlink to jump to a specific cell. If you create a spreadsheet that splits its data over multiple worksheets (a trick you'll pick up in Chapter 4), you can use a hyperlink to jump from worksheet to worksheet.

To create a hyperlink that uses a worksheet location as its target, follow these steps:

1. **Make note of the location you want to use for your target.**

   A worksheet hyperlink points to a specific cell.

   If you're a bit more advanced, you can create a link that moves to a new worksheet. For example, the reference Sheet2!A1 moves to cell A1 in Sheet2. (To learn more about worksheets and how to create them, see Chapter 4.)

   If you're feeling really slick, you can create a link that points to a *named cell reference*—a descriptive cell shortcut that you create and name. (To learn how to create a named cell reference, see page 371.)

2. **Move to the cell where you want to place the hyperlink.**

3. **Choose Insert→Links→Hyperlink (or press the shortcut key Ctrl+K).**

   The Insert Hyperlink dialog box appears.

4. **Click the "Place in This Document" option on the left side of the dialog box.**

   Excel displays a tree that represents the layout of the current workbook. When you link to another location in the workbook, you need to supply a cell reference (see Figure 2-21).

5. **If you want to jump to another worksheet, select the worksheet from the list.**

   All the worksheet names in your workbook appear under the Cell Reference heading.

6. **Type in the cell reference in the "Type the cell reference" text box.**

   Excel jumps to this location when somebody clicks the link.

7. **At the top of the dialog box, in the "Text to display" text box, enter whatever you would like the link to say.**

   Dealer's choice here.

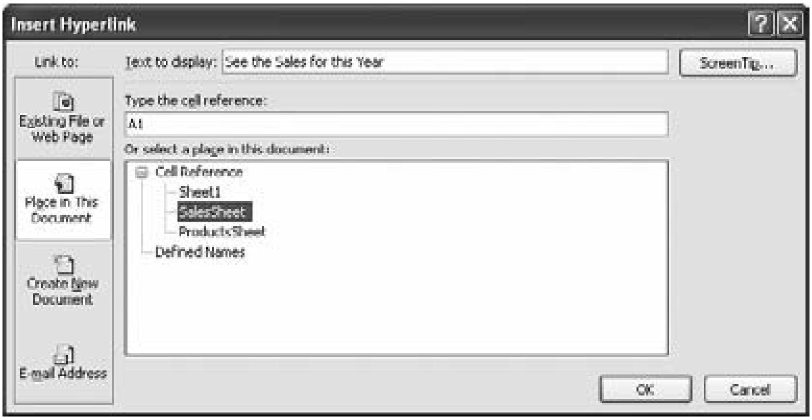

Figure 2-21: *Excel gives you a convenient tree that represents your workbook. You can choose any of the worksheets in your workbook, or you can choose a predefined named range (page 371).*

8. **If you want to set a custom tooltip for this hyperlink, then click the ScreenTip button. Type in your message and click OK.**

9. **Click OK to insert the hyperlink.**

   Admire your handiwork.

Tip: To edit a hyperlink, just move to the cell and choose Insert→Links→Hyperlink again (or press Ctrl+K). The same window appears, although now it has the title Edit Hyperlink.

CHAPTER
3

Moving Data

Simple spreadsheets are a good way to get a handle on Excel. But in the real world, you often need a spreadsheet that's more sophisticated—one that can grow and change as you start to track more information. For example, on the expenses worksheet you created in Chapter 1, perhaps you'd like to add information about which stores you've been shopping in. Or maybe you'd like to swap the order in which your columns appear. To make changes like these, you need to add a few more skills to your Excel repertoire.

This chapter covers the basics of spreadsheet modification, including how to select cells, how to move data from one place to another, and how to change the structure of your worksheet. What you learn here will make you a master of spreadsheet manipulation.

Selecting Cells

First things first: Before you can make any changes to an existing worksheet, you need to select the cells you want to modify. Happily, selecting cells in Excel—try saying that five times fast—is easy. You can do it many different ways, and it's worth learning them all. Different selection techniques come in handy in different situations, and if you master all of them in conjunction with the formatting features described in Chapter 5, you'll be able to transform the look of any worksheet in seconds.

Making Continuous Range Selections

Simplest of all is selecting a *continuous range* of cells. A continuous range is a block of cells that has the shape of a rectangle (high-school math reminder: a square is a kind of rectangle), as shown in Figure 3-1. The easiest way to select a continuous range is to click the top-left cell you want to select. Then drag to the right (to select more columns) or down (to select more rows). As you go, Excel highlights the selected cells in blue. Once you've highlighted all the cells you want, release the mouse button. Now you can perform an action, like copying the cells' contents, formatting the cells, or pasting new values into the selected cells.

Figure 3-1:
Top: The three selected cells (A1, B1, and C1) cover the column titles.

Bottom: This selection covers the nine cells that make up the rest of the worksheet. Notice that Excel doesn't highlight the first cell you select. In fact, Excel knows you've selected it (as you can see by the thick black border that surrounds it), but it has a white background to indicate that it's the active cell. When you start typing, Excel inserts your text in this cell.

In the simple expense worksheet from Chapter 1, for example, you could first select the cells in the top row and then apply bold formatting to make the column titles stand out. (Once you've selected the top three cells, press Ctrl+B, or chose Home→Font→Bold.)

Note: When you select some cells and then press an arrow key or click into another cell *before* you perform any action, Excel clears your selection.

You have a few useful shortcuts for making continuous range selections (some of these are illustrated in Figure 3-2):

- Instead of clicking and dragging to select a range, you can use a two-step technique. First, click the top-left cell. Then hold down Shift and click the cell at the bottom-right corner of the area you want to select. Excel highlights all the cells in between automatically. This technique works even if both cells aren't visible at the same time; just scroll to the second cell using the scroll bars, and make sure you don't click any other cell on your way there.
- If you want to select an entire column, click the header at the top of the column (as shown in Figure 3-2). For example, if you want to select the second column, then click the gray "B" box above the column. Excel selects all the cells in this column, right down to row 1,048,576.
- If you want to select an entire row, click the numbered row header on the left edge of the row. For example, you can select the second row by clicking the gray "2" box to the left of the row. All the columns in this row will be selected.

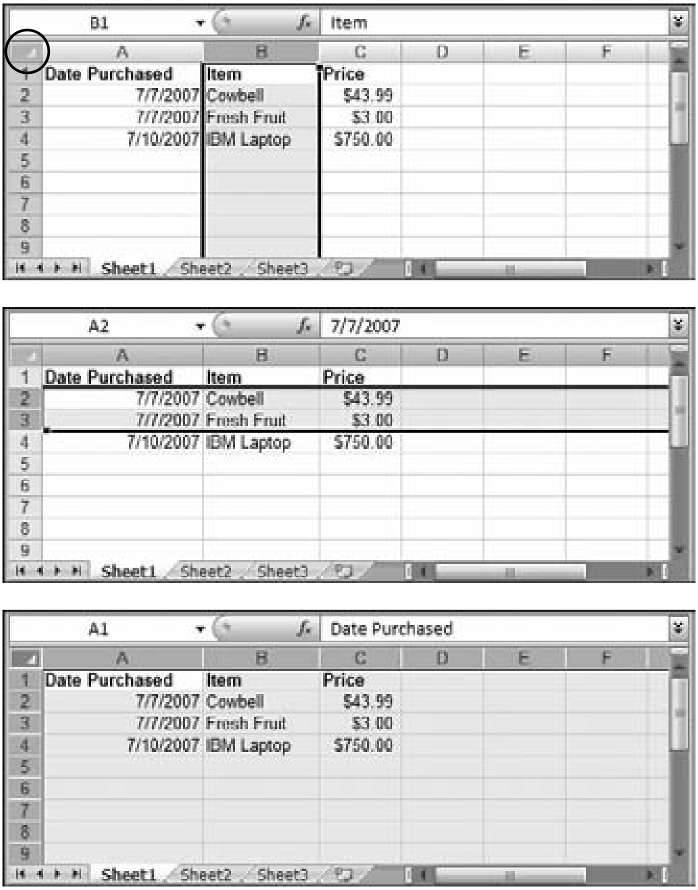

Figure 3-2:
Top: Click a column header to select that entire column.

Middle: Click a row number to select that entire row.

Bottom: To select every cell in the worksheet, click the gray box near the top-left corner of the worksheet (circled).

- If you want to select multiple adjacent columns, click the leftmost column header and then drag to the right until all the columns you want are selected. As you drag, a tooltip appears indicating how many columns you've selected. For example, if you've selected three columns, you'll see a tooltip with the text 3C (C stands for "column").
- If you want to select multiple adjacent rows, click the topmost row header and then drag down until all the rows you want are selected. As you drag, a tooltip appears indicating how many rows you've selected. For example, if you've selected two rows, you'll see a tooltip with the text 2R (R stands for "row").
- If you want to select all the cells in the entire worksheet, click the blank gray box that's just outside the top-left corner of the worksheet. This box is immediately to the left of the column headers and just above the row headers.

Tip: When you're selecting multiple rows or columns, make sure you click *between* the column header's left and right edges, not on either edge. When you click the edge of the column header, you end up resizing the column instead of making a selection.

TIME-SAVING TIP

A Truly Great Calculation Trick

Excel provides a seriously nifty calculation tool in the status bar. Just select two or more cells and look down to the status bar, where you'll see the number of cells you've selected (the count) along with their sum and their average (shown in Figure 3-3).

To choose what calculations appear in the status bar, right-click anywhere on the status bar and then, in the menu that appears, choose one of the following options:

- **Average**. The average of all the selected numbers or dates.
- **Count**. The number of selected cells that aren't blank.
- **Numerical Count**. The number of selected cells that contain numbers or dates.
- **Minimum**. The selected number or date with the smallest value (for dates this means the earliest date).
- **Maximum**. The selected number or date with the largest value (for dates this means the latest date).
- **Sum**. The sum of all selected numbers. Although you can use Sum with date values, because of the way Excel stores date values, adding dates together generates meaningless results.

Most of the status bar calculations don't work properly if you select both date and numeric information. For example, when you're attempting to add up a list of numbers and dates, Excel computes the value using both date values—which it stores internally as numbers, as explained on page 70—and the ordinary numbers; Excel then displays the final count using the formatting of the first selected cell. That adds up, alas, to a number that doesn't really mean anything.

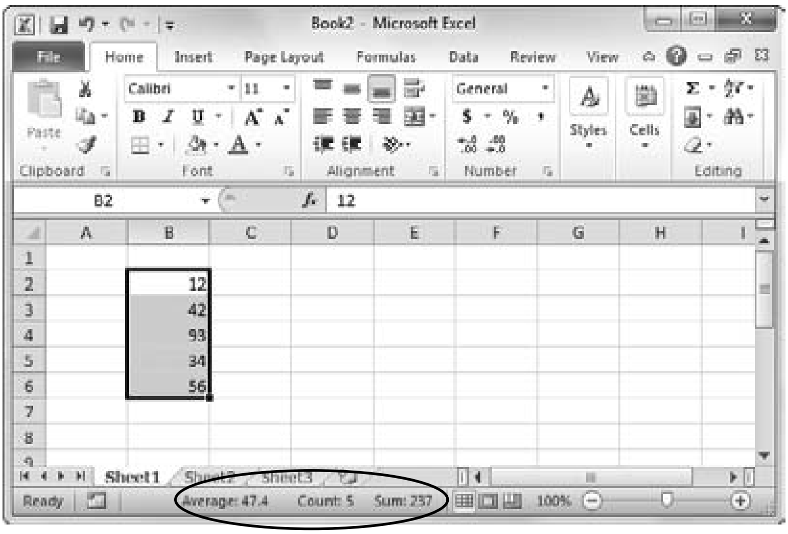

Figure 3-3: *The nicest detail about the status bar's quick calculations is that you can mix and match several at a time. Here, you see the count, average, and sum of the selected cells.*

Making Noncontiguous Selections

In some cases, you may want to select cells that are *noncontiguous* (also known as nonadjacent), which means they don't form a neat rectangle. For example, you might want to select columns A and C, but not column B. Or you might want to select a handful of cells scattered throughout the worksheet.

The trick to noncontiguous cell selection is using the Ctrl key. All you need to do is select the cells you want while holding down Ctrl. You can select individual cells by Ctrl-clicking them, or you can select multiple blocks of cells on different parts of the sheet by clicking and dragging in several different places while holding down Ctrl. You can also combine the Ctrl key with any of the shortcuts discussed earlier to select entire columns or rows as a part of your selection. Excel highlights in blue the cells you select (except for the last cell selected, which, as shown in Figure 3-4, isn't highlighted because it becomes the active cell).

Figure 3-4: *This figure shows a noncontiguous selection that includes four cells (A1, B2, C3, and B4). The last selected cell (B4) isn't highlighted because it's the active cell. This behavior is a little bit different from a continuous selection, in which the first selected cell is always the active cell. With a noncontiguous selection, the last selected cell becomes the active cell. Either way, the active cell is still a part of the selection.*

Automatically Selecting Your Data

Excel provides a nifty shortcut that can help you select a series of cells without dragging or Shift-clicking anything. It's called AutoSelect, and its special power is to select all the data values in a given row or column until it encounters an empty cell.

To use AutoSelect, follow these steps:

1. **Move to the first cell that you want to select.**

   Before continuing, decide which direction you want to extend the selection.

2. **Hold down Shift. Double-click whichever edge of the active cell corresponds to the direction you want to AutoSelect.**

   For example, if you want to select the cells below the active cell, then double-click its bottom edge. (You'll know you're in the right place when the mouse pointer changes to a four-way arrow.)

3. **Excel completes your selection automatically.**

   AutoSelection selects every cell in the direction you choose until it reaches the first blank cell. The blank cell (and any cells beyond it) won't be selected.

Making Selections with the Keyboard

The mouse can be an intuitive way to navigate around a worksheet and select cells. It can also be a tremendous time-suck, especially for nimble-fingered typists who've grown fond of the keyboard shortcuts that let them speed through actions in other programs.

Fortunately, Excel is ready to let you use the keyboard to select cells in a worksheet. Just follow these steps:

1. **Start by moving to the first cell you want to select.**

   Whichever cell you begin on becomes the anchor point from which your selected area grows. Think of this cell as the corner of a rectangle you're about to draw.

2. **Hold down Shift and move to the right or left (to select columns) and down or up (to select more rows), using the arrow keys.**

   As you move, the selection grows. Instead of holding down Shift, you can also just press F8 once, which turns on extend mode. When extend mode is on, you'll see the text Extend Selection in the status bar. As you move, Excel selects cells just as though you were holding down Shift. You can turn off extend mode by pressing F8 once you've finished marking your range.

Tip: If you really want to perform some selection magic, you can throw in one of Excel's powerful keyboard shortcuts. Use Ctrl+Space to select an entire column or Shift+Space to select the entire row. Or use the remarkable Ctrl+Shift+Space, which selects a block that includes the current cell and all the nearby contiguous cells (stopping only at the edges where it finds a blank cell). Finally, you can hit Ctrl+Shift+Space twice in a row to select the entire worksheet.

Making a noncontiguous selection is almost as easy. The trick is you need to switch between extend mode and another mode called add mode. Just follow these steps:

1. **Move to the first cell you want to select.**

   You can add cells to a noncontiguous range one at a time or add multiple continuous ranges. Either way, you start with the first cell you want to select.

2. **Press F8.**

   This key turns on extend mode. You'll see the text Extend Selection appear in the Status bar to let you know extend mode is turned on.

3. **If you want to select more than one cell, use the arrow keys to extend your selection.**

   If you just want to select the currently active cell, do nothing; you're ready to go onto the next step. When you want to add a whole block of cells, you can mark out your selection now. Remember, at this point you're still selecting a continuous range. In the steps that follow, you can add several distinct continuous ranges to make a noncontiguous selection.

4. **Press Shift+F8 to add the highlighted cells to your noncontiguous range.**

   When you hit Shift+F8, you switch to add mode, and you see the text "Add to Selection" appear in the status bar.

5. **You now have two choices: You can repeat steps 1 to 4 to add more cells to your selection, or you can perform an action with the current selection, like applying new formatting.**

   You can repeat steps 1 to 4 as many times as necessary to add more groups of cells to your noncontiguous range. These new cells (either individuals or groups) don't need to be near each other or in any way connected to the other cells you've selected. If you change your mind and decide you don't want to do anything with your selection after all, press F8 twice—once to move back into extend mode and then again to return to normal mode. Now, the next time you press an arrow key, Excel releases the current selection.

POWER USERS' CLINIC

Selecting Cells with the Go To Feature

In Chapter 1 (on page 28), you learned how you could use the Go To feature to jump from one position in a cell to another. A little-known Excel secret also allows you to use the Go To feature to select a range of cells.

It works like this: Start off at the top-left cell of the range you want to select. Then, open the Go To window by selecting Home→Editing→Find & Select→Go To or pressing Ctrl+G. Type in the address of the bottom-right cell in the selection you want to highlight. Now, here's the secret: Hold down Shift when you click the OK button. This action tells Excel to select the range of cells as it moves to the new cell.

For example, if you start in cell A1 and use the Go To window to jump to B3, then you'll select a block of six cells: A1, A2, A3, B1, B2, and B3.

Tip: You can also use the keyboard to activate AutoSelect. Just hold down the Shift key and use one of the shortcut key combinations that automatically jumps over a range of cells. For example, when you hold down Shift and then press Ctrl+↓, you'll automatically jump to the last occupied cell in the current row with all the cells in between selected. For more information about the shortcut keys, refer to Table 1-1 on page 27.

Moving Cells Around

One of the most common reasons to select groups of cells on a worksheet is to copy or move them from one place to another. Excel is a champion of the basic cut-and-paste feature, and it also gives you worthwhile enhancements that let you do things like drag and drop blocks of cells and copy multiple selections to the clipboard at the same time.

Before you get started shuffling data from one place to another, here are a few points to keep in mind:

- Excel lets you cut or copy a single cell or a continuous range of cells. Ordinarily, when you cut or copy a cell, *everything* goes with it, including the data and the current formatting. However, Excel also has tricks that let you copy data without formatting (or even just the formatting). You'll learn about those on page 106.
- When you paste cells onto your worksheet, you have two basic choices. You can paste the cells into a new, blank area of the worksheet, or you can paste the cells in a place that already contains data. In this second case, Excel overwrites the existing cells with the new pasted data.
- Cutting and copying cells works almost exactly the same way. The only difference you'll see is that when you perform a cut-and-paste operation (as opposed to a copy-and-paste operation), Excel erases the source data once the operation is complete. However, Excel doesn't remove the source cells from the worksheet. Instead, it just leaves them empty. (Page 115 shows you what to do if you do want to remove or insert cells, not just the data they contain.)

A Simple Cut-and-Paste or Copy-and-Paste

Here's the basic procedure for any cut-and-paste or copy-and-paste operation:

1. **Select the cells you want to cut or copy.**

   You can use any of the tricks you learned in the previous section to highlight a continuous range of cells. (You can't cut and paste noncontiguous selections.)

   When you want to cut or copy only a single cell, just move to the cell—you don't actually need to select it.

2. **If you want to cut your selection, choose Home→Clipboard→Cut (or Ctrl+X). When you want to copy your selection, choose Home→Clipboard→Copy (or Ctrl+C).**

   Excel highlights your selection with a *marquee border* (Figure 3-5), so called because the border blinks like the twinkling lights around an old-style movie theater marquee. At the same time, the text "Select destination and press ENTER or choose Paste" appears in the Status bar (if it fits).

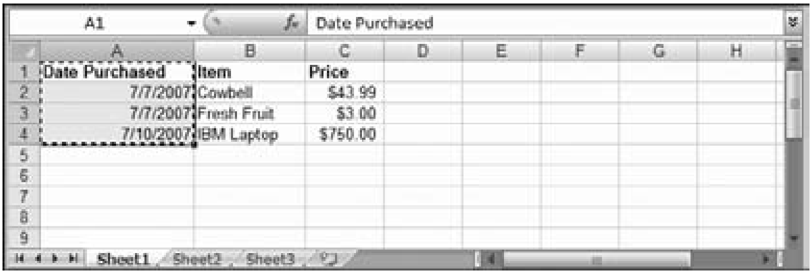

Figure 3-5: *In this example, cells A1 to A4 have been copied. The next step is to move to the place where you want to paste the cells and then press Enter to complete the operation. Excel treats cut and copy operations in the same way. In both cases, the selection remains on the spreadsheet, surrounded by the marquee border. When you perform a cut operation, Excel doesn't empty the cells until you paste them somewhere else.*

3. **Move to the new location in the spreadsheet where you want to paste the cells.**

   If you selected one cell, move to the new cell where you want to place the data. If you selected multiple cells, then move to the top-left corner of the area where you want to paste your selection. If you have existing data below or to the right of this cell, Excel overwrites it with the new content you're pasting.

   It's perfectly acceptable to paste over part of the data you're copying. For example, you could make a selection that consists of columns A, B, and C and paste that selection starting at column B. In this case, the pasted data appears in columns B, C, and D, and Excel overwrites the original content in these columns (although the original content remains in column A).

Tip: In some cases, you want to paste without overwriting part of your worksheet. For example, you might want to paste a column in a new position and shift everything else out of the way. To pull this trick off, you need the Insert Copied Cells command, which is described on page 114.

4. **Paste the data by selecting Home→Clipboard→Paste (or press Ctrl+V or Enter on the keyboard).**

   If you're performing a cut-and-paste, Excel removes the original data from the spreadsheet just before pasting it in the new location.

   If you're performing a copy-and-paste, a tiny clipboard icon appears in the bottom-right corner of your pasted cells, with the text "(Ctrl)" next to it. Click this icon and you'll get a menu of specialized paste options (described on page 107).

A Quicker Cut-and-Paste or Copy-and-Paste

If you want a really quick way to cut and paste data, you can use Excel's drag-and-drop feature. It works like this:

1. **Select the cells you want to move.**

   Just drag your pointer over the block of cells you want to select.

2. **Click the border of the selection box and don't release the mouse button.**

   You'll know that you're in the right place when the mouse pointer changes to a four-way arrow. You can click any edge, but *don't* click in the corner.

3. **Drag the selection box to its new location. If you want to copy (not move) the text, hold down the Ctrl key while you drag.**

   As you drag, a light-gray box shows you where Excel will paste the cells.

4. **Release the mouse button to move the cells.**

   If you drop the cells into a region that overlaps with other data, Excel prompts you to make sure that you want to overwrite the existing cells. This convenience isn't provided with ordinary cut-and-paste operations. (Excel uses it for drag-and-drop operations because it's all too easy to inadvertently drop your cells in the wrong place, especially while you're still getting used to this feature.)

Tip: Excel has a hidden dragging trick that impresses even the most seasoned users. To use it, follow the steps listed above but click on the border of the selection box with the *right* mouse button instead of the left. When you release the mouse button to finish the operation, a pop-up menu appears with a slew of options. Using this menu, you can perform a copy instead of a move, shift the existing cells out of the way, or use a special pasting option to copy values, formats, or links (page 107).

FREQUENTLY ASKED QUESTION

The Mysterious Number Signs

What does it mean when I see ####### in a cell?

A series of number signs is Excel's way of telling you that a column isn't wide enough to display the number or date that it contains (see Figure 3-6). Sometimes these signs appear when you're copying a big number into smaller cell.

The problem here is that Excel needs a certain amount of space to show your number. It's not acceptable to show just the first two digits from the number 412, for example, because that will look like the completely different number 41. However, Excel will trim off decimal places, if it can, which means that it will show 412.22344364 as 412.223 in a narrow column, while storing the full value with all the decimal places behind the scenes. But when the column is too narrow to show the whole number part of your value, or if you've set a required number of decimal places for your cell (page 141) and the column is too narrow to accommodate them, Excel shows the number signs to flag the problem.

Fortunately, it's easy to solve this problem—just position the mouse pointer at the right edge of the cell header and drag it to the right to enlarge the column. Provided you've made the column large enough, the missing number reappears. For a quicker solution, double-click the right edge of the column to automatically make it large enough.

This error doesn't usually occur while you're entering information for the first time because Excel automatically resizes columns to accommodate any numbers you type in. The problem is more likely to crop up if you shrink a column afterward, or if you cut some numeric cells from a wide column and paste them into a much narrower column. To verify the source of your problem, just move to the offending cell and then check the formula bar to see your complete number or date. Excel doesn't use the number signs with text cells—if those cells aren't large enough to hold their data, the words simply spill over to the adjacent cell (if it's blank) or become truncated (if the adjacent cell has some content).

There's one other situation that can cause a cell to display #######. If you create a formula that subtracts one time from another (as described in Chapter 11), and the result is a *negative* time value, you see the same series of number signs. But, in this case, column resizing doesn't help.

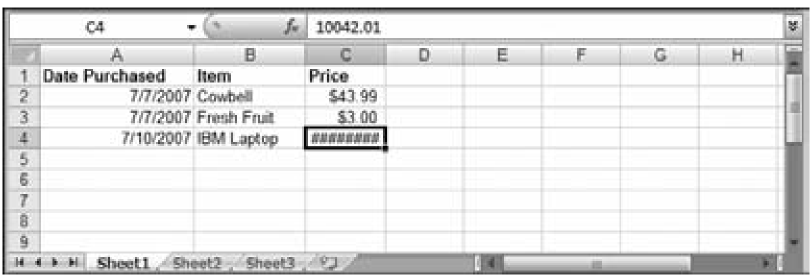

Figure 3-6: *Cell C4 has a wide number in an overly narrow column. You can see the mystery number if you move to the cell and check out the formula bar (it's 10,042.01), or you can expand the column to a more reasonable width.*

Fancy Pasting Tricks

When you copy cells, *everything* comes along for the ride, including text, numbers, and formatting. For example, if you copy a column that has one cell filled with bold text and several other cells filled with dollar amounts (including the dollar sign), when you paste this column into its new location, the numbers will still have the dollar sign and the text will still have bold formatting. If you want to change this behavior, you can use one of Excel's fancy paste options.

In the past, these options could be intimidatingly complex. But Excel 2010 improves life with a new *paste preview* feature. This allows you to preview what your pasted cells will look like before you've actually pasted them into your worksheet.

Here's how to try it out. First, copy your cells in the normal way. (Don't cut them, or the Paste Special feature won't work.) Then, move to where you want to paste the information, go to the Home→Clipboard section of the ribbon, and click the drop-down arrow at the bottom of the Paste button. You'll see a menu full of tiny pictures, each of which represents a different type of paste (see Figure 3-7).

Here's where things get interesting. When you hover over one of these pictures (but don't click it), the name of the paste option pops up, and Excel shows you a preview of what the pasted data will look in your worksheet. If you're happy with the result, click the picture to finish the paste. Or move your mouse over a different option to preview *its* results. And if you get cold feet, you can call the whole thing off by clicking any cell on the worksheet, in which case the preview will disappear and the worksheet will return to its previous state.

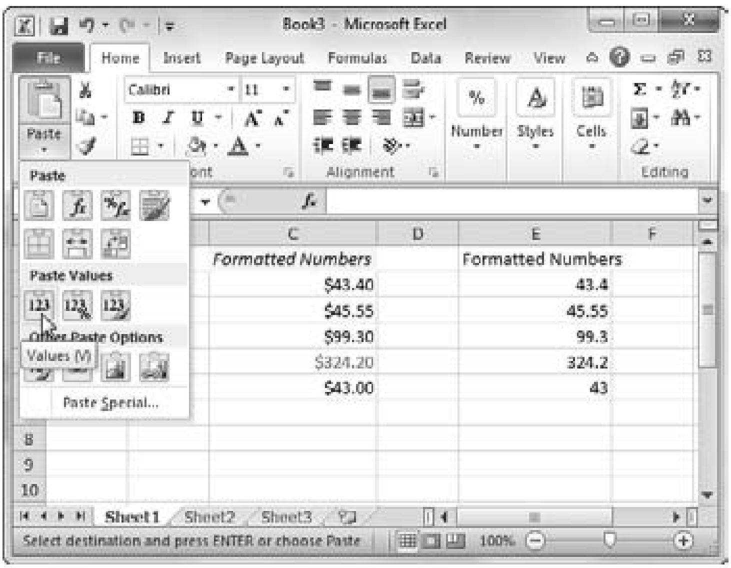

Figure 3-7: *In this example, the original data is in cells C1 to C6. The paste preview is shown in cells E1 to E6. Here, the preview is for the Values paste option, which copies all the numbers but none of the formatting. You can move to a different paste option and get a different preview or click a cell in the worksheet to banish the preview and cancel the paste operation.*

When you copy Excel cells (as opposed to data from another program), the list of paste options includes fourteen choices arranged into three groups. In the first group, named Paste, you'll find these choices:

- **Paste**. This option is the same as a normal paste operation, and it pastes both formatting and numbers.
- **Formulas**. This option pastes only cell content—numbers, dates, and text—without any formatting. If your source range includes any formulas, Excel also copies the formulas.
- **Formulas and Number Formatting**. This option is the same as Formulas, except it also copies the settings that control how numbers appear. (In other words, numbers that have currency signs, percentage signs, or thousands separators will remain formatted. However, you'll lose other formatting details, like fancy fonts, colors, and borders.)
- **Keep Source Formatting**. This option copies all the data and the formatting. In fact, it's the same as the ordinary Paste option, which is an unexplained Excel quirk.
- **No Borders**. This option copies all the data and formatting (just like an ordinary paste), except it ignores any borders that you've drawn around the cells. Page 167 describes to add borders between your cells.
- **Keep Source Column Widths**. This option copies all the data and formatting (just like an ordinary paste), but it also adjusts the columns in the paste region so that it has the same widths as the source columns.
- **Transpose**. This option inverts your information before it pastes it, so that all the columns become rows and the rows become columns. Figure 3-8 shows an example.

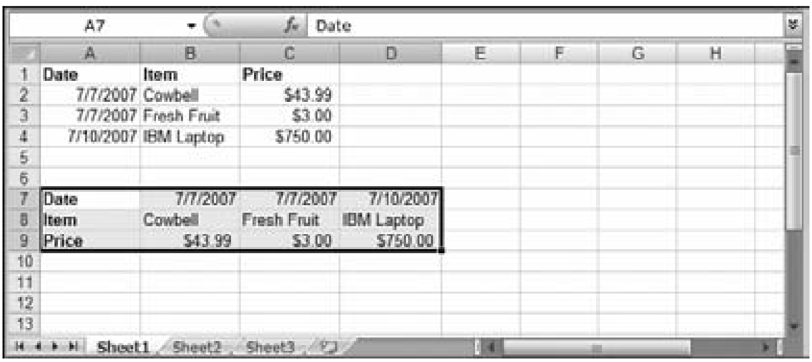

Figure 3-8: *With the Transpose option (from the Paste Special dialog box), Excel pastes the table at the top and transposes it on the bottom.*

The second group of paste options, called Paste Values, includes three choices:

- **Values**. This option pastes only cell content—numbers, dates, and text—without any formatting. If your source range includes any formulas, Excel pastes the *result* of those formulas (the calculated number) but not the actual formulas. (You'll learn everything you need to know about formulas in Chapter 8.)
- **Values and Number Formatting**. This option pastes the cell content and the formatting settings that control how numbers appear. If your source range includes any formulas, Excel pastes the calculated result of those formulas but not the actual formulas.
- **Values and Source Formatting**. This option is the same as a normal paste operation, except it doesn't copy formulas. Instead, it pastes the calculated result of any formulas.

The third group of paste options, called Other Paste Options, includes four choices that are little more specialized and a little less common:

- **Formatting**. This option applies the formatting from the source selection, but it doesn't actually copy any data.
- **Paste Link**. This option pastes a *link* in each cell that refers to the original data. (By comparison, an ordinary paste creates a duplicate *copy* of the source content.) If you use this option and then modify a value in one of the source cells, Excel automatically modifies the copy, too. (In fact, if you take a closer look at the copied cells in the formula bar, you'll find that they don't contain the actual data. Instead, they contain a formula that points to the source cell. For example, if you paste cell A2 as a link into cell B4, the cell B4 contains the reference =A2. You'll learn more about cell references and get to the bottom of this strange trick in Chapter 8.)
- **Picture.** This option pastes a *picture* of your cell, which is more than a little odd. The picture is placed right in the worksheet, with the formatting and borders you'd expect. In fact, if you don't look closely, this picture looks almost exactly like ordinary Excel data. The only way you'll know that it isn't is to click it. Unlike ordinary Excel data, you can't edit the data in a picture; instead, you're limited to resizing it, dragging it around your worksheet, and changing its borders. (You'll pick up these picture manipulation skills in Chapter 19.)

Note: Although it might make sense to copy a picture of your worksheet into other programs (a feat you'll master in Chapter 27), there's little reason to use the picture-pasting feature inside an Excel spreadsheet.

- **Linked Picture.** This option is the same as Picture, except Excel regenerates the picture whenever you modify the values of formatting in the source cells. This way, the picture always matches the source cells. Excel experts sometimes use this feature to create a summary that shows the important parts of a massive spreadsheet in one place. But in the wrong hands, this feature is a head-scratching trick that confuses everyone.

At the very bottom of the menu with paste options is a command named Paste Special. This brings up another window, with more esoteric pasting options. You'll take a peek in the next section.

Once you become familiar with the different paste options, you don't need to rely on the ribbon to use them. Instead, you can use them after a normal copy-and-paste. After you insert your data (by pressing Enter or using the Ctrl+V shortcut), look for the small paste icon that appears near the bottom-right corner of the pasted region. (Excel nerds know this icon as a *smart tag*.) If you click this icon (or press the Ctrl key), Excel pops open a menu (Figure 3-9) with the same set of paste options you saw earlier.

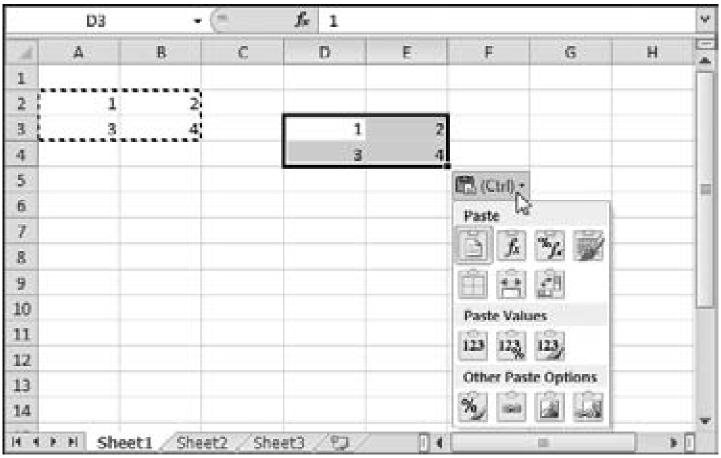

Figure 3-9: *The paste icon appears following the completion of every paste operation, letting you control a number of options, including whether the formatting matches the source or destination cells. You can change the type of paste several times, until you get exactly the result you want. The only disadvantage is that this menu doesn't have the same preview feature that the ribbon offers.*

Note: The paste icon appears only after a copy-and-paste operation, not a cut-and-paste operation. If you paste cells from the Clipboard panel, the paste icon still appears, but it provides just two options: keeping the source formatting or pasting the data only.

Paste Special

The paste options in the ribbon are practical and powerful. But Excel has even more paste options for those who have the need. To see them all, choose Home→Clipboard→Paste→Paste Special to pop open a dialog box with a slew of options (Figure 3-10).

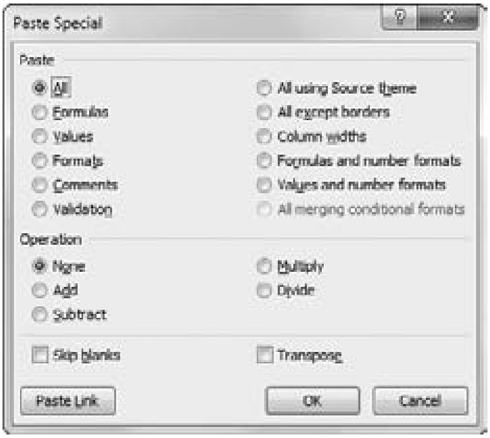

Figure 3-10:
The Paste Special window allows you to choose exactly what Excel will paste, and it also lets you apply a few other settings. The big drawback is that you don't get a preview, and some of the options are less than clear. In this example, Excel will perform an ordinary paste with a twist; it won't bother to copy any blank cells.

Paste Special is a bit of a holdover from the past. Many of its options are duplicated in the ribbon's drop-down Paste list. However, the Paste Special dialog box allows you to do a few things that aren't possible from the ribbon, including:

- **Paste comments**. Choose Comments in the Paste section, and click OK. This leaves all the text and formatting behind but copies any comments you've added to the cells. (You'll learn about comments on page 724.)
- **Paste validation**. Choose Validation in the Paste section and click OK. This leaves all the text and formatting behind but copies any validation settings that you've applied to the cells. (You'll learn about validation on page 694.)
- **Combine source and destination cells**. Choose All in the Paste section, choose Add, Subtract, Multiply, or Divide from the Operation section, and then click OK. For example, if you choose Subtract and paste the number 4 into a cell that currently has the number 6, Excel will change the cell to 2 (because *6–2=4*). It's an intriguing idea, but few people use the Operation settings, because they're not intuitive.
- **Refrain from copying blank cells.** Choose All in the Paste section, click the "Skip blanks" checkbox at the bottom of the dialog box and click OK. Now, if any of the cells you're copying are blank, Excel ignores them and leaves the current contents of the destination cell intact. (With an ordinary paste, Excel would overwrite the existing value, leaving a blank cell.)

Copying Multiple Items with the Clipboard

In Windows' early days, you could copy only a single piece of information at a time. If you copied two pieces of data, only the most recent item you copied would remain in the clipboard, a necessary way of life in the memory-starved computing days of yore. But nowadays, Excel boasts the ability to hold 24 separate cell selections in the Office clipboard. This information remains available as long as you have at least one Office application open.

Note: Even though the Office clipboard holds 24 pieces of information, you won't be able to access all this information in Windows applications that aren't part of the Office suite. If you want to paste Excel data into a non-Office application, you'll be able to paste only the data that was added to the clipboard most recently.

When you use the Home→Clipboard→Paste command (or Ctrl+V), you're using the ordinary Windows clipboard. That means you always paste the item most recently added to the clipboard. But if you fire up the Office clipboard, you can hold a lot more. Go to the Home→Clipboard section of the ribbon and then click the dialog box launcher (the small arrow-in-a-square icon in the bottom-right corner) to open the Clipboard panel. Now Excel adds all the information you copy to *both* the Windows clipboard and the more capacious Office clipboard. Each item that you copy appears in the Clipboard panel (Figure 3-11).

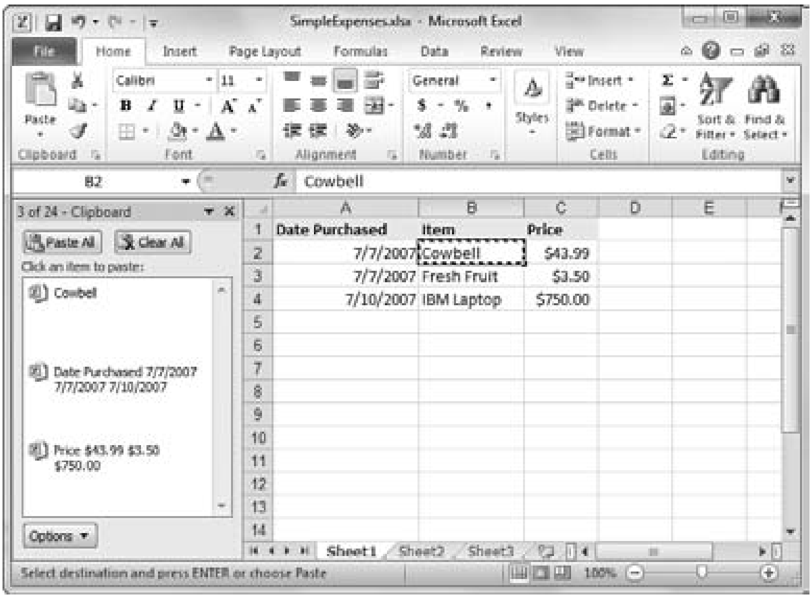

***Figure 3-11:** The Clipboard panel shows a list of all the items you've copied to it since you opened it (up to a limit of 24 items). Each item shows the combined content for all the cells in the selection. For example, the first item in this list includes four cells: the Price column title followed by the three prices. If you're using multiple Office applications at the same time, you may see scraps of Word documents, PowerPoint presentations, or pictures in the clipboard along with your Excel data. The icon next to the item always tells you which program the information came from.*

Using the Clipboard panel, you can perform the following actions:

- Click Paste All to paste all the selections into your worksheet. Excel pastes the first selection into the current cell and then begins pasting the next selection starting in the first row underneath that, and so on. As with all paste operations, the pasted cells overwrite any existing content in your worksheet.
- Click Clear All to remove all the selections from the clipboard. This is a useful approach if you want to add more data to the Clipboard but don't want to confuse this information with whatever selection you previously copied.
- Click a selection in the list to paste it into the current location in the worksheet.
- Click the drop-down arrow at the right of a selection item to show a menu that allows you to paste that item or remove it from the clipboard.

Depending on your settings, the Clipboard panel may automatically spring into action. To configure this behavior, click the Options button at the bottom of the Clipboard panel to display a menu of options. These include:

- **Show Office Clipboard Automatically**. If you turn on this option, the Clipboard panel automatically appears if you copy more than one piece of information to the clipboard. (Remember, without the Clipboard panel, you can access only the last piece of information you've copied.)
- **Show Office Clipboard When Ctrl+C Pressed Twice**. If you turn on this option, the Clipboard panel appears if you press the Ctrl+C shortcut twice in a row without doing anything else in between.
- **Collect Without Showing Office Clipboard**. If you turn on this option, it overrides the previous two settings, ensuring that the Clipboard panel never appears automatically. You can still call up the Clipboard panel manually, of course.
- **Show Office Clipboard Icon on Taskbar**. If you turn on this option, a clipboard icon appears in the system tray at the right of the taskbar. You can double-click this icon to show the Clipboard panel while working in any Office application. You can also right-click this icon to change clipboard settings or tell the Office clipboard to stop collecting data.
- **Show Status Near Taskbar When Copying**. If you turn on this option, you'll see a tooltip near the Windows system tray whenever you copy data. (The *system tray* is the set of notification icons at the bottom-right corner of your screen, in the Windows taskbar.) The icon for the Office clipboard shows a clipboard icon, and it displays a message like "4 of 24 -Item Collected" (which indicates you have just copied a fourth item to the clipboard).

UP TO SPEED

Cutting or Copying Part of a Cell

Excel's cut-and-paste and copy-and-paste features let you move data in one or more cells. But what if you simply want to take a snippet of text from a cell and transfer it to another cell or even another application? Excel makes this operation possible, but you need to work a little differently.

First, move to the cell that contains the content you want to cut or copy and then place it in edit mode by double-clicking it with the mouse, clicking in the formula bar, or pressing F2. You can now scroll through the cell content using the arrow keys. Move to the position where you want to start chopping or copying, hold down Shift and then arrow over to the right. Keep moving until you've selected all the text you want to cut or copy. Then, hit Ctrl+C to copy the text, or Ctrl+X to cut it. (When you cut text, it disappears immediately, just like in other Windows applications.) Hit Enter to exit edit mode once you're finished.

The final step is to paste your text somewhere else. You can move to another cell that has data in it already, press F2 to enter edit mode again, move to the correct position in that cell, and then press Ctrl+V. However, you can also paste the text directly into a cell by just moving to the cell and pressing Ctrl+V without placing it into edit mode. In this case, the data you paste overwrites the current content in the cell.

Adding and Moving Columns or Rows

The cut-and-paste and copy-and-paste operations let you move data from one cell (or group of cells) to another. But what happens if you want to make some *major* changes to your worksheet itself? For example, imagine you have a spreadsheet with 10 filled columns (A to J) and you decide you want to add a new column between columns C and D. You could cut all the columns from D to J and then paste them starting at E. That would solve the problem and leave the C column free for your new data. But the actual task of selecting these columns can be a little awkward, and it only becomes more difficult as your spreadsheet grows in size.

A much easier option is to use two dedicated Excel commands designed for inserting new columns and rows into an existing spreadsheet. If you use these features, you won't need to disturb your existing cells at all.

Inserting Columns and Rows

To insert a new column, follow these steps:

1. **Find the column immediately to the *right* of where you want to place the new column.**

   That means that if you want to insert a new, blank column between columns A and B, start by looking at the existing column B.

2. **Right-click the column header (the gray box with the column letter in it) and choose Insert.**

   Excel inserts a new column and automatically moves all the columns to the right of column A (so column B becomes column C, column C becomes column D, and so on).

Inserting Rows

Inserting rows is just as easy as inserting new columns. Just follow these steps:

1. **Find the row that's immediately *below* where you want to place the new row.**

   That means that if you want to insert a new, blank row between rows 6 and 7, start by looking at the existing row 7.

2. **Right-click on the row header (the numbered box at the far left of the row) and choose Insert.**

   Excel inserts a new row, and all the rows beneath it are automatically moved down one row.

Note: In the unlikely event that you have data at the extreme right edge of the spreadsheet, in column XFD, Excel doesn't let you insert a new column anywhere in the spreadsheet because the data would be pushed off into the region Beyond The Spreadsheet's Edges. Similarly, if you have data in the very last row (row 1,048,576), Excel doesn't let you insert more rows. If you do have data in either of these spots and try to insert a new column or row, Excel displays a warning message.

Inserting Copied or Cut Cells

Usually, inserting entirely new rows and columns is the most straightforward way to change the structure of your spreadsheet. You can then cut and paste new information into the blank rows or columns. However, in some cases, you may simply want to insert cells into an *existing* row or column.

To do so, begin by copying or cutting a cell or group of cells and then select the spot you want to paste into. Next, choose Home→Cells→Insert→Insert Copied Cells from the menu (or Home→Cells→Insert→Insert Cut Cells if you're performing a cut instead of a copy operation). Unlike the cut-and-paste feature, when you insert cells, you won't overwrite the existing data. Instead, Excel asks you whether the existing cells should be shifted down or to the right to make way for the new cells (as shown in Figure 3-12).

You need to be careful when you use the Insert Copied Cells feature. Because you're shifting only certain parts of your worksheet, it's possible to mangle your data, splitting the information that should be in one row or one column into multiple rows or columns! Fortunately, you can always back out of a tight spot using Undo (page 79). Figure 3-13 shows the sort of problem you could encounter.

Figure 3-12:
When you insert copied cells, Excel asks whether it should move the existing cells down or to the right.

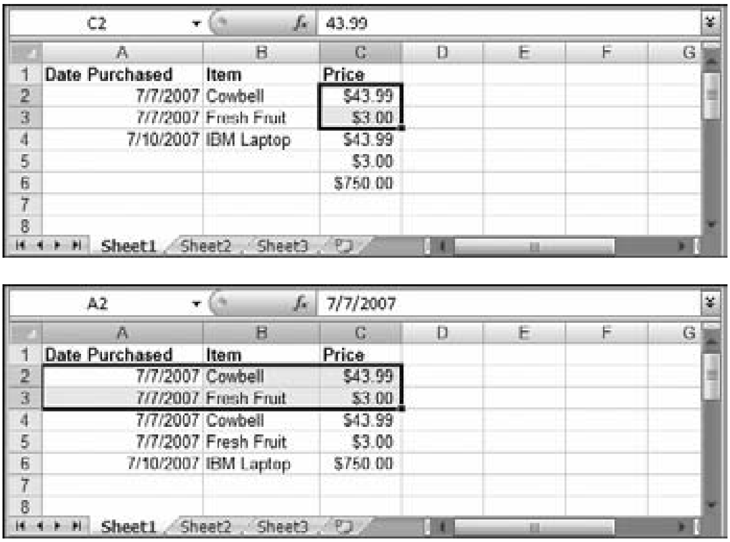

Figure 3-13:
Top: Here, two price cells ($43.99 and $3.00) were copied and pasted before the picture was taken, and the existing price cells were shifted down to accommodate the new entries. But the prices now no longer line up with the appropriate item names, which is probably not what you want.

Bottom: It makes much more sense to use the Insert Copied Cells command when you're copying a whole row's worth of data. Here's worksheet where two new rows have been pasted, and Excel politely moves the original set of items out of the way.

Deleting Columns and Rows

In Chapter 1, you learned that you can quickly remove cell values by moving to the cell and hitting the Delete key. You can also delete an entire range of values by selecting multiple cells and hitting the Delete key. Using this technique, you can quickly wipe out an entire row or column.

However, using delete simply clears the cell content. It doesn't remove the cells or change the structure of your worksheet. If you want to simultaneously clear cell values and adjust the rest of your spreadsheet to fill in the gap, you need to use the Home→Cell→Delete command.

For example, if you select a column by clicking the column header, you can either clear all the cells (by pressing the Delete key) or remove the column (by choosing Home→Cells→Delete). Deleting a column in this way is the reverse of inserting a column. All the columns to the right are automatically moved one column to the left to fill in the gap left by the column you removed. Thus, if you delete column B, column C becomes the new column B, column D becomes column C, and so on. If you take out row 3, row 4 moves up to fill the void, row 5 becomes row 4, and so on.

Usually, you'll use Home→Cells→Delete to remove entire rows or columns. However, you can also use it just to remove specific cells in a column or row. In this case, Excel prompts you with a dialog box asking whether you want to fill in the gap by moving cells in the current column up or by moving cells in the current row to the left. This feature is the reverse of the Insert Copied Cells feature, and you'll need to take special care to make sure you don't scramble the structure of your spreadsheet when you use this approach.

CHAPTER

4

Managing Worksheets

So far, you've learned how to create a basic worksheet with a table of data. That's great for getting started, but as power users, professional accountants, and other Excel jockeys quickly learn, some of the most compelling reasons to use Excel involve *multiple* tables that share information and interact with each other.

For example, say you want to track the performance of your company: you create one table summarizing your firm's yearly sales, another listing expenses, and a third analyzing profitability and making predictions for the coming year. If you create these tables in different spreadsheet files, you must copy shared information from one location to another, all without misplacing a number or making a mistake. And what's worse, with data scattered in multiple places, you're missing the chance to use some of Excel's niftiest charting and analytical tools. But cramming a bunch of tables onto the same worksheet page isn't the solution. Not only are you likely to lose your spot in the avalanche of data, you'll also face a host of formatting and cell management problems.

Fortunately, a better solution exists. Excel lets you create spreadsheets with multiple pages of data, each of which can conveniently exchange information with other pages. Each page is called a worksheet, and a collection of one or more worksheets is called a *workbook* (which is also sometimes called a *spreadsheet file*).

In this chapter, you'll learn how to manage the worksheets in a workbook. You'll also take a look at Find and Replace, an Excel tool for digging through worksheets in search of specific data.

Worksheets and Workbooks

Many workbooks contain more than one table of information. For example, you might have a list of your bank account balances and a list of items repossessed from your home in the same financial planning spreadsheet. You might find it a bit challenging to arrange these different tables. You could stack them (Figure 4-1) or place them side by side (Figure 4-2), but neither solution is perfect.

	A	B	C	D
1	2007 Purchases			
2	Date	Item	Price	
3	7/7/2007	Cowbell	$43.99	
4	7/7/2007	Fresh Fruit	$3.00	
5	7/10/2007	IBM Laptop	$750.00	
6				
7	2006 Purchases			
8	Date	Item	Price	
9	1/8/2006	Gorilla Suit	$140.25	
10	5/30/2006	Crayons	$2.99	
11	8/30/2006	Sofa	$599.00	
12				

Figure 4-1:
Stacking tables on top of each other is usually a bad idea. If you need to add more data to the first table, then you must move the second table. You'll also have trouble properly resizing or formatting columns because each column contains data from two different tables.

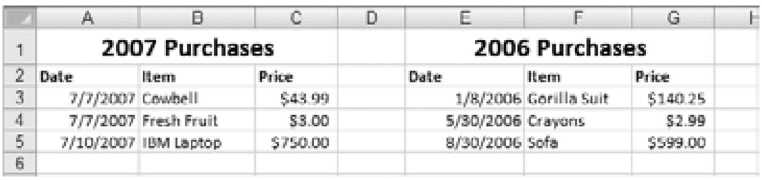

	A	B	C	D	E	F	G	H
1	2007 Purchases				2006 Purchases			
2	Date	Item	Price		Date	Item	Price	
3	7/7/2007	Cowbell	$43.99		1/8/2006	Gorilla Suit	$140.25	
4	7/7/2007	Fresh Fruit	$3.00		5/30/2006	Crayons	$2.99	
5	7/10/2007	IBM Laptop	$750.00		8/30/2006	Sofa	$599.00	
6								

Figure 4-2:
You're somewhat better off putting tables side by side, separated by a blank column, than you are stacking them, but this method can create problems if you need to add more columns to the first table. It also makes for a lot of side-to-side scrolling.

Most Excel masters agree that the best way to arrange different tables of information is to use separate worksheets for each table. When you create a new workbook, Excel automatically fills it with three blank worksheets named Sheet1, Sheet2, and Sheet3. Often, you'll work exclusively with the first worksheet (Sheet1), and not even realize that you have two more blank worksheets to play with—not to mention the ability to add plenty more.

To move from one worksheet to another, you have a few choices:

- Click the worksheet tabs at the bottom of Excel's grid window (just above the status bar), as shown in Figure 4-3.
- Press Ctrl+Page Down to move to the next worksheet. For example, if you're currently in Sheet1, this key sequence jumps you to Sheet2.

- Press Ctrl+Page Up to move to the previous worksheet. For example, if you're currently in Sheet2, this key sequence takes you back to Sheet1.

Figure 4-3: *Worksheets provide a good way to organize multiple tables of data. To move from one worksheet to another, click the appropriate Worksheet tab at the bottom of the grid. Each worksheet contains a fresh grid of cells—from A1 all the way to XFD1048576.*

Excel keeps track of the active cell in each worksheet. That means if you're in cell B9 in Sheet1, and then move to Sheet2, when you jump back to Sheet1 you'll automatically return to cell B9.

Tip: Excel includes some interesting viewing features that let you look at two different worksheets at the same time, even if these worksheets are in the same workbook. You'll learn more about custom views in Chapter 7.

Adding and Removing Worksheets

When you start a fresh workbook in Excel, you automatically get three blank worksheets in it. You can easily add more worksheets. Just click the Insert Worksheet button, which appears immediately to the right of your last worksheet tab (Figure 4-4). You can also use the Home→Cells→Insert→Insert Sheet command, which works the same way but inserts a new worksheet immediately to the *left* of the current worksheet. (Don't panic: page 122 shows how you can rearrange worksheets after the fact.)

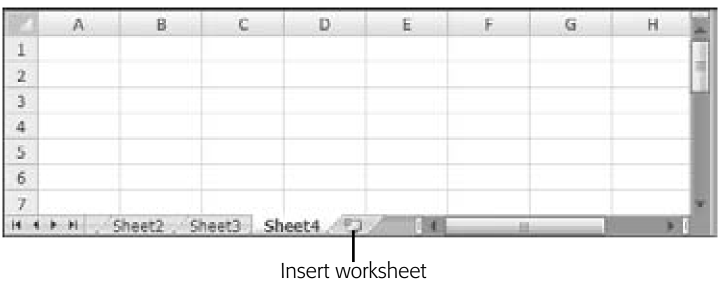

Figure 4-4: *Every time you click the Insert Worksheet button, Excel inserts a new worksheet after your existing worksheets and assigns it a new name. For example, if you start with the standard Sheet1, Sheet2, and Sheet3 and click the Insert Worksheet button, then Excel adds a new worksheet named—you guessed it—Sheet4.*

If you continue adding worksheets, you'll eventually find that all the worksheet tabs won't fit at the bottom of your workbook window. If you run out of space, you need to use the scroll buttons (which are immediately to the left of the worksheet tabs) to scroll through the list of worksheets. Figure 4-5 shows the scroll buttons.

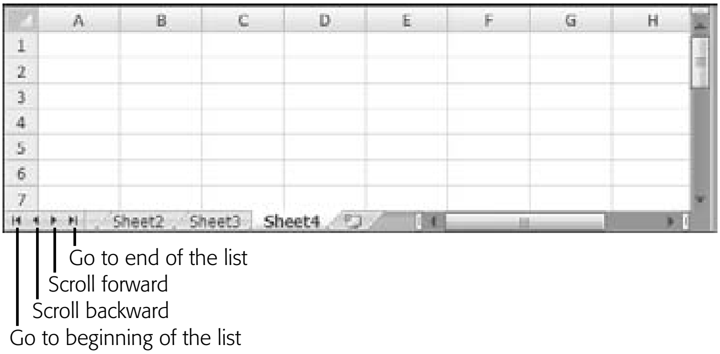

Figure 4-5:
Using the scroll buttons, you can move between worksheets one at a time or jump straight to the first or last tab. These scroll buttons control only which tabs you see–you still need to click the appropriate tab to move to the worksheet you want to work on.

Tip: If you have a huge number of worksheets and they don't all fit in the strip of worksheet tabs, there's an easier way to jump around. Right-click the scroll buttons to pop up a list with all your worksheets. You can then move to the worksheet you want by clicking it in the list.

Removing a worksheet is just as easy as adding one. Simply move to the worksheet you want to get rid of, and then choose Home→Cells→Delete→Delete Sheet (you can also right-click a worksheet tab, and then choose Delete). Excel won't complain if you ask it to remove a blank worksheet, but if you try to remove a sheet that contains any data, it presents a warning message asking for your confirmation. Also, if you're down to one last worksheet, Excel won't let you remove it. Doing so would create a tough existential dilemma for Excel—a workbook that holds no worksheets—so the program prevents you from taking this step.

Warning: Be careful when deleting worksheets, as you can't use Undo (Ctrl+Z) to reverse this change!

Excel starts you off with three worksheets for each workbook, but changing this setting's easy. You can configure Excel to start with fewer worksheets (as few as one), or many more (up to 255). Select File→Options, and then choose the General section. Under the heading "When creating new workbooks", change the number in the "Include this many sheets" box, and then click OK. This setting takes effect the next time you create a new workbook.

Note: Although you're limited to 255 sheets in a new workbook, Excel doesn't limit how many worksheets you can add *after* you've created a workbook. The only factor that ultimately limits the number of worksheets your workbook can hold is your computer's memory. However, modern-day PCs can easily handle even the most ridiculously large, worksheet-stuffed workbook.

Hiding Worksheets

Deleting worksheets isn't the only way to tidy up a workbook or get rid of information you don't want. You can also choose to *hide* a worksheet temporarily.

When you hide a worksheet, its tab disappears, but the worksheet itself remains part of your spreadsheet file, available whenever you choose to unhide it. Hidden worksheets also don't appear on printouts. To hide a worksheet, right-click the worksheet tab, and then choose Hide. (Or, for a more long-winded approach, choose Home→Cells→Format→Hide & Unhide→Hide Sheet.)

To redisplay a hidden worksheet, right-click any worksheet tab, and then choose Unhide. The Unhide dialog box appears along with a list of all hidden sheets, as shown in Figure 4-6. Select a sheet from the list, and then click OK to unhide it. (Once again, the ribbon can get you the same window—just point yourself to Home→Cells→Format→Hide & Unhide→Unhide Sheet.)

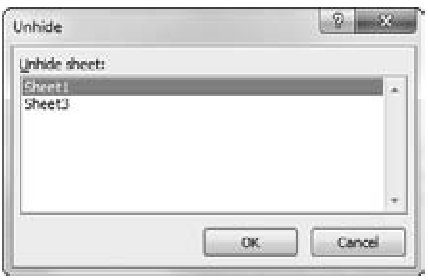

Figure 4-6:
This workbook contains two hidden worksheets. To restore one, just select it from the list, and then click OK. Unfortunately, if you want to show multiple hidden sheets, you must use the Unhide Sheet command multiple times. Excel has no shortcut for unhiding multiple sheets at once.

Naming and Rearranging Worksheets

The standard names Excel assigns to new worksheets—Sheet1, Sheet2, Sheet3, and so on—aren't very helpful for identifying what they contain. And they become even less helpful if you start adding new worksheets, since the new sheet numbers don't necessarily indicate the position of the sheets, just the order in which you created them.

For example, if you're on Sheet 3 and you add a new worksheet (by choosing Home→Cells→Insert→Insert Sheet), then the worksheet tabs read: Sheet1, Sheet2, Sheet4, Sheet3. (That's because the Insert Sheet command inserts the new sheet just before your current sheet.) Excel doesn't expect you to stick with these auto-generated names. Instead, you can rename them by right-clicking the worksheet tab and selecting Rename, or just double-click the sheet name. Either way, Excel highlights the worksheet tab, and you can type a new name directly onto the tab. Figure 4-7 shows worksheet tabs with better names.

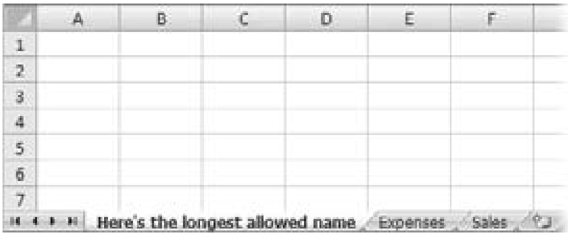

Figure 4-7:
Worksheet names can be up to 31 characters long and can include letters, numbers, some symbols, and spaces. Remember, though, the longer the worksheet name, the fewer worksheet tabs you'll be able to see at once, and the more you'll need to rely on the scroll buttons to the left of the worksheet tabs. For convenience's sake, try to keep your names brief by using titles like Sales04, Purchases, and Jet_Mileage.

Note: Excel has a small set of reserved names that you can never use. To witness this problem, try to create a worksheet named History. Excel doesn't let you, because it uses the History worksheet as part of its change tracking features (page 729). Use this Excel oddity to impress your friends.

Sometimes Excel refuses to insert new worksheets exactly where you'd like them. Fortunately, you can easily rearrange any of your worksheets just by dragging their tabs from one place to another, as shown in Figure 4-8.

Figure 4-8:
When you drag a worksheet tab, a tiny page appears beneath the arrow cursor. As you move the cursor around, you'll see a black triangle appear, indicating where the worksheet will land when you release the mouse button.

Tip: You can use a similar technique to create *copies* of a worksheet. Click the worksheet tab and begin dragging, just as you would to move the worksheet. However, before releasing the mouse button, press the Ctrl key (you'll see a plus sign [+] appear). When you let go, Excel creates a copy of the worksheet in the new location. The original worksheet remains in its original location. Excel gives the new worksheet a name with a number in parentheses. For example, a copy of Sheet1 is named Sheet1 (2). As with any other worksheet tab, you can change this name.

GEM IN THE ROUGH

Colorful Worksheet Tabs

Names aren't the only thing you can change when it comes to newly added worksheets. Excel also lets you modify a worksheet tab's background color. This minor convenience has no effect on your data or your printout, but it can help you quickly find an important worksheet if it has lots of neighbors.

To change the background color of a worksheet tab, right-click the tab, and then select Tab Color (or move to the appropriate worksheet and Home→Cells→Format→Tab Color). A list of color choices appears; make your selection by clicking the color you want.

Moving Worksheets from One Workbook to Another

Once you get the hang of creating different worksheets for different types of information, your Excel files can quickly fill up with more sheets than a linens store. What happens when you want to shift some of these worksheets around? For instance, you may want to move (or copy) a worksheet from one Excel file to another. Here's how:

1. **Open both spreadsheet files in Excel.**

   The file that contains the worksheet you want to move or copy is called the *source* file; the other file (where you want to move or copy the worksheet to) is known as the *destination* file.

2. **Go to the source workbook.**

   Remember, you can move from one window to another using the Windows task bar, or by choosing the file's name from the ribbon's View→Windows→Switch Windows list.

3. **Right-click the worksheet you want to transfer, and then, from the shortcut menu that appears, choose Move or Copy.**

   If you want, you can transfer multiple worksheets at once. Just hold down the Ctrl key, and then select all the worksheets you want to move or copy. Excel highlights all the worksheets you select (and groups them together). Right-click the selection, and then choose Move or Copy.

   When you choose Move or Copy, the Move or Copy dialog box appears (as shown in Figure 4-9).

4. **Choose the destination file from the "To book" list.**

   The "To book" drop-down list shows all the currently open workbooks (including the source workbook).

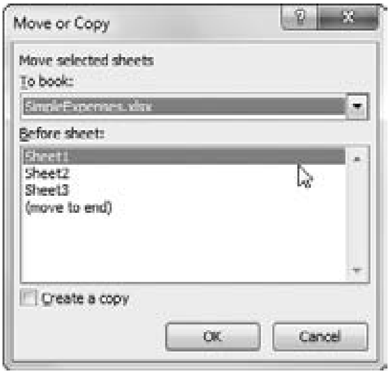

Figure 4-9:
Here, the selected worksheet is about to be moved into the SimpleExpenses.xlsx workbook. (The source workbook isn't shown.) The SimpleExpenses workbook already contains three worksheets (named Sheet1, Sheet2, and Sheet3). Excel inserts the new worksheet just before the first sheet. Because the "Create a copy" checkbox isn't turned on, Excel removes the worksheet from the source workbook when it completes the transfer.

Tip: Excel also lets you move your worksheets to a new workbook, which it automatically creates for you. To move them to a new workbook, choose the "(new book)" item in the "To book" list. The new workbook won't have the standard three worksheets. Instead, it'll have only the worksheets you've transferred.

5. **Specify the position where you want the worksheet inserted.**

   Choose a destination worksheet from the "Before sheet" list. Excel places the copied worksheets just *before* the worksheet you select. If you want to place the worksheets at the end of the destination workbook, select "(move to end)." Of course, you can always rearrange the worksheets after you transfer them, so you don't need to worry too much about getting the perfect placement.

6. **If you want to copy the worksheet, turn on the "Create a copy" checkbox at the bottom of the window.**

   If you don't turn this option on, then Excel copies the worksheet to the destination workbook and removes it from the source workbook. If you *do* turn this option on, you'll end up with a copy of the worksheet in both workbooks.

7. **Click OK.**

   This final step closes the Move or Copy dialog box and transfers the worksheet (or worksheets).

Note: If there are any worksheet name conflicts, Excel adds a number in parentheses after the moved sheet's name. For example, if you try to copy a worksheet named Sheet1 to a workbook that already has a Sheet1, Excel names the copied worksheet Sheet1 (2).

Grouping Worksheets

As you've seen in previous chapters, Excel lets you work with more than one column, row, or cell at a time. The same holds true for worksheets. You can select multiple worksheets and perform an operation on all of them at once. This process of selecting multiple sheets is called *grouping*, and it's helpful if you need to hide or format several worksheets (for example, if you want to make sure all your worksheets start with a bright yellow first row), and you don't want the hassle of selecting them one at a time. Grouping sheets doesn't let you do anything you couldn't do ordinarily—it's just a nifty timesaver.

Here are some operations—all of which are explained in detail below—that you can simultaneously perform on worksheets that are grouped together:

- Move, copy, delete, or hide the worksheets.
- Apply formatting to individual cells, columns, rows, or even entire worksheets.
- Enter new text, change text, or clear cells.
- Cut, copy, and paste cells.
- Adjust some page layout options, like paper orientation (on the Page Layout tab).
- Adjust some view options, like gridlines and the zoom level (on the View tab).

To group worksheets, hold down Ctrl while clicking multiple worksheet tabs. When you're finished making your selections, release the Ctrl key. Figure 4-10 shows an example.

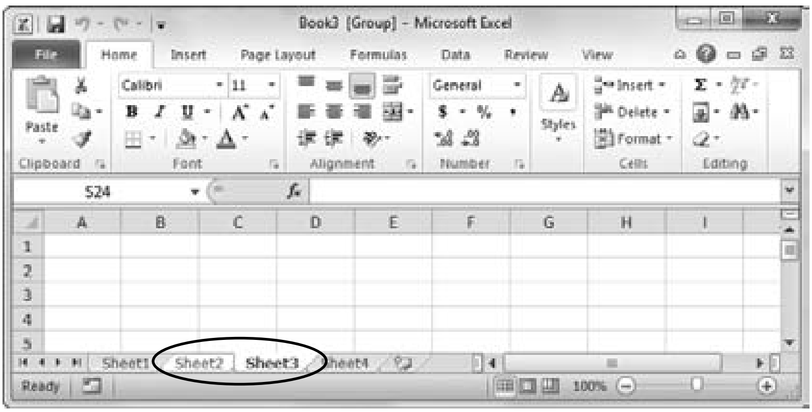

Figure 4-10: *In this example, Sheet2 and Sheet3 are grouped. When worksheets are grouped, their tab colors change from gray to white. Also, in workbooks with groups, the title bar of the Excel window includes the word [Group] at the end of the file name.*

Tip: As a shortcut, you can select all the worksheets in a workbook by right-clicking any tab, and then choosing Select All Sheets.

To ungroup worksheets, right-click one of the worksheet tabs and select Ungroup Sheets, or just click one of the worksheet tabs that isn't in your group. You can also remove a single worksheet from a group by clicking it while holding down Ctrl. However, this technique works only if the worksheet you want to remove from the group is *not* the currently active worksheet.

Managing grouped sheets

As your workbook grows, you'll often need better ways to manage the collection of worksheets you've accumulated. For example, you might want to temporarily hide a number of worksheets, or move a less important batch of worksheets from the front (that is, the left side) of the worksheet tab holder to the end (the right side). And if a workbook's got way too many worksheets, you might even want to relocate several worksheets to a brand new workbook.

It's easy to perform an action on a group of worksheets. For example, when you have a group of worksheets selected, you can drag them en masse from one location to another in the worksheet tab holder. To delete or hide a group of sheets, just right-click one of the worksheet tabs in your group, and then choose Delete or Hide. Excel then deletes or hides *all* the selected worksheets (provided that action will leave at least one visible worksheet in your workbook).

Note: Excel keeps track of print and display settings on a per-worksheet basis. In other words, when you set the zoom percentage to 50 percent in one worksheet so you can see more data, it doesn't affect the zoom in another worksheet. However, when you make the change for a *group* of worksheets, they're all affected in the same way.

Formatting grouped sheets

When you format cells inside *one* grouped worksheet, it triggers the same changes in the cells in the *other* grouped worksheets. So you have another tool you can use to apply consistent formatting over a batch of worksheets. It's mainly useful when your worksheets are all structured in the same way.

For example, imagine you've created a workbook with 10 worksheets, each one representing a different customer order. If you group all 10 worksheets together, and then format just the first one, Excel formats all the worksheets in exactly the same way. Or say you group Sheet1 and Sheet2, and then change the font of column B in Sheet2—Excel automatically changes the font in column B in Sheet1, too. The same is true if you change the formatting of individual cells or the entire worksheet—Excel replicates these changes across the group. (To change the font in the currently selected cells, just select the column and, in the Home→Font section of the ribbon, make a new font choice from the font list. You'll learn much more about the different types of formatting you can apply to cells in Chapter 5.)

Note: It doesn't matter which worksheet you modify in a group. For example, if Sheet1 and Sheet2 are grouped, you can modify the formatting of both in either worksheet. Excel automatically applies the changes to the other sheet.

Entering data in grouped sheets

With grouped worksheets, you can also modify the contents of individual cells, including entering or changing text and clearing cell contents. For example, if you enter a new value in cell B4 in Sheet2, Excel enters the same value into cell B4 in the grouped Sheet1. Even more interesting, if you modify a value in a cell in Sheet2, the same value appears in the same cell in Sheet1, even if Sheet1 didn't previously have a value in that cell. Similar behavior happens when you delete cells.

Editing a group of worksheets at once isn't as useful as moving and formatting them, but it does have its moments. Once again, it makes most sense when all the worksheets have the same structure. For example, you could use this technique to put the same copyright message in cell A1 on every worksheet, or to add the same column titles to multiple tables (assuming they're arranged in *exactly* the same way). One example where grouped sheets make sense is if you have a different worksheet for every month of a year, but each one has the same overall structure.

Warning: Be careful to remember the magnified power your keystrokes possess when you're operating on grouped worksheets. For example, imagine that you move to cell A3 on Sheet1, which happens to be empty. If you click Delete, you see no change. However, if cell A3 contains data on other worksheets that are grouped, these cells are now empty. Grouper beware.

Cut and paste operations work the same way as entering or modifying grouped cells. Whatever action you perform on one grouped sheet, Excel also performs on other grouped sheets. For example, consider what happens if you've grouped together Sheet1 and Sheet2, and you copy cell A1 to A2 in Sheet1. The same action takes place in Sheet2—in other words, Excel copies the contents of cell A1 (in Sheet2) to cell A2 (also in Sheet2). Obviously, the contents of cell A1 and A2 in Sheet1 may be different from the contents of cell A1 and A2 in Sheet2—the grouping simply means that whatever was in cell A1 will now also be in cell A2.

Find and Replace

When you're dealing with great mounds of information, you may have a tough time ferreting out the nuggets of data you need. Fortunately, Excel's find feature is great for helping you locate numbers or text, even when they're buried within massive workbooks holding dozens of worksheets. And if you need to make changes to a bunch of identical items, the find-and-replace option can be a real timesaver.

The Find and Replace feature includes both simple and advanced options. In its basic version, you're only a quick keystroke combo away from a word or number you *know* is lurking somewhere in your data pile. With the advanced options turned on, you can do things like search for cells that have certain formatting characteristics and apply changes automatically. The next few sections dissect these features.

The Basic Find

Excel's Find feature is a little like the Go To tool described in Chapter 1, which lets you move across a large expanse of cells in a single bound. The difference is that Go To moves to a *known* location, using the cell address you specify. The Find feature, on the other hand, searches every cell until it finds the content you've asked Excel to look for. Excel's search works similarly to the search feature in Microsoft Word, but it's worth keeping in mind a few additional details:

- Excel searches by comparing the content you enter with the content in each cell. For example, if you searched for the word *Date*, Excel identifies as a match a cell containing the phrase *Date Purchased*.
- When searching cells that contain numeric or date information, Excel always searches the cell *content*, not the display text. (For more information about the difference between the way Excel displays a numeric value and the underlying value Excel actually stores, see page 143.)

  For example, say a cell displays dates using the day-month-year format, like *2-Dec-10*. Internally, Excel stores the date as *12/2/2010*, which you'll see if you move to the cell and look in the formula bar. Thus, if you perform a search for *2010* or *12/2* you'll find the cell, because your search text matches part of the stored content. But if you search for *Dec* or *2-Dec-10*, you won't find a match. A similar behavior happens with numbers. For example, the search string *$3* won't match the currency value *$3.00*, because the dollar sign isn't part of the stored cell value—it's just a formatting detail. You can change this behavior and start searching the cell display text using the "Look in" setting described on page 132.
- Excel searches one cell at a time, from left to right. When it reaches the end of a row, it moves to the first column of the next row.

To perform a find operation, follow these steps:

1. **Move to the cell where you want the search to begin.**

   If you start halfway down the worksheet, for example, the search covers the cells from there to the end of the worksheet, and then "loops over" and starts at cell A1. If you select a group of cells, Excel restricts the search to just those cells. You can search across a set of columns, rows, or even a noncontiguous group of cells.

2. **Choose Home→Editing→Find & Select→Find, or press Ctrl+F.**

   The Find and Replace window appears, with the Find tab selected.

Note: To assist frequent searches, Excel lets you keep the Find and Replace window hanging around (rather than forcing you to use it or close it, as is the case with many other dialog boxes). You can continue to move from cell to cell and edit your worksheet data even while the Find and Replace window remains visible.

3. **In the "Find what" combo box, enter the word, phrase, or number you're looking for.**

   If you've performed other searches recently, you can reuse these search terms. Just choose the appropriate search text from the "Find what" drop-down list.

4. **Click Find Next.**

   Excel jumps to the next matching cell, which becomes the active cell. However, Excel doesn't highlight the matched text or in any way indicate *why* it decided the cell was a match. (That's a bummer if you've got, say, 200 words crammed into a cell.) If it doesn't find a matching cell, Excel displays a message box telling you it couldn't find the requested content.

   If the first match isn't what you're looking for, you can keep looking by clicking Find Next again to move to the next match. Keep clicking Find Next to move through the worksheet. When you reach the end, Excel resumes the search at the beginning of your worksheet, potentially bringing you back to a match you've already seen. When you're finished with the search, click Close to get rid of the Find and Replace window.

Find All

One of the problems with searching in Excel is that you're never quite sure how many matches there are in a worksheet. Sure, clicking Find Next gets you from one cell to the next, but wouldn't it be easier for Excel to let you know right away how many matches it found?

Enter the Find All feature. With Find All, Excel searches the entire worksheet in one go, and compiles a list of matches, as shown in Figure 4-11.

The Find All button doesn't lead you through the worksheet like the Find feature. It's up to you to select one of the results in the list, at which point Excel automatically moves you to the matching cell.

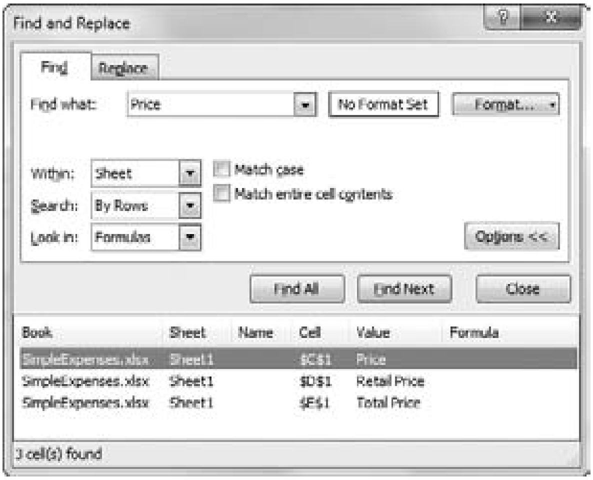

Figure 4-11:
In the example shown here, the search for "Price" matched three cells in the worksheet. The list shows you the complete text in the matching cell and the cell reference (for example, $C$1, which is a reference to cell C1).

The Find All list won't automatically refresh itself: After you've run a Find All search, if you *add* new data to your worksheet, you need to run a new search to find any newly added terms. However, Excel does keep the text and numbers in your found-items list synchronized with any changes you make in the worksheet. For example, if you change cell D5 to Total Price, the change appears in the Value column in the found-items list *automatically*. This tool is great for editing a worksheet because you can keep track of multiple changes at a single glance.

Finally, the Find All feature is the heart of another great Excel guru trick: it gives you another way to change multiple cells at once. After you've performed the Find All search, select all the entries you want to change from the list by clicking them while you hold down Ctrl (this trick lets you select several at once). Click the formula bar, and then start typing the new value. When you're finished, hit Ctrl+Enter to apply your changes to every selected cell. Voilà—it's like Find and Replace, but you're in control!

More Advanced Searches

Basic searches are fine if all you need to find is a glaringly unique phrase or number (*Pet Snail Names* or *10,987,654,321*). But Excel's advanced search feature gives you lots of ways to fine-tune your searches or even search more than one worksheet. To conduct an advanced search, begin by clicking the Options button in the Find and Replace window, as shown in Figure 4-12.

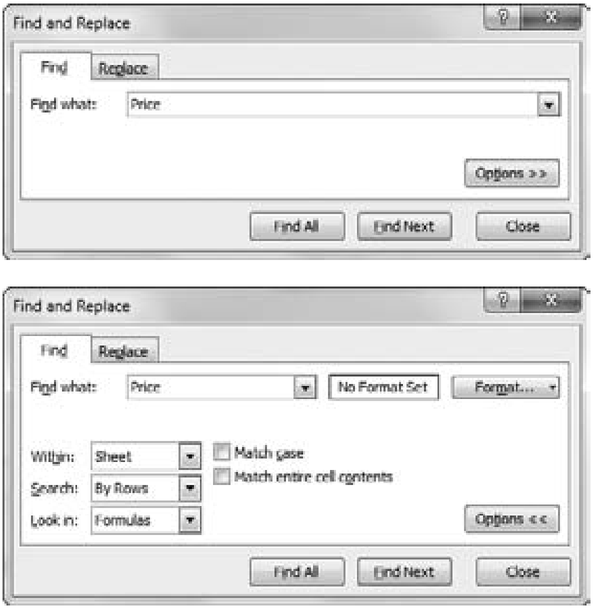

Figure 4-12:
In the standard Find and Replace window (top), when you click Options, Excel gives you a slew of additional settings (bottom) so you can configure things like search direction, case sensitivity, and format matching.

You can set any or all of the following options:

- The Within box controls the span of your search. The standard option, Sheet, searches all the cells in the currently active worksheet. If you want to continue the search in the other worksheets in your workbook, choose Workbook. When performing a workbook search, Excel examines your worksheets from left to right, starting with the current worksheet. When it finishes searching the last worksheet in your workbook, it loops back and starts again at the first worksheet in the workbook.

- The Search box chooses the direction of the search. The standard option, By Rows, completely searches each row before moving on to the next one. That means that if you start in cell B2, Excel searches C2, D2, E2, and so on. Once it's moved through every column in the second row, it moves onto the third row and searches from left to right.

  On the other hand, if you choose By Columns, Excel searches all the rows in the current column before moving to the next column. That means that if you start in cell B2, Excel searches B3, B4, and so on until it reaches the bottom of the column, and then starts at the top of the next column (column C).

Note: The search direction determines which path Excel follows when it's searching. However, the search will still ultimately traverse every cell in your worksheet (or the current selection).

- The "Look in" box tells Excel what to examine in each cell. If you choose Formulas (the standard option), Excel tries to make a match between your search text and the cell *content (for example, the number 3.5 or the date 12/2/2010)*. If you choose Values, Excel tries to make a match between your search text and the cell *display text* (for example, the formatted number $3.50 or the formatted date 2-Dec-10). And if you choose Comments, Excel searches the text of any attached comment boxes (page 724), but ignores the actual cell content.
- The "Match case" option specifies whether capitalization is important. If you select "Match case", Excel finds only words or phrases whose capitalization matches. Thus, searching for *Date* matches the cell value *Date*, but not *date*.
- The "Match entire cell contents" option lets you restrict your searches to the entire contents of a cell. Excel ordinarily looks to see if your search term is contained *anywhere* inside a cell. So, if you specify the word *Price*, Excel finds cells containing text like *Current Price* and even *Repriced Items*. Similarly, a number like 32 will match cell values like 3253, 10032, and 1.321. Turning on the "Match entire cell contents" option forces Excel to be precise.

Note: Remember, Excel searches for numbers as they're *displayed* (as opposed to looking at the underlying values that Excel uses to store numbers internally). That means that if you're searching for a number formatted using the dollar Currency format ($32.00, for example), and you've turned on the "Match entire cell contents" checkbox, you'll need to enter the number exactly as it appears on the worksheet. Thus, $32.00 would work, but 32 alone won't help you.

POWER USERS' CLINIC

Using Wildcards

Sometimes you sorta, kinda know what you're looking for—for example, a cell with some version of the word "date" in it (as in "date" or "dated" or "dating"). What you really need is a search tool that's flexible enough to keep its eyes open for results that are *similar* but not exactly alike. Power searchers will be happy to know that Excel lets you use *wildcards* in your searches. Wildcards are search symbols that let you search for variations on a word.

The asterisk (*) wildcard represents a group of one or more characters. A search for s*nd finds any word that begins with the letter s and ends with the letters nd; for example, it would find words like sand, sound, send, or even the bizarre series of characters *sgrthdnd*. The question mark (?) wildcard represents any single character. For example, a search for *f?nd* turns up *find* or *fund*, but not *friend*.

Wildcards are particularly useful when you're using the "Match entire cell contents" option. For example, if you turn on the "Match entire cell contents" option and enter the search term *date** you'll find any cell that starts with the word *date*. In contrast, if you performed the same search without turning the "Match entire cell contents" option on, you'd find any cell *containing* the word *date*.

If you happen to want to search for special characters like the asterisk or the question mark, you'll need to use the tilde (~) before the wildcard. For example, the search string ~* searches for cells that contain the asterisk symbol.

Finding Formatted Cells

Excel's Find and Replace is an equal opportunity search tool: It doesn't care what the contents of a cell look like. But what if you know, for example, that the data you're looking for is formatted in bold, or that it's a number that uses the Currency format? You can use these formatting details to help Excel find the data you want and ignore cells that aren't relevant.

To use formatting details as part of your search criteria, follow these steps:

1. **Launch the Find tool.**

   Choose Home→Editing→Find & Select→Find, or press Ctrl+F. Make sure that the Find and Replace window is showing the advanced options (by clicking the Options button).

2. **Decide how you want to specify the formatting.**

   You have two options, and they both involve the Format button that appears next to the "Find what" search box.

   The quickest approach is to copy all the format information from another cell. To do this, click the arrow on the right-side of the Format button to pop open a menu with additional options, and then click Choose Format From Cell. The mouse pointer changes to a plus symbol with an eyedropper next to it. Next, click any cell that has the formatting you want to match. Keep in mind that when you use this approach, you copy *all* the format settings.

   A more controlled approach is to specify the exact formatting settings you want to hunt down. To do this, click the Format button. The Find Format dialog box appears (Figure 4-13). It contains the same options as the Format Cells dialog box discussed on page 139. Using the Find Format dialog box, you can specify any combination of settings for number format, alignment, font, fill pattern, and borders. Chapter 5 explains all these settings in detail. You can also search for protected and locked cells, which are described in Chapter 24. When you're finished, click OK to return to the Find and Replace window.

3. **Review your formatting and start your search.**

   Next to the "Find what" search box, a preview appears indicating the formatting of the cells that you're searching for, as shown in Figure 4-14. If everything checks out, click Find All or Find Next to get started.

   To remove these formatting restrictions in subsequent searches, click the arrow on the right of the Format button, and then choose Clear Find Format.

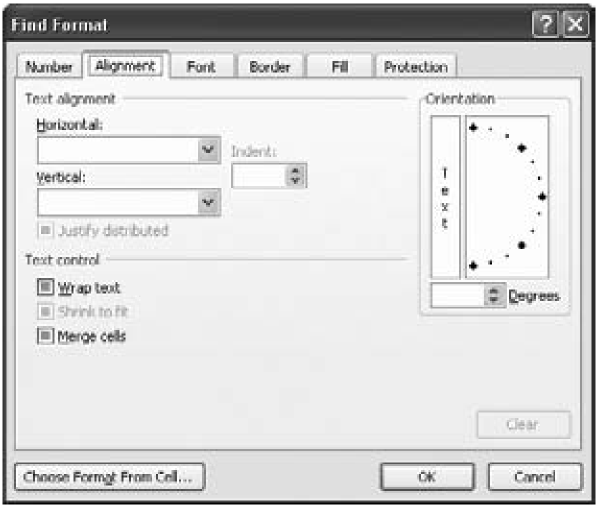

Figure 4-13:
In the Find Format dialog box, Excel won't use any formatting option that's blank or grayed out as part of its search criteria. For example, here, Excel won't search based on alignment. Checkboxes are a little trickier. In some versions of Windows, it looks like the checkbox is filled with a solid square. In other versions of Windows, it looks like the checkbox is dimmed and checked at the same time (as with the "Merge cells" setting in this example). Either way, this visual cue indicates that Excel won't use the setting as part of its search.

Figure 4-14:
The Find Format dialog box shows a basic preview of your formatting choices. In this example, the search will find cells containing the word "price" that also use white lettering, a black background, and the Bauhaus font.

Finding and Replacing Values

You can use Excel's search muscles to find not only the information you're interested in, but also to modify cells quickly and easily. Excel lets you make two types of changes using its *Replace* tool:

- **You can automatically change cell content**. For example, you can replace the word *Colour* with *Color* or the number *$400* with *$40*.
- **You can automatically change cell formatting**. For example, you can search for every cell that contains the word *Price* or the number *$400* and change the fill color. Or, you can search for every cell that uses a specific font, and modify these cells so they use a new font.

Here's how to perform a replace operation. Once you've mastered the technique, check out the box on page 136, which describes some super handy tricks you can do with this process.

1. **Move to the cell where the search should begin.**

   Remember, if you don't want to search the entire spreadsheet, just select the range of cells you want to search.

2. **Choose Home→Editing→Find & Select→Replace, or press Ctrl+H.**

   The Find and Replace window appears, with the Replace tab selected, as shown in Figure 4-15.

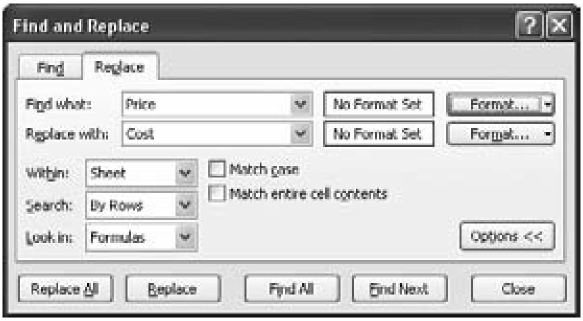

Figure 4-15:
The Replace tab looks pretty similar to the Find tab. Even the advanced options are the same. The only difference is that you also need to specify the text you want to use as a replacement for the search terms you find.

3. **In the "Find what" box, enter your search term. In the "Replace with" box, enter the replacement text.**

   Type the replacement text exactly as you want it to appear. If you want to set any advanced options, click the Options button (see the earlier sections "More Advanced Searches" and "Finding Formatted Cells" for more on your choices).

4. **Perform the search.**

   You've got four different options here. *Replace All* immediately changes all the matches your search identifies. *Replace* changes only the first matched item (you can then click Replace again to move on to subsequent matches or to select any of the other three options). *Find All* works just like the same feature described on page 129. *Find Next* moves to the next match, where you can click Replace to apply your specified change, or click any of the other three buttons. The replace options are good if you're confident you want to make a change; the find options work well if you first want to see what changes you're about to make (although you can reverse either option using Ctrl+Z to fire off the Undo command).

Note: It's possible for a single cell to contain more than one match. In this case, clicking Replace replaces every occurrence of that text in the entire cell.

POWER USERS' CLINIC

Mastering the Art of Replacement

You can use the Find and Replace feature in many imaginative ways. Here are just a few examples:

- **You can automatically delete a specific piece of text**. Just enter the appropriate "Find what" text, and leave the "Replace with" box blank.
- **You can change the formatting used in specific cells**. Just type the same text in both the "Find what" and "Replace with" text, and then click the Format button next to the "Replace with" combo box to set some formatting attributes. (You don't need to specify any formatting settings for your "Find what" search criteria.)
- **You can change the formatting used in a series of cells**. For example, imagine you have a worksheet that has several cells bolded. Say you want to adjust the formatting of these cells to use a new font. To perform this operation, leave both the "Find what" and "Replace with" boxes blank. Then, set the formatting search criteria to look for the bold font attribute, and set the replacement formatting to use the new font. Click Replace All, and all the cells that currently have bold formatting acquire the new font. You might find mastering this technique tricky, but it's one of the most powerful formatting tricks around.

CHAPTER

5

Formatting Cells

When you create a basic workbook, you've taken only the first step toward mastering Excel. If you plan to print your data, email it to colleagues, or show it off to friends, you need to think about whether you've formatted your worksheets in a viewer-friendly way. The careful use of color, shading, borders, and fonts can make the difference between a messy glob of data and a worksheet that's easy to work with and understand.

But formatting isn't just about deciding, say, where and how to make your text bold. Excel also lets you control the formatting of numerical values. In fact, two aspects of formatting are fundamental in any worksheet:

- **Cell appearance**. Cell appearance formatting is all about cosmetic details like color, typeface, alignment, and borders. When most people think of formatting, they think of cell appearance first.
- **Cell values**. Cell value formatting controls the way Excel displays numbers, dates, and times. For numbers, this includes details like whether to use scientific notation, the number of decimal places displayed, and the use of currency symbols, percent signs, and commas. With dates, cell value formatting determines what parts of the date the cell displays, and in what order.

In many ways, cell value formatting is more significant than cell appearance formatting, because it can change the meaning of your data. For example, even though 45%, $0.45, and 0.450 are all the same number (just formatted differently), your spreadsheet readers will see a failing test score, a cheap price for chewing gum, and a world-class batting average.

Note: Keep in mind that regardless of how you *format* your cell values, Excel maintains an unalterable *value* for every number entered. For more on how Excel internally stores numbers, see the box on page 143.

In this chapter, you'll learn about cell value formatting, and then unleash your inner artist with cell appearance formatting.

Formatting Cell Values

The basic principle behind cell value formatting is this: the cell value that Excel *stores* doesn't necessarily match the cell value it *displays*. This gives you the best of both worlds. Your cells can store super-accurate values, but you don't need to clutter your worksheet with numbers that have 13 decimal places.

To make your worksheet's as clear and readable as possible, you need to make sure that the display value is in a form that makes sense for your spreadsheet. Figure 5-1 shows how Excel can show the same number in a variety of ways.

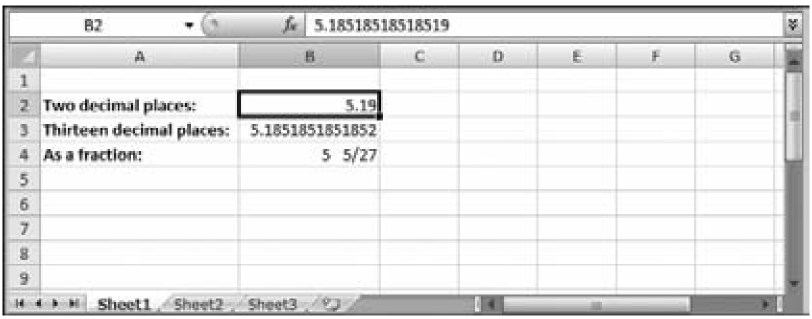

***Figure 5-1:** This worksheet shows how different formatting can affect the appearance of the same data. Each of the cells B2, B3, and B4 contain the same number: 5.18518518518519. In the formula bar, Excel always displays the exact number it's storing, as you see here with cell B2. However, in the worksheet itself, each cell's appearance differs depending on how you format the cell.*

The first time you type a number or date into a blank cell, Excel makes an educated guess about the number format you want. For example, if you type in a currency value like $34.99, Excel assumes you want a number format that uses the dollar sign. If you then type in a new number in the same cell without a dollar sign (say, 18.75), Excel adds the dollar sign automatically (making it $18.75).

Changing the Cell Value Format

Before long, you'll need to change a cell value format, or you'll want to fine-tune it. The basic process unfolds like this:

1. **Select the cells you want to format.**

   You can apply formatting to individual cells or a collection of cells. Usually, you'll want to format an entire column at once because all the values in a column typically contain the same type of data. Remember, to select a column, you simply need to click the column header (the gray box at the top with the column letter) or press Ctrl+Space.

Note: Technically, a column contains *two* types of data: the values you're storing within the actual cells and the column header in the topmost cell (where the text is). However, you don't need to worry about unintentionally formatting the column title because Excel applies number formats only to numeric cells (cells that contain dates, times, or numbers). Excel doesn't use the number format for the column header cell, because it contains text.

---

2. **Select Home→Cells→Format→Format Cells, or just right-click the selection, and then choose Format Cells.**

   In either case, the Format Cells dialog box appears, as shown in Figure 5-2.

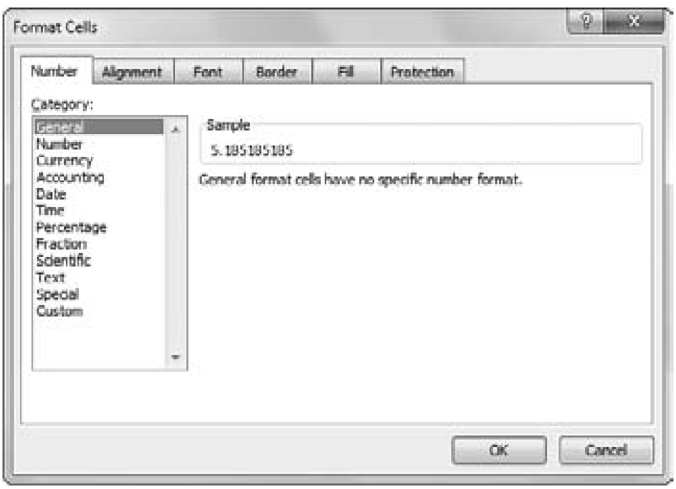

Figure 5-2: *The Format Cells dialog box provides one-stop shopping for cell value and cell appearance formatting. The first tab, Number, lets you configure numeric value formatting. You can use the Alignment, Font, Border, and Fill tabs to control the cell's appearance. Finally, the Protection tab lets you prevent changes and hide formulas. (You'll learn about worksheet protection features in Chapter 24.)*

3. **Set the format options.**

   The Number tab's options let you choose how Excel translates the cell value into a display value. For example, you can change the number of decimal places that Excel uses to show the number. (The next section covers number formatting choices in much more detail.)

   Most of the Format Cells dialog box's other tabs are for cell appearance formatting, which is covered later in this chapter.

Note: Once you apply formatting to a cell, it retains that formatting even if you clear the cell's contents (by selecting it, and then pressing Delete). In addition, formatting comes along for the ride if you copy a cell, so if you copy the content from cell A1 to cell A2, the formatting comes with it. Formatting includes both cell value formatting *and* cell appearance.

The only way to remove formatting is to highlight the cell and select Home→Editing→Clear→Clear Formats. This command removes the formatting, restoring the cell to its original, General number format (which you'll learn more about next), but it doesn't remove any of the cell's content.

4. **Click OK.**

   Excel applies your formatting changes and changes how it displays the values of the selected cells accordingly.

You'll spend a lot of time in this chapter at the Format Cells dialog box. As you've already seen, the most obvious way to get there is to choose Home→Format→Cells→Format Cells. However, your mouse finger's sure to tire out with that method. Fortunately, there's a quicker route—you can use one of three *dialog box launchers*. Figure 5-3 shows the way.

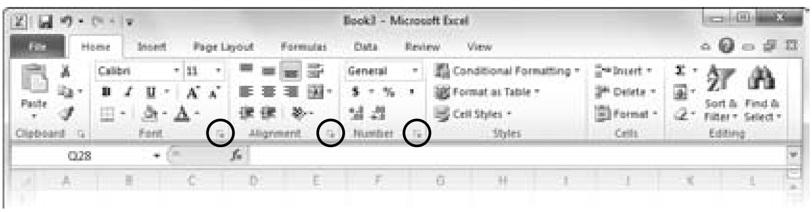

Figure 5-3: *The ribbon's Home tab gives you a quick way to open the Format Cells dialog box from three different spots: the Font tab, the Alignment tab, or the Number tab.*

Tip: If you don't want to take your fingers off the keyboard, you can use the shortcut Ctrl+1 to show the Format Cells dialog box at any time.

Formatting Numbers

In the Format Cells dialog box, the Number tab lets you control how Excel displays numeric data in a cell. Excel gives you a lengthy list of predefined formats (as shown in Figure 5-4), and also lets you design your own formats. Remember, Excel uses number formats when the cell contains only numeric information. Otherwise, Excel simply ignores the number format. For example, if you enter *Half past* 12 in a column full of times, Excel considers it plain ol' text—although, under the hood, the cell's numerical formatting stays put, and Excel uses it if you change the cell content to a time.

When you create a new spreadsheet, every cell starts out with the same number format: General. This format comes with a couple of basic rules:

- If a number has any decimal places, Excel displays them, provided they fit in the column. If the number has more decimal places than Excel can display, it leaves out the ones that don't fit. (It rounds up the last displayed digit, when appropriate.) If you change a column width, Excel automatically adjusts the amount of digits it displays.
- Excel removes leading and trailing zeros. Thus, 004.00 becomes 4. The only exception to this rule happens with numbers between –1 and 1, which retain the 0 before the decimal point. For example, Excel displays the number .42 as 0.42.

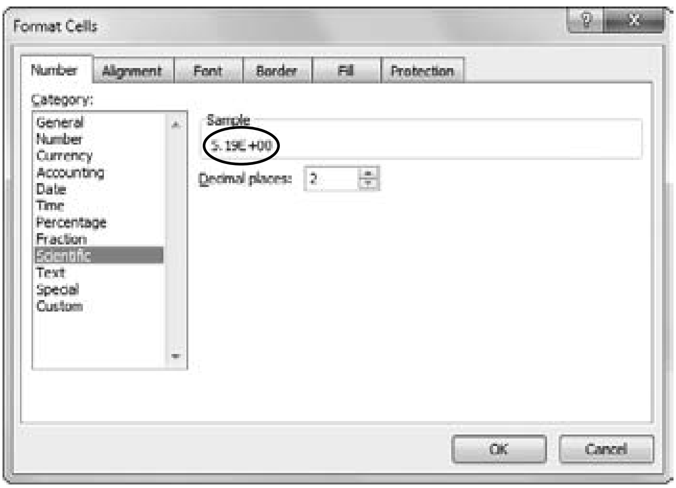

Figure 5-4:
You can learn about the different number formats by selecting a cell that already has a number in it, and then choosing a new number format from the Category list (Home→Cells→Format→ Format Cells). When you do so, Excel uses the Format Cells dialog box to show how it'll display the number if you apply that format. In this example, you see that the cell value, 5.18518518518519, will appear as 5.19E+00, which is scientific notation with two decimal places.

As you saw in Chapter 2, the way you type in a number can change a cell's formatting. For example, if you enter a number with a currency symbol, the number format of the cell changes automatically to Currency. Similarly, if you enter three numbers separated by dashes (-) or forward slashes (/), Excel assumes you're entering a date, and adjusts the number format to Date.

However, rather than rely on this automatic process, it's far better just to enter ordinary numbers and set the formatting explicitly for the whole column. This approach prevents you from having different formatting in different cells (which can confuse even the sharpest spreadsheet reader), and it makes sure you get exactly the formatting and precision you want. You can apply formatting to the column before or after you enter the numbers. And it doesn't matter if a cell is currently empty; Excel still keeps track of the number format you've applied.

Different number formats provide different options. For example, if you choose the Currency format, then you can choose from dozens of currency symbols. When you use the Number format, you can choose to add commas (to separate groups of three digits) or parentheses (to indicate negative numbers). Most number formats let you set the number of decimal places.

The following sections give a quick tour of the predefined number formats available in the Format Cells dialog box's Number tab. Figure 5-5 gives you an overview of how different number formats affect similar numbers.

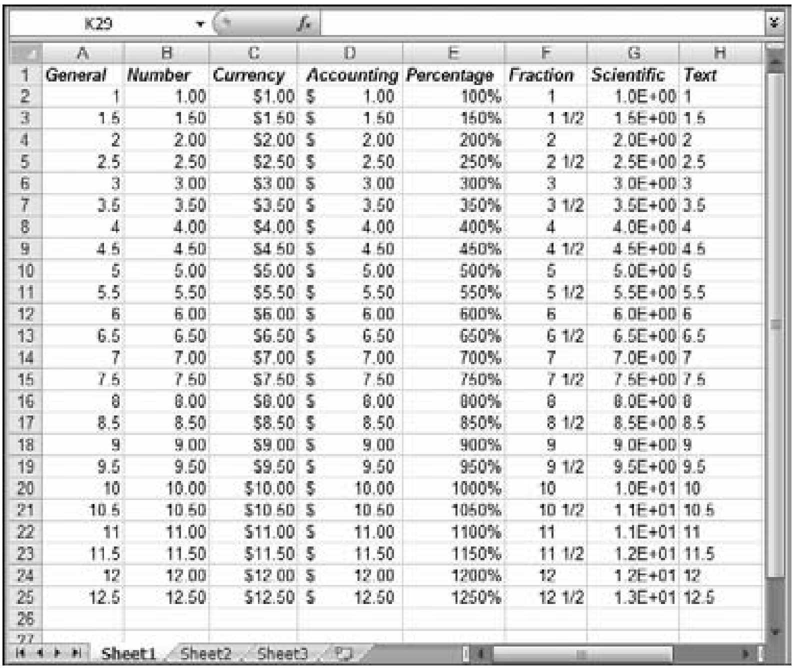

	A	B	C	D	E	F	G	H
1	General	Number	Currency	Accounting	Percentage	Fraction	Scientific	Text
2	1	1.00	$1.00	$ 1.00	100%	1	1.0E+00	1
3	1.5	1.50	$1.50	$ 1.50	150%	1 1/2	1.5E+00	1.5
4	2	2.00	$2.00	$ 2.00	200%	2	2.0E+00	2
5	2.5	2.50	$2.50	$ 2.50	250%	2 1/2	2.5E+00	2.5
6	3	3.00	$3.00	$ 3.00	300%	3	3.0E+00	3
7	3.5	3.50	$3.50	$ 3.50	350%	3 1/2	3.5E+00	3.5
8	4	4.00	$4.00	$ 4.00	400%	4	4.0E+00	4
9	4.5	4.50	$4.50	$ 4.50	450%	4 1/2	4.5E+00	4.5
10	5	5.00	$5.00	$ 5.00	500%	5	5.0E+00	5
11	5.5	5.50	$5.50	$ 5.50	550%	5 1/2	5.5E+00	5.5
12	6	6.00	$6.00	$ 6.00	600%	6	6.0E+00	6
13	6.5	6.50	$6.50	$ 6.50	650%	6 1/2	6.5E+00	6.5
14	7	7.00	$7.00	$ 7.00	700%	7	7.0E+00	7
15	7.5	7.50	$7.50	$ 7.50	750%	7 1/2	7.5E+00	7.5
16	8	8.00	$8.00	$ 8.00	800%	8	8.0E+00	8
17	8.5	8.50	$8.50	$ 8.50	850%	8 1/2	8.5E+00	8.5
18	9	9.00	$9.00	$ 9.00	900%	9	9.0E+00	9
19	9.5	9.50	$9.50	$ 9.50	950%	9 1/2	9.5E+00	9.5
20	10	10.00	$10.00	$ 10.00	1000%	10	1.0E+01	10
21	10.5	10.50	$10.50	$ 10.50	1050%	10 1/2	1.1E+01	10.5
22	11	11.00	$11.00	$ 11.00	1100%	11	1.1E+01	11
23	11.5	11.50	$11.50	$ 11.50	1150%	11 1/2	1.2E+01	11.5
24	12	12.00	$12.00	$ 12.00	1200%	12	1.2E+01	12
25	12.5	12.50	$12.50	$ 12.50	1250%	12 1/2	1.3E+01	12.5
26								
27								

Figure 5-5: Each column contains the same list of numbers. Although this worksheet shows an example for each number format (except dates and times), it doesn't show all your options. Each number format has its own settings (like the number of decimal places) that affect how Excel displays data.

General

The General format is Excel's standard number format; it applies no special formatting other than the basic rules described on page 66. General is the only number format (other than Text) that doesn't limit your data to a fixed number of decimal places. That means if you want to display numbers that differ wildly in precision (like 0.5, 12.334, and 0.120986398), it makes sense to use General format. On the other hand, if your numbers have a similar degree of precision (for example, if you're logging the number of miles you run each day), the Number format makes more sense.

Number

The Number format is like the General format but with three refinements. First, it uses a fixed number of decimal places (which you set). That means that the decimal point always lines up (assuming you've formatted an entire column). The Number format also lets you use commas as a separator between groups of three digits, which is handy if you're working with really long numbers. Finally, you can choose to have negative numbers displayed with the negative sign, in parentheses, or in red lettering.

UP TO SPEED

The Relationship Between Formatting and Values

The format that you choose for a number doesn't affect Excel's internal storage of that number. For example, if a cell contains the fraction 1/3, then Excel stores this value as 0.333333333333333. (The exact number of decimal places varies, depending on the number you've entered, due to the slight approximations computers need to make when converting fractional numbers into 0s and 1s.) When deciding how to format a cell, you may choose to show only two decimal places, in which case the number appears in your worksheet as 0.33. Or maybe you choose just one decimal place, in which case the number is simply 0.3. In both cases, Excel still keeps the full 15 or so decimal places on hand. To tell the difference between the displayed number and the real number that Excel stores behind the scenes, just move to the cell. Then look at the formula bar, which always shows you the real deal.

Because of this difference between the stored value and the displayed number, there may be some situations where it looks like Excel's making a mistake. For example, imagine you have three cells, and each stores 0.333333333333333 but displays only 0.3. When you add these three cell values together, you won't end up with 0.3 + 0.3 + 0.3 = 0.9. Instead, you'll add the more precise stored values and end up with a number that's infinitesimally close to, but not quite, 1. Excel rounds this number up to 1.

This is almost always the way you want Excel to work, because you know full well that if you add up 1/3 three times you end up with 1. But, if you need to, you can change this behavior.

To change what Excel does, select File→Options, choose the Advanced section, and then scroll down to the "When calculating this workbook" group of settings. A "Set precision as displayed" checkbox appears. When you turn on this checkbox, Excel adjusts all the values in your current spreadsheet so that the stored value matches the displayed value. Unfortunately, with this choice, you get less precise data. For example, if you use this option with the 1/3 example, Excel stores the display value 0.3 instead of 0.333333333333333. Because you can't reverse this change, Excel warns you and asks for a final confirmation when you try to apply the "Precision as displayed" setting.

Currency

The Currency format closely matches the Number format, with two differences. First, you can choose a currency symbol (like the dollar sign, pound symbol, euro symbol, and so on) from an extensive list; Excel displays the currency symbol before the number. Second, the Currency format always includes commas. The Currency format also supports a fixed number of decimal places (chosen by you), and it lets you customize how negative numbers are displayed.

Accounting

The Accounting format is modeled on the Currency format. It also lets you choose a currency symbol, uses commas, and has a fixed number of decimal places. The difference is that the Accounting format uses a slightly different alignment. The currency symbol's always at the far left of the cell (away from the number), and there's always an extra space that pads the right side of the cell. Also, the Accounting format always shows negative numbers in parentheses, which is an accounting standard. Finally, the Accountancy format never shows the number 0. Instead, a dash (–) is displayed in its place. There's really no reason to prefer the Currency or the Accounting format. Think of it as a personal decision, and choose whichever looks nicest on your worksheet. The only exception is if you happen to *be* an accountant, in which case you really have no choice in the matter—stick with your namesake.

Percentage

The Percentage format displays fractional numbers as percentages. For example, if you enter 0.5, that translates to 50 percent. You can choose the number of decimal places to display.

There's one trick to watch out for with the Percentage format. If you forget to start your number with a decimal, Excel quietly "corrects" your numbers. For example, if you type 4 into a cell that uses the Percentage format, Excel interprets this as 4 percent. As a result, it actually stores the value 0.04. A side-effect of this quirkiness is that if you want to enter percentages larger than 100 percent, you can't enter them as decimals. For example, to enter 200 percent, you need to type in 200 (not 2.00).

Fraction

The Fraction format displays your number as a fraction instead of a number with decimal places. The Fraction format doesn't mean you must enter the number as a fraction (although you can if you want by using the forward slash, like 3/4). Instead it means that Excel converts any number you enter and displays it as a fraction. Thus, to have 1/4 appear you can either enter .25 or 1/4.

Note: If you try to enter 1/4 and you *haven't* formatted the cell to use the Fraction number format, you won't get the result you want. Excel assumes you're trying to enter a date (in this case, January 4 of the current year). To avoid this misunderstanding, change the number format *before* you type in your fraction. Or, enter it as *0 1/4* (zero and one quarter).

People often use the Fraction format for stock market quotes, but it's also handy for certain types of measurements (like weights and temperatures). When using the Fraction format, Excel does its best to calculate the closest fraction, which depends on a few factors including whether an exact match exists (entering .5 always gets you 1/2, for example) and what type of precision level you've picked when selecting the Fraction formatting.

You can choose to have fractions with three digits (for example, 100/200), two digits (10/20), or just one digit (1/2), using the top three choices in the Type list. For example, if you enter the number 0.51, then Excel shows it as 1/2 in one-digit mode, and the more precise 51/100 in three-digit mode. In some cases, you may want all numbers to use the same denominator (the bottom number in the fraction) so that it's easy to compare different numbers. (Don't you wish Excel had been around when you were in grammar school?) In this case, you can choose to show all fractions as halves (with a denominator of 2), quarters (a denominator of 4), eighths (8), sixteenths (16), tenths (10), and hundredths (100). For example, the number 0.51 displays as 2/4 if you choose quarters.

Tip: Entering a fraction in Excel can be awkward, because Excel may attempt to convert it to a date. To prevent this confusion, always start by entering 0, and then a space. For example, instead of typing 2/3, enter 0 2/3 (which means zero and two-thirds). If you have a whole number and a fraction, like 1 2/3, you'll also be able to duck the date confusion.

FREQUENTLY ASKED QUESTION

Just How Precise Are Excel Numbers, Anyway?

Can I enter a number with 10 decimal places? How about 20?

Here's a handy way to find out: type the fraction *2/3* into a cell, and then check the formula bar, which shows you the number Excel has stored. Turns out Excel thinks of 2/3 as 0.666666666666667.

This test shows that Excel is limited to about 15 significant digits, and it rounds up the last digit. You may be slightly unnerved by the word "about," but in the binary world of computers, fractional numbers don't have a fixed number of digits and may just be approximations with very slight rounding errors. You can find a good (but technical) explanation of this phenomenon on the online encyclopedia Wikipedia: *http://en.wikipedia.org/wiki/Floating_point.*

Because Excel doesn't store fractions precisely as they exist in the world of real math, you may occasionally experience minor rounding errors in calculations with more than 14 significant digits. (Recall from high-school math that the number of significant digits is the number of digits starting with the first nonzero digit and ending with the last nonzero digit. Essentially, the significant digits hold all the information in your number.) This behavior shouldn't cause you to panic—it's a limitation of nearly all computer software, based on the way that computers manipulate numbers.

Scientific

The Scientific format displays numbers using scientific notation, which is ideal when you need to handle numbers that range widely in size (like 0.0003 and 300) *in the same column*. Scientific notation displays the first non-zero digit of a number, followed by a fixed number of digits, and then indicates what power of 10 that number needs to be multiplied by to generate the original number. For example, 0.0003 becomes 3.00×10^{-4} (displayed in Excel as 3.00E-04). The number 300, on the other hand, becomes 3.00×10^{2} (displayed in Excel as 3.00E02). Scientists—surprise, surprise—like the Scientific format for doing things like recording experimental data or creating mathematical models to predict when an incoming meteor will strike the Earth.

Text

Few people use the Text format for numbers, but it's certainly possible to do so. The Text format simply displays a number as though it were text, although you can still perform calculations with it. Excel shows the number exactly as it's stored internally, positioning it against the left edge of the column. You can get the same effect by placing an apostrophe before the number (although that approach won't let you use the number in calculations).

TIMESAVING TIP

Shortcuts in the Ribbon

You don't need to waste hours jumping between your worksheet and the Format Cells dialog box. The ribbon gets you to some of the most commonly used number formats in the Home→Number section.

The Home→Number section's most prominent part is the drop-down list of number formats (Figure 5-6). Just underneath are buttons that let you apply one of the three most common formats: Accounting, Percent, or Number. Just to the right are two buttons that let you increase or decrease the number of decimal places that you see at once.

One of the neatest features is the list of currency options for the Accounting button. If you click the drop-down arrow on the Accounting button (which looks like a dollar sign), you see a list with different currency symbols you can choose (like Pounds, Euros, Chinese Yuan, and so on). But if you click the other portion of the Accounting button (not the arrow), you get the currency symbol that's appropriate, based on your computer's regional settings.

Formatting Dates and Times

Excel gives you lots of options here. You can use everything from compact styles like 3/13/10 to longer formats that include the day of the week, like Saturday, March 13, 2010. Time formats give you a similar range of options, including the ability to use a 12-hour or 24-hour clock, show seconds, show fractional seconds, and include the date information.

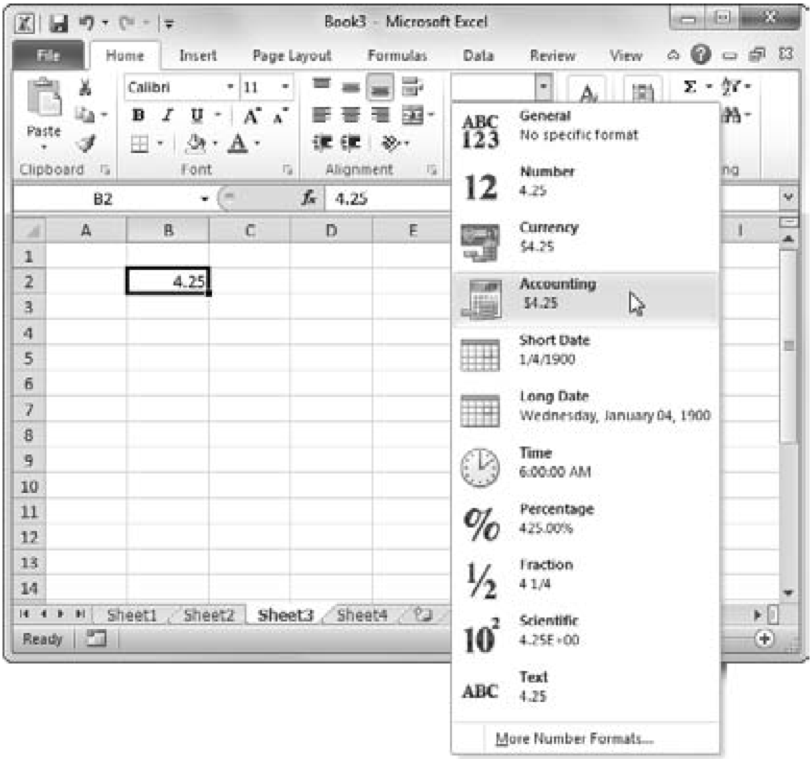

Figure 5-6: *The all-around quickest way to apply a number format is to select some cells, and then, from the number format list, choose an option. Best of all, you see a small preview of what the value in the first selected cell will look like if you apply the format.*

To format dates and times, first open the Format Cells dialog box shown in Figure 5-7 (Home→Cells→Format→Format Cells). Choose Date or Time from the column on the left, and then choose the format from the list on the right. Date and Time both provide a slew of options.

Excel has essentially two types of date and time formats:

- **Formats that take the regional settings of the spreadsheet viewer's computer into account**. With these formats, dates display differently depending on the computer that's running Excel. This choice is a good one, because it lets everyone see dates in just the way they want to, which means no time-consuming arguments about month-day-year or day-month-year ordering.
- **Formats that *ignore* the regional settings of individual computers**. These formats define a fixed pattern for month, day, year, and time components, and display date-related information in exactly the same way on all computers. If you need to absolutely make sure a date is in a certain format, you should use this choice.

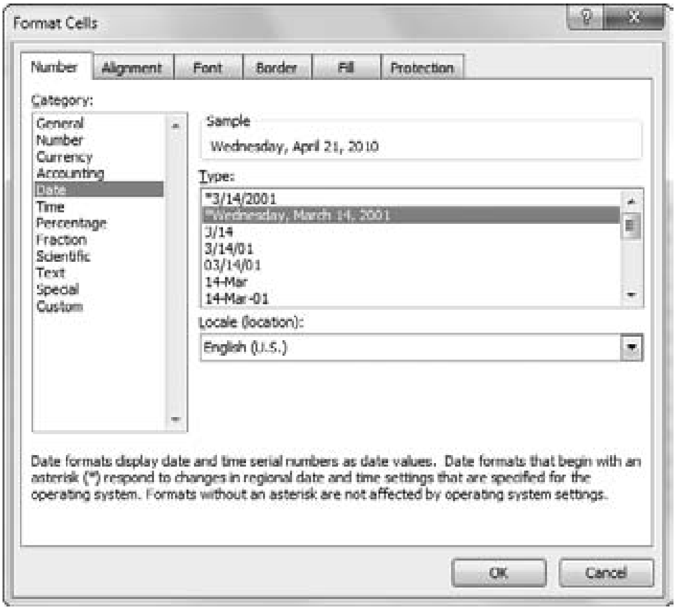

Figure 5-7:
Excel gives you dozens of different ways to format dates and times. You can choose between formats that modify the date's appearance depending on the regional settings of the computer viewing the Excel file, or you can choose a fixed date format. When using a fixed date format, you don't need to stick to the U.S. standard. Instead, choose the appropriate region from the Locale list box. Each locale provides its own set of customized date formats.

The first group (the formats that rely on a computer's regional settings) is the smallest. It includes two date formats (a compact, number-only format and a long, more descriptive format) and one time format. In the Type list, these formats are at the top and have an asterisk next to them.

The second group (the formats that are independent of a computer's regional settings) is much more extensive. In order to choose one of these formats, you first select a region from the Locale list, and then you select the appropriate date or time format (that isn't preceded by an asterisk). Some examples of locales include "English (U.S.)" and "English (U.K.)."

If you enter a date without specifically formatting the cell, Excel usually uses the short region-specific date format. That means that the order of the month and year vary depending on the regional settings of the current computer. If you incorporate the month name (for example, January 1, 2010), instead of the month number (for example, 1/1/2010), Excel uses a medium date format that includes a month abbreviation, like 1-Jan-2010.

Note: You may remember from Chapter 2 that Excel stores a date internally as the cumulative number of days that have elapsed since a certain long-ago date. You can take a peek at this internal number using the Format Cells dialog box. First, enter your date. Then, format the cell using one of the number formats (like General or Number). The underlying date number appears in your worksheet where the date used to be.

Special Formats for Special Numbers

You wouldn't ever want to perform mathematical operations with some types of numeric information. For example, it's hard to image a situation where you'd want to add or multiply phone numbers or Social Security numbers.

When entering these types of numbers, therefore, you may choose to format them as plain old text. For example, you could enter the text (555) 123-4567 to represent a phone number. Because of the parentheses and the dash (–), Excel won't interpret this information as a number. Alternatively, you could just precede your value with an apostrophe (') to explicitly tell Excel that it should be treated as text (you might do this if you don't use parentheses or dashes in a phone number).

But whichever solution you choose, you're potentially creating more work for yourself, because you must enter the parentheses and the dash for each phone number you enter (or the apostrophe). You also increase the likelihood of creating inconsistently formatted numbers, especially if you're entering a long list of them. For example, some phone numbers may end up entered in slightly similar but somewhat different formats, like 555-123-4567 and (555)1234567.

To avoid these problems, apply Excel's Special number format (shown in Figure 5-8), which converts numbers into common patterns. And lucky you: In the Special number format, one of the Type options is Phone Number (other formats are for Zip codes and Social Security numbers).

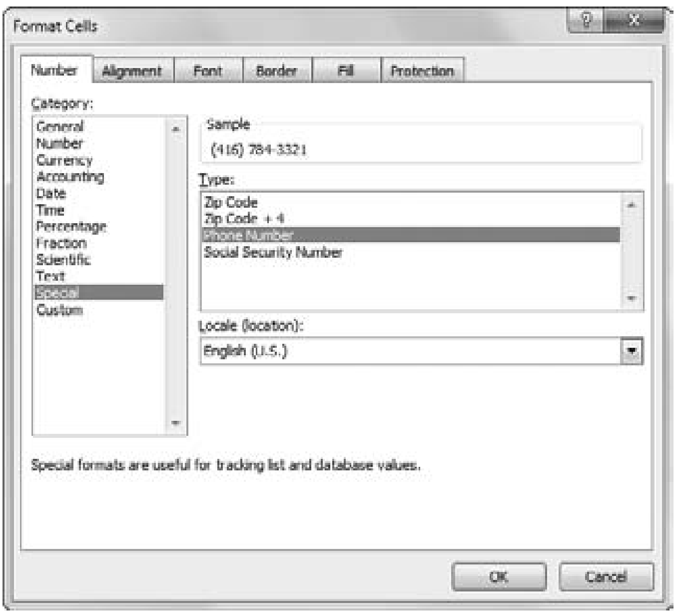

Figure 5-8:
Special number formats are ideal for formatting sequences of digits into a common pattern. For example, in the Type list, if you choose Phone Number, then Excel converts the sequence of digits 5551234567 into the proper phone number style–(555) 123-4567–with no extra work required on your part.

The Special format is a good idea, but it's limited. Out of the box, Excel provides only a small set of special types that you can use. However, there's no reason you can't handle similar problems by creating your own custom formats, as you'll see in the next section.

Custom Formats

As versatile as Excel is, it can't read your mind. You'll find some situations when you want to format numbers in a specialized way that Excel just doesn't expect. For example, you may want to use the ISO (International Organization for Standardization) format for dates, which is used in a wide range of scientific and engineering documents. This format is year-month-day (as in 2010-12-25). Although it's fairly straightforward, Excel doesn't provide this format as a standard option.

Or maybe you want to type in short versions of longer numbers. For example, say your company, International Pet Adventures, uses an employee number to identify each worker, in the format 0521-1033. It may be that 0521- is a departmental identification code for the Travel department. To save effort, you want to be able to enter 1033 and have Excel automatically insert the leading 0521- in your worksheets.

The solution lies in creating your own *custom formats*. Custom formats are a powerful tool for taking control of how Excel formats your numbers. Unfortunately, they aren't exactly easy to master.

The basic concept behind custom formats is that you define the format using a string of special characters and placeholders. This *format string* tells Excel how to format the number or date, including details like how many decimal places it should include, and how it should treat negative numbers. You can also add fixed characters that never change, like the employee number format just described.

Creating a custom format

Here's the easiest way to apply a custom format:

1. **Select the cells you want to format.**

   This selection can include any combination of cells, columns, rows, and so on. To make life easier, make sure the first cell you select contains a value you want to format. That way, you'll be able to use the Format Cells dialog box to preview the effect of your custom format.

Note: Excel saves any custom format strings you create as part of your workbook file. Once you've perfected a format string you like, you can apply it to as many cells as you want.

2. **Select Home→Cells→Format→Format Cells, or just right-click the selection, and then choose Format Cells.**

   The Format Cells dialog box appears, as shown earlier in Figure 5-2.

3. **Choose a format that's similar to the format you want to use.**

   For example, if you want to apply a custom date format, begin by selecting the Date number format, and then choosing the appropriate style. If you want to apply a custom currency format, begin by selecting the Currency number format, and then specifying the appropriate options (like the number of decimal places).

   To create the International Pet Adventures employee code, it makes sense to first select the Number format, and then choose 0 decimal places (because the number format you're looking to model—0521-1033—doesn't use any decimal places).

4. **At the bottom of the Category list, click Custom.**

   Now you see a list of different custom number strings. At the top of this list is a highlighted format string that's based on the format you chose in step 3. Now, you just need to modify this string to customize the format. (Make sure you don't accidentally select another format before you click Custom, or you won't end up with the right format string.)

   If you're creating the International Pet Adventures employee code, you'll see a 0. This 0 means you can use any number without a decimal place. However, what you really want in this situation is to create an employee number that always starts with 0521- and then has four more digits. You'll specify your new format in the next step.

5. **Enter your custom string.**

   Type your custom string into the box below the Type label, as shown in Figure 5-9. The correct format string for the International Pet Adventures example is as follows:

   ```
   "0521-"0000
   ```

   This string tells Excel to begin all cells that are formatted using this Custom format with the digits 0521-, and then follow them with whatever four numbers are entered into the cell (if no numbers are entered, four zeros follow the 0521-). The following sections explain all the ingredients you can use in your custom format.

Tip: Remember, you can preview your custom format string in the Format Cells dialog box's Sample section. As you adjust the custom format string, the Sample box shows you what the current cell would look like if you applied the custom format to it.

6. **Click OK to commit your changes.**

   If the results aren't quite what you want, you can start over again. But this time, skip step 3 because you want to change the current format string rather than replace it with a new format string.

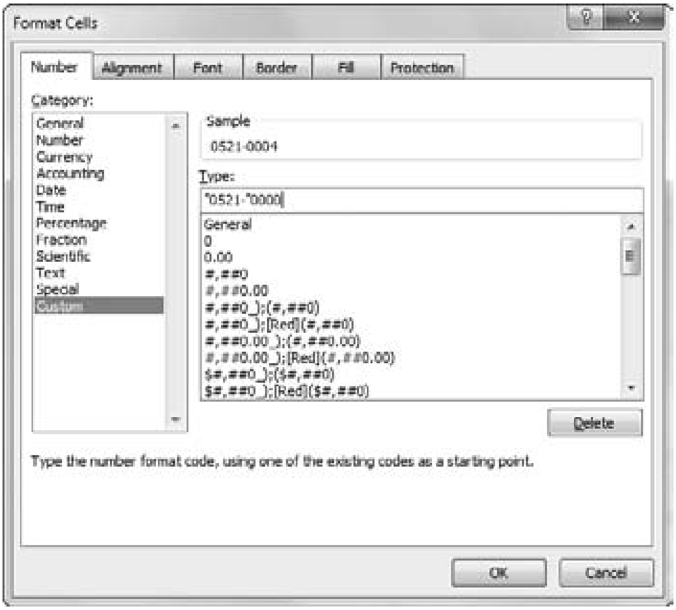

Figure 5-9:
Custom number strings let you do almost anything with a number format, but you need to spell it out explicitly using the cryptic code Excel provides. In the example shown here, the custom format string is "0521-"0000. The "0521-" is a fixed string of characters that's added to the beginning of every number. The four zeros indicate that the cell's expecting four digits. If you provide a one-, two-, or three-digit number, Excel adds the zeroes needed to make a four-digit number. For example, Excel automatically displays the number 4 as the employee code 0521-0004.

7. **To use the Custom format you've created, select one or more cells, show the Format Cells dialog box (by right-clicking the cells, and then choosing Format Cells), and then select your new Custom format.**

   Excel lists newly created Custom formats in the Custom category, at the bottom of the Type list. If you wanted to use the new International Pet Adventures employee code, click OK after selecting your new format, and then begin entering the four digits specific to each employee. For example, if you format a cell with the new Custom format, and then type 6754 into the cell, you'll see *0521-6754*. Remember, despite your craft formatting, the actual number is still 6754, which is important to know if you use the number in a calculation.

Custom format string codes

The tricky part about Custom formats is creating the right format string. To the untrained eye, the format string looks like a cryptic jumble of symbols—which it is. But these symbols, or *formatting codes* in Excel lingo, actually have very specific and clear meanings.

For example, the format string $#,##0.00 translates into the following series of instructions:

$ tells Excel to add a currency symbol before the number.

#,## tells Excel to use commas to separate thousands.

0.00 tells Excel to always include a single digit and two decimal places, no matter what the number is.

In fact, $#,##0.00 is the format string for the basic Currency format. Once you understand what the codes stand for and how they work together, you can create some really useful Custom format strings.

You have three types of codes at your disposal for creating format strings: those used to format dates and times; those used to format numbers; and those used to format ordinary text. The following three sections tackle each type of format code.

Date and time format strings

Date and time format strings are built out of pieces. Each piece represents a single part of the date, like the day, month, year, minute, hour, and so on. You can combine these pieces in whatever order you want, and you can insert your own custom text along with these values.

Note: Keep in mind that none of these formatting codes actually generate or insert the date in your worksheet for you. That is, simply formatting an empty cell with one of these custom strings isn't going to cause the date to appear. Instead, these format strings take the dates that ***you*** enter and make sure that they all appear in a uniform style.

Table 5-1 shows the basic ingredients for a date or time format string. These strings are placeholders that represent the different parts of the date. If you want to include fixed text along with the date, put it in quotation marks.

Table 5-1. *Date and time formatting codes*

Code	Description	Example value displayed on worksheet
d	The day of the month, from 1 to 31, with the numbers between 1 and 9 appearing without a leading 0.	7
dd	The day of the month, from 01 to 31 (leading 0 included from 1 to 9).	07
ddd	A three-letter abbreviation for the day of the week.	Thu
dddd	The full name of the day of the week.	Thursday
m	The number value, from 1 to 12, of the month (no leading 0 used).	1

mm	The number value, from 01 to 12, of the month (leading 0 used for 01 to 09).	01
mmm	A three-letter abbreviation for the month.	Jan
mmmm	The full name of the month.	January
yy	A two-digit abbreviation of the year.	10
yyyy	The year with all four digits.	2010
h	The hour, from 0 to 23 (no leading 0 used).	13
hh	The hour, from 00 to 23 (leading 0 used from 00 to 09).	13
:m	The minute, from 0 to 59.	5
:mm	The minute, from 0 to 59 (leading 0 used for 00 to 09).	05
:s	The second, from 0 to 59 (no leading 0 used). If you want to add tenths, hundredths, or thousandths of a second, follow this with .0 or .00 or .000, respectively. For example, :s.0	5
:ss	The second, from 0 to 59 (leading 0 used from 00 to 09). If you want to add tenths or hundredths of a second, follow this with .0 or .00, respectively.	05
AM/PM	Tells Excel to use a 12-hour clock, including the AM or PM tag.	PM
A/P	Tells Excel to use a 12-hour clock, with an A or P tag.	P
[]	Tells Excel that a given time component (hour, minute, or second) shouldn't "roll over." For example, Excel's standard approach is to have seconds become minutes once they hit the 60 mark, and minutes become hours at the 60 mark. Similarly, hours roll over into a new day when they hit 24. But if you don't want this to happen—for example, when tracking total time on a CD playlist—you could use a format string like [mm]:ss or [h]:mm:ss. The first format string shows the total time in minutes and seconds (without rolling over to hours), while the second shows the total time in hours, minutes, and seconds (without rolling over to days).	133:12 (for the format string [mm]:ss and a value just over two hours)

For example, consider the following format string:

```
yyyy-mm-dd
```

If you apply this format string to a cell that contains a date, you'll end up with the following in your worksheet (assuming you entered the date January 15, 2010): *2010-01-15*.

Note: Regardless of how you type in the date, once you've formatted a cell using a Custom format string, that always overrides the format you use when you type in the date. In other words, it doesn't matter whether you type *1/15/10* or *January 15, 2010* in the cell—Excel still displays it as 2010-01-15, if that's what your custom format dictates.

Now if you format the same value with this format string:

```
"Day "yyyy-mm-dd
```

You'll see this in your worksheet:

```
Day 2010-01-15
```

And remember, whatever information you choose to display or hide, Excel always stores the same date internally.

Note: You'll learn much more about date and time calculations in Chapter 11.

Number format strings

Custom number formats are more challenging than Custom date formats, because Excel gives you lots of flexibility when it comes to customizing number formats. Table 5-2 shows the different codes you can use. The most important of these codes are the digit placeholders 0, ?, and #. You use these to tell Excel where it should slot in the various digits of the number that's currently in the cell (or that you're typing in). For example, a format string that looks like this:

```
#,###.00
```

displays the numbers 45 and 4,500 like this:

```
   45.00
4,500.00
```

In this format string, the # character is a placeholder that lets you put the comma wherever you want it. The 0 character is a placeholder that makes sure the trailing zeroes appear, even when you're showing a whole number. Table 5-2 reveals many more tricks of the trade.

Table 5-2. *Number formatting codes*

Code	Description	Example
0	This digit placeholder forces a zero to appear whenever a number isn't provided.	0.000 displays .3 as 0.300.
?	This digit placeholder forces a space to appear whenever a number isn't provided.	?.??? displays .3 as " .3 " (quotations used to indicate spacing).

#	This digit placeholder indicates where you can place a number, but doesn't automatically insert a zero or a space if there isn't a number in this position. You can use this symbol to set the precision of decimal values or to indicate where commas should go.	### displays .3 as .3. #,### displays 9999 as 9,999.
.	The period, or decimal point, determines where the decimal place will go. It's used in conjunction with the digit placeholders 0, ?, and #.	#.## truncates 1.23456 to 1.23.
,	The comma causes the thousands separator to appear. It's used with the digit placeholders 0, ?, and #. You can also use it to scale a number. For example, if you place one comma at the end of a format string, then Excel displays the number rounded to thousands. Add two commas to the end, and Excel displays the number only in millions (and so on).	#,### displays 3000 as 3,000. #,###,###, displays 123456789 as 123,457 (rounded to thousands).
/	The forward slash formats a number as a fraction. You use this symbol in conjunction with the digit placeholders ? and # to indicate the number of digits you want in your fraction.	?/? displays 1.75 as 7/4, while # ?/? formats the same number as 1 3/4.
E+	This code formats numbers using scientific notation. You use this symbol in conjunction with the digit placeholders 0, ?, and #.	#.## E+## means Excel displays 12345 as 1.23 E+4.
[color]	Applies a specified color to the text that follows the closing bracket. The color name goes inside the square brackets. Excel supports eight colors: [black], [blue], [cyan], [green], [magenta], [red], [white], and [yellow].	[red]#,### displays the number that follows the brackets in red lettering.
_	The underscore character, when followed by any other character, inserts a space equal to the width of that other character. This code is occasionally used when aligning complex formatting codes.	_W inserts a space as wide as the capital letter W.
*	The asterisk, when followed by any other character, inserts that other character in a cell until the cell is filled.	#,###*- displays 9999 as 9,999------ (with the dashes appearing until the cell is filled).

Note: Excel uses custom number formats to decide how to round off displayed numbers, and how to format them (by adding commas, currency symbols, and so on). But no matter what format string you use, you can't coax Excel into shaving off digits that appear to the *left* of the decimal place—and for good reason: doing so would mangle your numbers beyond recognition.

It helps to keep a few pointers in mind when you use the number formatting codes listed in Table 5-2:

- Use 0 to indicate a number that *must* be wherever the 0 is placed—if it's not, Excel automatically inserts a 0. For example, the format string 0.00 would display the number .3 as 0.30. And the format string 00.00 would format the same value as 00.30.
- The question mark (?) works similarly, but it turns into spaces instead of zeros, ensuring that multiple numbers wind up aligned in a column. For example, ??.?? displays the number 3 as " 3. " (without the quote marks).
- The # symbol lets you indicate where a number *can* exist but doesn't *have to* exist. For example, the format string 0.0# indicates that the first digit before the decimal place and the first digit after the decimal place must be present (that's what the zeros tell Excel). However, the second number after the decimal place is optional. With this format string, Excel rounds additional digits starting with the third decimal place. Thus, this format string displays the value .3 as 0.3, .34 as 0.34, and .356 as 0.36. You can also use the # symbol to indicate where commas should go, as in the format string #,###.00. This string displays the value 3639 as 3,639.00.

Note: Remember, Custom format strings control how Excel displays values. These strings aren't meant to control what values someone can enter in a cell. To set rules for allowed data, you need a different feature–data validation, described on page 694.

Excel also lets you use codes that apply currency symbols, percent symbols, and colors. As with date values, you can insert fixed text—also known as *literals*—into a number formatting string using quotation marks. For example, you could add "USD" at the end of the format string to indicate that a number's denominated in U.S. dollars. Excel automatically recognizes some characters as literals, including currency symbols, parentheses, plus (+) and minus (–) symbols, backward slashes (\), and spaces, which means you don't need to use quotation marks to have those characters appear.

Finally, the last thing you should know about Custom number format strings is that if you'd like your worksheet to display different types of values (for example, negative versus positive) differently, you can actually create a collection of four different format strings, each of which formats different types of numbers, depending on what values you type into the cell. Collectively, these four format strings tell Excel how to deal with positive values, negatives values, zero, and text values. The format strings must always appear in this order and be separated by semicolons. Here's an example:

```
#,###; [red]#,###; "---"; @
```

Excel uses the first format string (#,###) if the cell contains a positive number. Excel uses the second format string ([red]#,###) to display negative numbers. This format is the same as for positive numbers, except it displays the text in red. The third format string ("---") applies to zero values. It inserts three dashes into the cell when the cell contains the number 0. (If the cell is empty, no format string is used, and the cell remains blank.) Finally, Excel uses the last format if you enter text into the cell. The @ symbol simply copies any text into the cell as it's entered.

Tip: For a real trick, use the empty format string ;;; to puzzle friends and coworkers. This format string specifies that no matter what the content in the cell (positive number, negative number, zero, or text), Excel shouldn't display it. You can add information to the cell (and see it in the formula bar), but it doesn't appear on the worksheet or in your printouts.

Text format strings

Good news: Text format strings are extremely simple. Usually, you use a text format string to repeatedly insert the same text in a large number of cells. For example, you may want to add the word NOTE before a collection of entries. To do this, your format string needs to define the literal text you want to use—in this case, the word "NOTE"—and place the text within quotation marks (including any spaces you wish to appear). Use the @ symbol to indicate which side of the string the cell contents should go. For example, if you set the format string:

```
"NOTE: "@
```

and then you type *Transfer payment* into the cell, Excel displays it as *NOTE: Transfer payment*. However, if the cell is completely empty it stays blank, and Excel doesn't add the NOTE text.

Formatting Cell Appearance

Formatting cell values is important, because it helps maintain consistency among your numbers. But to really make your spreadsheet readable (and even beautiful), you'll need to enlist some of Excel's tools for controlling things like alignment, color, and borders and shading.

You can format a cell's appearance in two ways. You can find the button you need on the Home tab of the ribbon, or you can go back to the more comprehensive Format Cells dialog box. Just select the single cell or group of cells that you want to work with, and then choose Home→Cells→Format→Format Cells. Or, right-click the selection, and then choose Format Cells. The following sections walk you through the options in the Format Cells dialog box.

Tip: Even a small amount of formatting can make a worksheet easier to interpret by drawing the viewer's eye to important information. Of course, as with formatting a Word document or designing a web page, a little goes a long way. Don't feel the need to bury your worksheet in exotic colors and styles just because you can.

Alignment and Orientation

As you learned in the previous chapter, Excel automatically aligns cells according to the type of information you've entered. But what if this automatic alignment isn't what you want? Fortunately, in the Format Cells dialog box, the Alignment tab lets you easily change alignment as well as control some other interesting settings, like the ability to rotate text.

Excel lets you control the position of content between a cell's left and right borders, known as the *horizontal alignment*. Excel offers the following choices for horizontal alignment, some of which are shown in Figure 5-10:

- **General** is the standard type of alignment; it aligns cells to the right if they hold numbers or dates and to the left if they hold text. You learned about this type of alignment in Chapter 2.
- **Left (Indent)** tells Excel to always line up content with the left edge of the cell. You can also choose an indent value to add some extra space between the content and the left border.
- **Center** tells Excel to always center content between the left and right edges of the cell.
- **Right (Indent)** tells Excel to always line up content with the right edge of the cell. You can also choose an indent value to add some extra space between the content and the right border.
- **Fill** copies content multiple times across the width of the cell, which is almost never what you want.
- **Justify** is the same as Left if the cell content fits on a single line. When you insert text that spans more than one line, Excel *justifies* every line except the last one, which means Excel adjusts the space between words to try and ensure that both the right and left edges line up.
- **Center Across Selection** is a bit of an oddity. When you apply this option to a single cell, it has the same effect as Center. If you select more than one adjacent cell in a row (for example, cell A1, A2, A3), this option centers the value in the first cell so that it appears to be centered over the full width of all cells. However, this happens only as long as the other cells are blank. This setting may confuse you a bit at first, because it can lead to cell values being displayed over cells in which they aren't stored. Another approach to centering large text titles and headings is to use cell merging (as described in the box on page 162), but Excel purists prefer Center Across Selection, because it doesn't muck with the worksheet's structure.
- **Distributed (Indent)** is the same as Center—if the cell contains a numeric value or a single word. If you add more than one word, then Excel enlarges the spaces between words so that the text content fills the cell perfectly (from the left edge to the right edge).

Vertical alignment controls the position of content between a cell's top and bottom border. Vertical alignment becomes important only if you enlarge a row's height so that it becomes taller than the contents it contains. To change the height of a row, click the bottom edge of the row header (the numbered cell on the left side of the worksheet), and drag it up or down. As you resize the row, the content stays fixed at the bottom. The vertical alignment setting lets you adjust the cell content's positioning.

Excel gives you the following vertical alignment choices, some of which are shown in Figure 5-10:

- **Top** tells Excel that the first line of text should start at the top of the cell.
- **Center** tells Excel to center the block of text between the top and bottom border of the cell.
- **Bottom** tells Excel that the last line of text should end at the bottom of the cell. If the text doesn't fill the cell exactly, then Excel adds some padding to the top.
- **Justify** is the same as Top for a single line of text. When you have more than one line of text, Excel increases the spaces between each line so that the text fills the cell completely from the top edge to the bottom edge.
- **Distributed** is the same as Justify for multiple lines of text. If you have a single line of text, this is the same as Center.

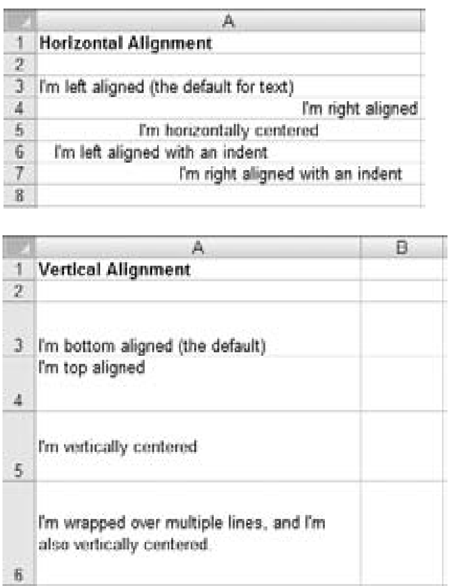

Figure 5-10:
Top: Horizontal alignment options in action.

Bottom: This sheet shows how vertical alignment and cell wrapping work with cell content.

If you have a cell containing a large amount of text, you may want to increase the row's height so you can display multiple lines. Unfortunately, you'll notice that enlarging a cell doesn't automatically cause the text to flow into multiple lines and fill the newly available space. But there's a simple solution: just turn on the "Wrap text" checkbox (on the Alignment tab of the Format Cells dialog box). Now, long passages of text flow across multiple lines. You can use this option in conjunction with the vertical alignment setting to control whether Excel centers a block of text, or lines it up at the bottom or top of the cell. Another option is to explicitly split your text into lines. Whenever you want to insert a line break, just press Alt+Enter, and then start typing the new line.

Tip: After you've expanded a row, you can shrink it back by double-clicking the bottom edge of the row header. Assuming you haven't turned on text wrapping, this action shrinks the row back to its standard single-line height.

Finally, the Alignment tab lets you rotate content in a cell up to 180 degrees, as shown in Figure 5-11. You can set the number of degrees in the Orientation box on the right of the Alignment tab. Rotating cell content automatically changes the size of the cell. Usually, you'll see it become narrower and taller to accommodate the rotated content.

Figure 5-11: *Here's a worksheet that shows one of Excel's most commonly used formatting tricks (merged cells), along with one of its most exotic (rotated text).*

Tip: You can use the Home→Alignment section of the ribbon to quickly change alignment, indenting, rotation, and wrapping, without opening the Format Cells dialog box.

FREQUENTLY ASKED QUESTION

Shrinking Text and Merging Cells So You Can Fit More Text into a Cell

I'm frequently writing out big chunks of text that I'd love to scrunch into a single cell. Do I have any options other than text wrapping?

You betcha. When you need to store a large amount of text in one cell, text wrapping is a good choice. But it's not your only option. You can also shrink the size of the text or merge multiple cells, both from the Format Cells dialog box's Alignment tab.

To shrink a cell's contents, select the "Shrink to fit" checkbox. Be warned, however, that if you have a small column that doesn't use wrapping, this option can quickly reduce your text to vanishingly small proportions.

Joining multiple cells together removes the cells' shared borders and creates one mega-sized cell. Usually, you'll do this to accommodate a large amount of content that can't fit in a single cell (like a long title that you want to display over every column). For example, if you merge cells A1, B1, and C1, you end up with a single cell named A1 that stretches over the full width of the A, B, and C columns, as shown in Figure 5-11.

To merge cells, select the cells you want to join, choose Home→Cells→Format→Format Cells, and then, on the Alignment tab, turn on the "Merge cells" checkbox. There's no limit to how many cells you can merge. (In fact, you can actually convert your entire worksheet into a single cell if you want to go crazy.) And if you change your mind, don't worry—you simply need to select the single merged cell, choose Home→Cells→Format→Format Cells again, and then turn off the "Merge cells" checkbox to redraw the original cells.

Fonts and Color

As in almost any Windows program, you can customize the text in Excel, applying a dazzling assortment of colors and fancy typefaces. You can do everything from enlarging headings to colorizing big numbers. Here are the individual font details you can change:

- **The font style**. (For example, Arial, Times New Roman, or something a little more shocking, like Futura Extra Bold.) Calibri is the standard font for new worksheets. If you have an old-school workbook created by Excel 2003, you'll notice that it uses 10-point Arial instead.
- **The font size, in points**. The standard point size is 11, but you can choose anything from a minuscule 1-point to a monstrous 409-point. Excel automatically enlarges the row height to accommodate the font.
- **Various font attributes, like italics, underlining, and bold**. Some fonts have complimentary italic and bold typefaces, while others don't (in which case Windows uses its own algorithm to make the font bold or italicize it).
- **The font color**. This option controls the color of the text. (Page 167 covers how to change the background fill color in a cell.)

To change font settings, first highlight the cells you want to format, choose Home→Cells→Format→Format Cells, and then click the Font tab (Figure 5-12).

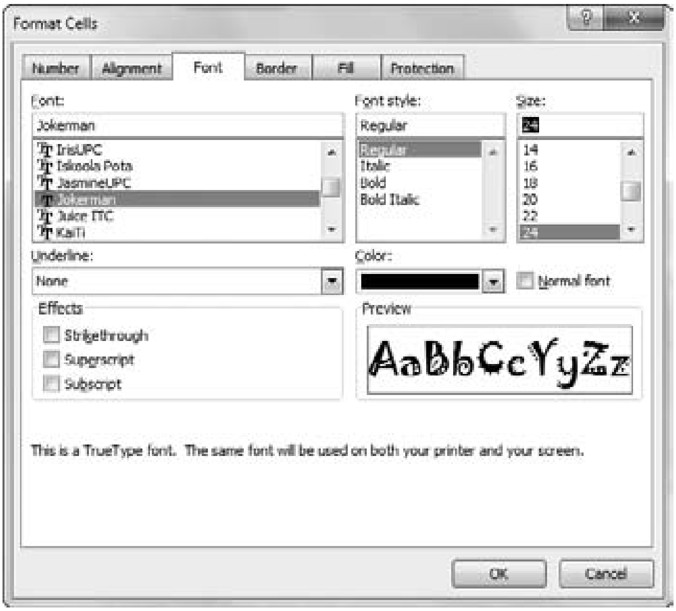

Figure 5-12: *Here's an example of how to apply an exotic font through the Format Cells dialog box. Keep in mind that when displaying data, and especially numbers, sans-serif fonts are usually clearer and look more professional than serif fonts. (Serif fonts have little embellishments, like tiny curls, on the ends of the letters; sans-serif fonts don't.) Calibri, the default spreadsheet font, is a sans-serif font. Adobe Minion Pro, the font used for the body text of this book, is clearly a serif font, which works best for large amounts of text.*

Tip: Thanks to Excel's handy Redo feature, you can repeatedly apply a series of formatting changes to different cells. After you make your changes in the Format Cells dialog box, simply select the new cell you want to format in the same way, and then hit Ctrl+Y to repeat the last action.

Rather than heading to the Format Cells dialog box every time you want to tweak a font, you can use the ribbon's handy shortcuts. The Home→Font section displays buttons for changing the font and font size. You also get a load of tiny buttons for applying basics like bold, italic, and underline, applying borders, and changing the text and background colors. (Truth be told, the formatting toolbar is way more convenient for setting fonts, because its drop-down menu shows a long list of font names, whereas the font list in the Format Cells dialog box is limited to showing an impossibly restrictive six fonts at a time. Scrolling through that cramped space is like reading the phone book on index cards.)

Without a doubt, the ribbon's most useful formatting feature is *live preview*, a frill that shows you the result of a change *before* you've even applied it. Figure 5-13 shows live preview in action.

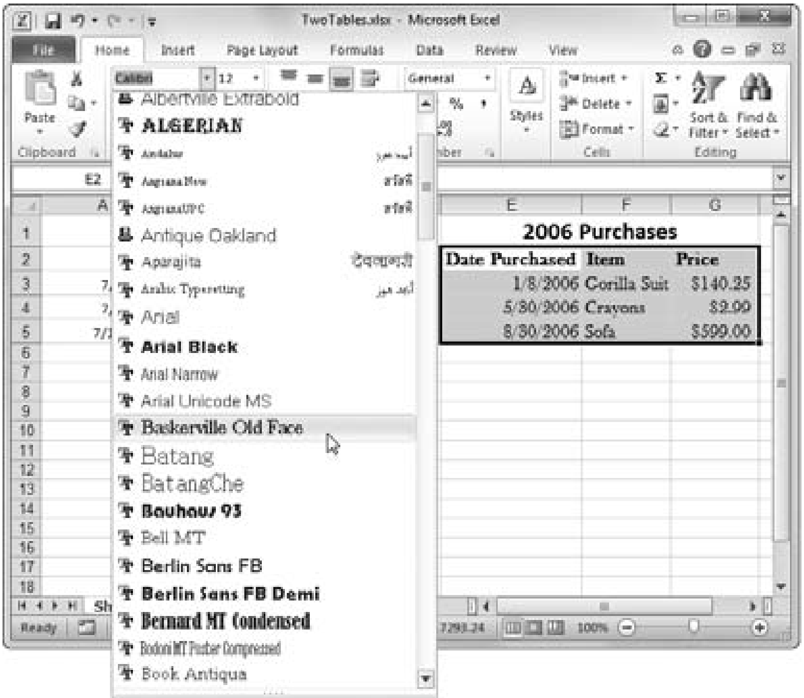

***Figure 5-13:** Right now, this spreadsheet's creator is just thinking about using the stylish Baskerville font for this table. However, the moment she hovers over Algerian (higher up in the font list), Excel switches the currently selected cells on the worksheet to that font, providing a preview of the change. The best part: When she moves the mouse pointer away, the formatting disappears instantaneously. To make the changes stick, all she needs to do is click the font. This live preview feature works with font names, font sizes, and colors.*

Note: No matter what font you apply, Excel, thankfully, always displays the cell contents in the formula bar in easy-to-read Calibri font. That makes things easier if you're working with cells that've been formatted using difficult-to-decipher script fonts, or really large or small text sizes.

Special characters

Most fonts contain not only digits and the common letters of the alphabet, but also some special symbols that you can type directly on your keyboard. One example is the copyright symbol ©, which you can insert into a cell by entering the text *(C)* and letting AutoCorrect do its work. Other symbols, however, aren't as readily available. One example is the special arrow character →. To use this symbol, you'll need the help of Excel's symbols. Simply follow these steps:

1. **Choose Insert→Symbols→Symbol.**

   The Symbol dialog box opens, as shown in Figure 5-14. Now it's time to hunt for the symbol you need.

POWER USERS' CLINIC

Formatting Individual Characters

The ribbon lets you perform one task that you can't with the Format Cells dialog box: applying formatting to just a part of a cell. For example, if a cell contains the text "New low price", you could apply a new color or bold format to the word "low."

To apply formatting to a portion of a cell, follow these steps:

1. **Move to the appropriate cell, and then put it into edit mode by pressing F2**.

   You can also put a cell into edit mode by double-clicking it, or by moving to it, and then clicking the text inside the formula bar.

2. **Select the text you want to format**.

   You can select the text by highlighting it with the mouse, or by holding down Shift while using the arrow keys to mark your selection.

3. **Choose a font option from the ribbon's Home Font section**.

   You can also change the size, the color, or the bold, italic, or underline settings. And if you don't want to waste time choosing the Home tab if you're currently somewhere else in the ribbon, then you can simply right-click the selected text to show a pop-up toolbar with font options.

Applying multiple types of text formatting to the same cell can get tricky. The formula bar doesn't show the difference, and, when you edit the cell, you may not end up entering text in the font you want. Also, be careful that you don't apply new font formatting to the cell later; if you do, you'll wipe out all the font information you've added to the cell.

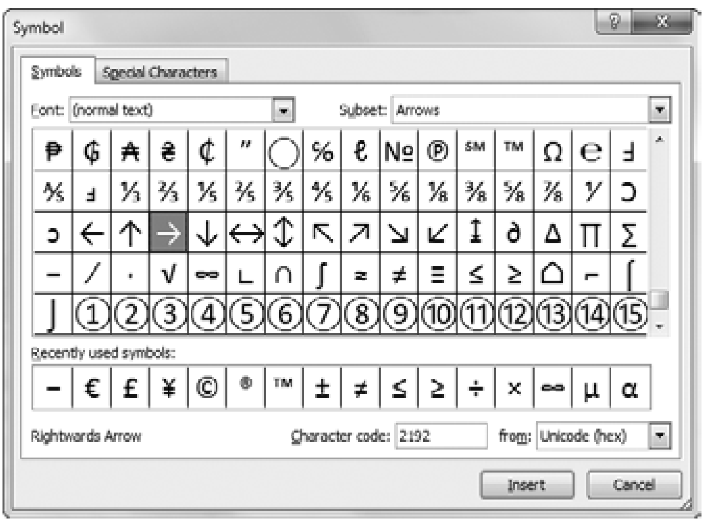

Figure 5-14: *The Symbol dialog box lets you insert one or more special characters. You can choose extended characters that are supported by most fonts (like currency symbols, non-English letters, arrows, and so on). Alternatively, you can use a font that's all about fancy characters, like the Wingdings font, which is chock full of tiny graphical icons.*

2. **Choose the font and subset (the group of symbols you want to explore).**

   If you're looking for a fairly common symbol (like a mathematical sign, an arrow, an accented letter, or a fraction), you probably don't need to change the font. In the Font box, keep the automatic selection of "(normal text)", and then, from the Subset box at the right, choose the type of default. For example, choose the Arrows subset to see arrow symbols that point in different directions.

   If you want funkier alternatives, choose a fancy font from the Font box on the left. You should be able to find at least one version of the Wingdings font in the list. Wingdings has the most interesting symbols to use. It's also the most likely to be on other people's computers, which makes a difference if you're planning to email your worksheet to other people. If you get your symbols from a really bizarre font that other people don't have, they won't be able to see your symbols.

Note: Wingdings is a special font included with Windows that's made up entirely of symbols like happy faces and stars, none of which you find in standard fonts. You can try to apply the Wingdings font on your own (by picking it from the font list), but you won't know which character to press on your keyboard to get the symbol you want. You're better off using Excel's Symbol dialog box.

3. **Select the character, and then click Insert.**

   Alternatively, if you need to insert multiple special characters, just double-click each one; doing so inserts each symbol right next to each other in the same cell without having to close the window.

Tip: If you're looking for an extremely common special character (like the copyright symbol), you can shorten this whole process. Instead of using the Symbols tab, just click over to the Special Characters tab in the Symbol dialog box. Then, look through the small list of commonly used symbols. If you find what you want, just select it, and then click Insert.

There's one idiosyncrasy that you should be aware of if you choose to insert symbols from another font. For example, if you insert a symbol from the Wingdings font into a cell that already has text, then you actually end up with a cell that has two fonts—one for the symbol character and one that's used for the rest of your text. This system works perfectly well, but it can cause some confusion. For example, if you apply a new font to the cell after inserting a special character, Excel adjusts the entire contents of the cell to use the new font, and your symbol changes into the corresponding character in the new font (which usually isn't what you want). These problems can crop up any time you deal with a cell that has more than one font.

On the other hand, if you kept the font selection on "(normal text)" when you picked your symbol, then you won't see this behavior. That's because you picked a more commonplace symbol that's included in the font you're already using for the cell. In this case, Excel doesn't need to use two fonts at once.

Note: When you look at the cell contents in the formula bar, you always see the cell data in the standard Calibri font. This consistency means, for example, that a Wingdings symbol doesn't appear as the icon that shows up in your worksheet. Instead, you see an ordinary letter or some type of extended non-English character, like æ.

Borders and Fills

The best way to call attention to important information isn't to change fonts or alignment. Instead, place borders around key cells or groups of cells and use shading to highlight important columns and rows. Excel provides dozens of different ways to outline and highlight any selection of cells.

Once again, the trusty Format Cells dialog box is your control center. Just follow these steps:

1. **Select the cells you want to fill or outline.**

   Your selected cells appear highlighted.

2. **Select Home→Cells→Format→Format Cells, or just right-click the selection, and then choose Format Cells.**

   The Format Cells dialog box appears.

3. **Head directly to the Border tab. (If you don't want to apply any borders, skip straight to step 4.)**

   Applying a border is a multistep process (see Figure 5-15). Begin by choosing the line style you want (dotted, dashed, thick, double, and so on), followed by the color. (Automatic picks black.) Both these options are on the left side of the tab. Next, choose where your border lines are going to appear. The Border box (the square that contains the word "Text") functions as a nifty interactive test canvas that shows you where your lines will appear. Make your selection either by clicking one of the eight Border buttons (which contain a single bold horizontal, vertical, or diagonal line), or click directly inside the Border box. If you change your mind, clicking a border line makes it disappear.

   For example, if you want to apply a border to the top of your selection, click the top of the Border box. If you want to apply a line between columns inside the selection, click between the cell columns in the Border box. The line appears indicating your choice.

Tip: The Border tab also provides two shortcuts in the tab's Presets section. If you want to apply a border style around your entire selection, select Outline after choosing your border style and color. Choose Inside to apply the border between the rows and columns of your selection. Choosing None removes all border lines.

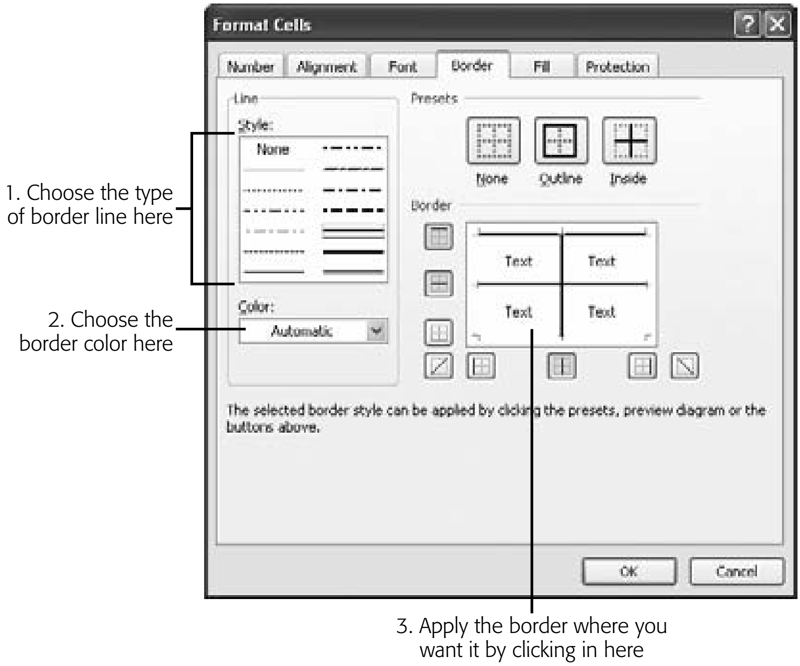

Figure 5-15: *Follow the numbered steps in this figure to choose the line style and color, and then apply the border. In this picture, Excel will apply a solid border between the columns and at the top edge of the selection.*

4. **Click the Fill tab.**

   Here you can select the background color, pattern color, and pattern style to apply shading to the cells in the selection (see Figure 5-16). Click the No Color box to clear any current color or pattern in the selected cells. When picking a pattern color, you may notice that certain colors are described as *theme colors*. These theme colors are a set of coordinated colors that change whenever you pick a new theme for your workbook, as described in on page 174.

   To get a really fancy fill, you can use a *gradient*, which is a blend of two colors. For example, with gradients you can create a fill that starts out white on one side of a cell and gradually darkens to blue on the other. To use a gradient fill, click the Fill Effects button, and then follow the instructions in Figure 5-17.

5. **Click OK to apply your changes.**

   If you don't like the modifications you've just applied, you can roll back time by pressing Ctrl+Z to trigger the indispensable Undo command.

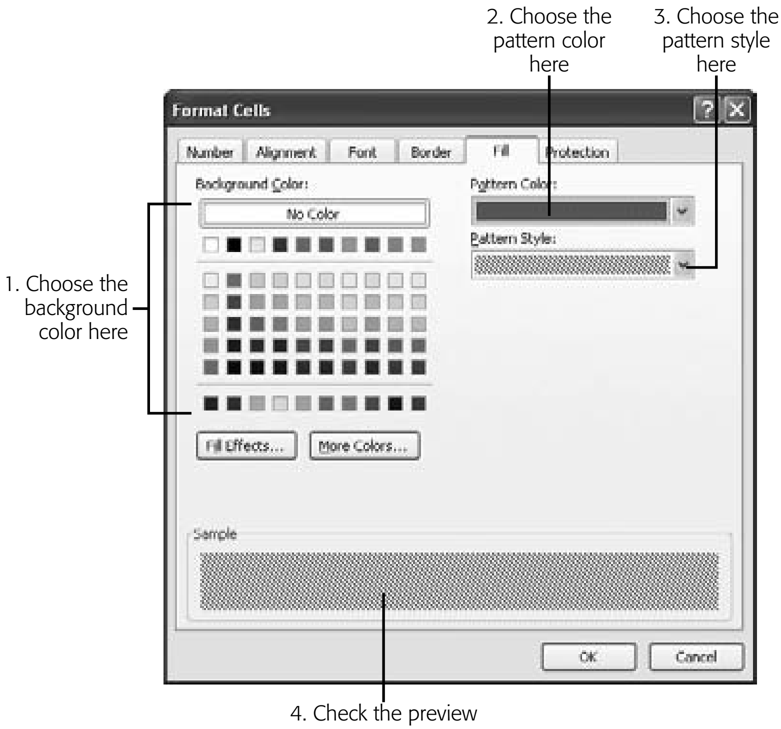

Figure 5-16: *Adding a pattern to selected cells is simpler than choosing borders. All you need to do is select the colors you want and, optionally, choose a pattern. The pattern can include a grid, dots, or the diagonal lines shown in this figure.*

Tip: You can remove a worksheet's gridlines, which is handy when you want to more easily see any custom borders you've added. To remove gridlines, select View→Show→Gridlines. (This action affects only the current file, and won't apply to new spreadsheets.)

Drawing Borders by Hand

If you need to add a border around a cell or group of cells, the Format Cells dialog box's Border tab does the trick (see Figure 5-15). However, you could have a hard time getting the result you want, particularly if you want to add a combination of different borders around different cells. In this situation, you have a major project on your hands that requires several trips back to the Format Cells dialog box.

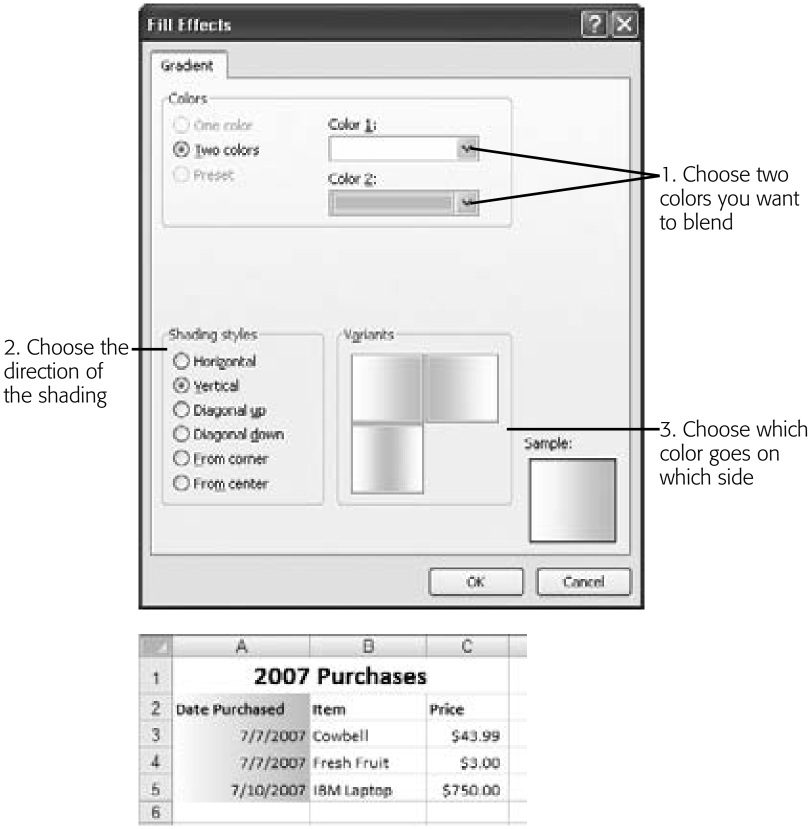

Figure 5-17:
Top: To create a gradient, you need to pick the two colors you want to use to create the blend, and you need to choose the way Excel does the blending (from one side to another, from the top to the bottom, and so on). When applying a gradient fill to a stack of cells, a vertical fill makes the most sense, because that way the gradients in each cell line up and look like one seamless shaded region. When applying a gradient fill to a row of cells, a horizontal fill looks better, for the same reason.

Bottom: A gradient fill on cells A2 to A5.

Fortunately, there's a little-known secret that lets you avoid the hassle: Excel's Draw Border feature. The Draw Border feature lets you draw border lines directly on your worksheet. This process is a little like working with a painting program. You pick the border style, color, and thickness, and then you drag to draw the line between the appropriate cells. When you draw, Excel applies the formatting settings to each affected cell, just as if you'd used the Borders tab.

Here's how it works:

1. **Look in the ribbon's Home→Font section for the border button.**

   The name of the border button changes to reflect whatever you used it for last. You can most easily find it by its position, as shown in Figure 5-18.

2. **Click the border button, choose Line Style, and then pick the type of line you want.**

   You can use dashed and solid lines of different thicknesses, just as you can in the Format Cells dialog box's Borders tab.

3. **Click the border button, choose Line Color, and then pick the color you want.**

   Now you're ready to start drawing.

4. **Click the border button, and then choose Draw Border.**

   When you choose Draw Border, your mouse pointer changes into a pencil icon.

5. **Using the border pencil, click a grid line where you want to place your border (Figure 5-19).**

   You can also drag side to side or up and down to draw a longer horizontal or vertical line. And if you drag your pointer down *and* to the side, you create an outside border around a whole block of cells.

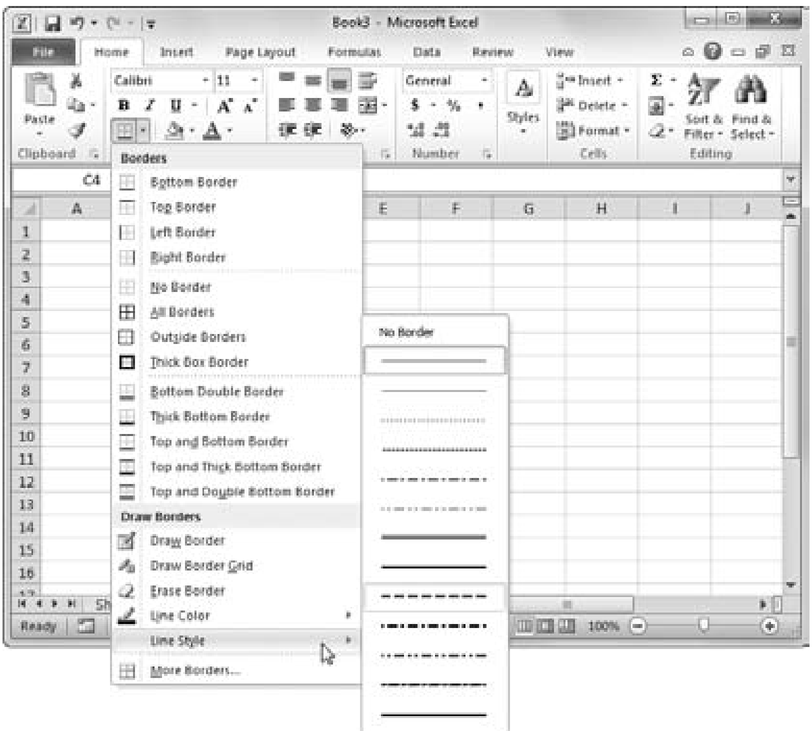

Figure 5-18: *The border button is at the bottom left of the Home→Font section. When you click it, you see a list of commands. Before you draw any borders, it makes sense to customize the border style. For example, you could choose Line Style, as shown here, and for color, you'd choose Line Color.*

Figure 5-19:
Here, a double-line border is being drawn between column A and column B.

6. **To stop drawing, head back to the border menu, and then choose Draw Border again.**

   If you make a mistake, you can even use an eraser to tidy it all up. Just click the border button, and then choose Erase Border. The mouse pointer changes to an eraser. Now you can click the border you want to remove.

Tip: If you don't want to use the Draw Border feature, you can still make good use of the border button. Just pick a line style and line color, select some cells, and then choose an option from the border menu. For example, if you pick Bottom Border, Excel applies a border with the color and style you chose to the bottom of the current cell selection.

DESIGN TIME

A Designer Worksheet

Cells aren't the only part of a worksheet you can tweak. A little-known feature in Excel lets you change the appearance of the entire worksheet by applying a custom *picture* as a background. Just select Page Layout→Page Setup→Background, and then choose a picture file. The picture can be in just about any format, including GIF, JPEG, BMP, or TIF, and Excel spreads it like tiles across your worksheet surface to fill the whole working area. To try this trick out, use some of the background pictures located in your windows folder. (These files usually have names like Feather Texture and Sante Fe Stucco.)

The designer background feature's really just for fun and Excel doesn't include it in the printout. To remove a background, just choose Page Layout→Page Setup→Delete Background.

CHAPTER
6

Smart Formatting Tricks

In the previous chapter, you took a comprehensive tour of Excel's formatting fundamentals. But of course, just because the features exist doesn't mean they're easy to use. Digging through the different options and applying a full range of formatting settings can be a tedious task. Fortunately, Excel also includes a few timesavers that let you speed up many formatting jobs.

In this chapter, you'll try out the essential formatting techniques that every Excel guru loves. They include:

- **The Format Painter**, which provides a quick and dirty way to transfer formatting from one cell to another.
- **Styles**, which let you standardize your favorite formatting choices so you can use them again.
- **Themes**, which give you a toolkit with a collection of ready-to-use styles that can jazz up the dullest worksheet.
- **Conditional formatting**, which gets Excel to do the hard work of finding values you're interested in and then highlighting them with custom formatting.

Once you master these four timesavers, you'll have the secret to making great-looking worksheets.

The Format Painter

The Format Painter is a simple yet elegant tool that lets you copy all of a cell's format settings—including fonts, colors, fill, borders, and even the number format—from one cell to another. (Apparently, the Excel team decided that the more accurate label "Format Copier" wasn't nearly as exciting as the name Format Painter.)

To use the Format Painter, follow these steps:

1. **Move to a cell that has the formatting you want to copy.**

   You can use the Format Painter to copy formatting from either one cell or a whole group of cells. For example, you could copy the format from two cells that use different fill colors, and paste that format to a whole range of new cells. The new cells would alternate between the two fill colors. Although this is a powerful trick, in most cases, it's easiest to copy the format from a single cell.

2. **Choose Home→Clipboard→Format Painter to switch into "format painting" mode.**

   The pointer changes so that it now includes a paintbrush icon, indicating that Excel is ready to copy the format.

3. **Click the cell where you want to apply the format.**

   The moment you release your mouse button, Excel applies the formatting and your pointer changes back to its normal appearance. If you want to copy the selected format to several cells at once, just drag to select a group of cells, rows, or columns, instead of clicking a single cell.

Excel doesn't let you get too carried away with format painting—as soon as you copy the format to a new cell or selection, you exit format painting mode. If you want to copy the desired format to another cell, you must backtrack to the cell that has your format, and start over again. However, there's a neat trick you can use if you know you're going to repeatedly apply the same format to a bunch of different cells. Instead of single-clicking the Format Painter button, double-click it. You'll remain in format painting mode until you click the Format Painter button again to switch it off.

Note: The Format Painter is a good tool for quickly copying formatting, but it's no match for another Excel feature called *styles*. With styles, you can define a group of formatting settings, and then apply them wherever you need them. Best of all, if you change the style after you've created it, Excel automatically updates all cells that you've formatted using that style. The next section describes styles in more detail.

Styles and Themes

Styles let you create a customized collection of format settings, give that collection a name, and store it in a spreadsheet file. You can then apply these settings anywhere you need them. For example, you could create a style called Great Big Header that uses the Cambria font, pumps up the font size to 46 points, and colors the text bright red.

Every Excel spreadsheet starts off with a collection of prebuilt styles. Microsoft designed these styles with two goals in mind: to give you quick access to most common and practical formatting choices, and to make great looking documents. To take a look at the styles waiting for you, choose Home→Styles→Cell Styles. Figure 6-1 shows the gallery of options that you'll see.

Note: If your Excel window is very wide, the Home→Styles→Cell Styles button disappears, and a scrollable list of styles in the Home→Styles section is visible instead. This change saves you a click, but doesn't alter the way that styles work.

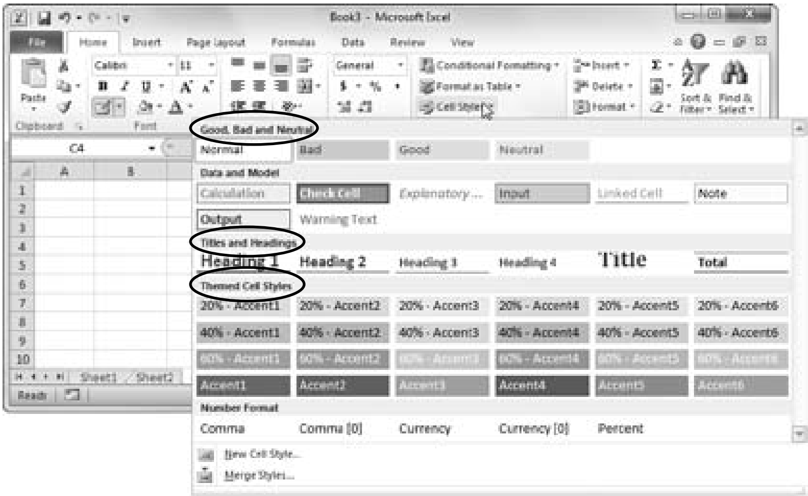

Figure 6-1: *Excel divides its built-in styles into separate categories, according to how you might use them. The Good, Bad, and Neutral category lets you separate the good news from the bad using the carefully shaded versions of the universal colors red, yellow, and green. The Titles and Headings category adds border formatting (page 167) to make great titles. And the Themed Cell Styles category gives you a range of differently colored, differently shaded cells that are chosen to match harmoniously with one another based on the current workbook theme (page 174).*

You can apply more than one style to the same cell to get a combination of formatting options. For example, you could use the Currency style to get the right number format, and then pick the Bad style to flag a huge debt with a light red background fill. (Bad is simply the name of a prebuilt style that applies a light red background fill and a dark red font color.) If you apply more than one style and they conflict (for example, both styles use a different background color), the style you applied last takes over.

Styles use Excel's live preview feature, which gives you try-before-you-buy formatting. When you select a group of cells and then hover over one of the styles in the ribbon, your selected cells change instantaneously to reflect that style. Run your mouse over the different style options, and you see a quickly changing series of different formatting choices. To actually apply a style, click it.

Note: Longtime Excel fans will recognize the Number Format group of styles shown at the bottom of Figure 6-1. These styles are the only ones that Excel included in previous versions. The Number Format styles simply apply a different number format (page 141), like Currency, Percentage, and so on. They don't change other appearance details. However, they're still useful. For example, you could use Currency style for all the dollar figures in your worksheet. Then, at some later point, you could modify that style to set a different currency symbol, alignment, or number of decimal places. This one change updates every cell that uses the Currency style–in this case, all the prices across your entire worksheet.

If this was all that styles offered—a handy way to reuse good-looking formatting presets—they'd still be quite useful. But wait, there's more: Excel styles come with two more invaluable features. First, you can create your own styles to reuse your own favorite formatting choices. Second, you can use *themes* to swap one set of styles for another. You'll learn how both these techniques work in the following sections.

TIMESAVING TIP

A Quicker Way to Apply Styles

The ribbon makes it fairly easy to work with styles. However, sometimes you may be in the middle of working with another ribbon tab and find it's just too inconvenient to jump back to the Home tab.

If you're in this situation, you can make your life more pleasant by adding the style gallery to the Quick Access toolbar. To do so, choose Home→Styles→Cell Styles, right-click any style, and then choose Add Gallery to Quick Access Toolbar. This adds a Cell Styles button to the Quick Access toolbar, where it's available at all times. Click this button, and you'll see the familiar style gallery.

If you get tired of the Cell Styles button on the Quick Access toolbar, right-click it, and then choose Remove from Quick Access Toolbar.

Custom Styles

Styles really shine in complex worksheets where you need to apply different formatting to different groups of cells. For example, say you've got a worksheet that's tracking your company's performance. You're confident that most of the data is reliable, but know that a few rows come from your notoriously hopeful sales department. To highlight these sales projections, you decide to use a combination of a bold font with a hot pink fill. And since these figures are estimated and aren't highly precise, you decide to use a number format without decimal places and precede the number with a tilde (~), the universal symbol for "approximately right."

You *could* implement all these changes manually, but that'll take fourscore and seven years. Better to set up a style that includes all these settings, and then apply it with a flick of the wrist whenever you need it. Styles are efficiency monsters in a few ways:

- They let you reuse your formatting easily, just by applying the style.
- They free you from worry about being inconsistent, because the style includes all the formatting you want.

- Excel automatically saves styles with your spreadsheet file, and you can transfer styles from one workbook to another.
- If you decide you need to change a style, you need just a few mouse clicks. Then, Excel automatically adjusts every cell that uses the modified style.

Here's how you create a new style:

1. **Begin by moving to a cell in your worksheet that has the formatting you want to use for your style.**

   The quickest way to create a new style is by using formatting you've already set up. However, you can also create a new style from scratch. In this case, you'd simply move to a blank, unformatted cell in your worksheet.

2. **Select Home→Styles→Cell Styles→New Cell Style.**

   The Style dialog box appears, which lets you design your own styles.

3. **In the "Style name" box, type a name for your new style.**

   For example, if you want to create a new style for column titles, enter the style name ColumnTitle. Each style name needs to be unique in your workbook.

4. **Choose the style options you want to set.**

   Styles don't need to set every aspect of a cell. For example, you might want to create a style that applies new font and fill settings, but keeps the current number format (page 141), alignment, and border details. When creating this style, you'd clear the Number, Alignment, and Border checkboxes, so these details aren't included in your style. Figure 6-2 shows an example.

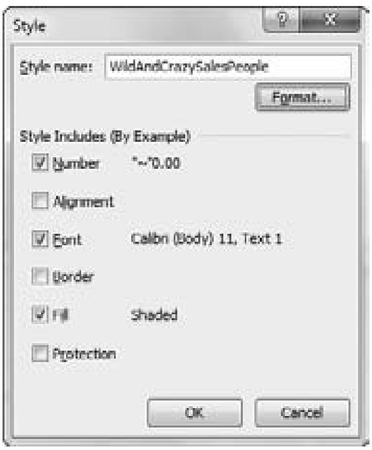

Figure 6-2:
Here, a new style, WildAndCrazySalesPeople, is about to be created. This style defines a number format as well as font and fill settings. If you don't want your style to include some of these settings, turn off the checkmark in the appropriate checkboxes. For example, if you want to create a style that applies a new font, fill, and border, but you want to keep the existing alignment and number format, turn off the Number and Alignment checkboxes. As a general rule, if you don't need to explicitly set a specific style characteristic, turn off the corresponding checkmark.

Note: When you select many style checkboxes, you create a pumped-up style that does a lot at once. When you select only one or two style checkboxes, you create a less powerful style, but one that's more flexible. You can use it in a variety of cells that already have customized formatting, without changing the formatting characteristics you want to keep.

5. **Click Format to specify the formatting options for the style.**

   When you click Format, the familiar Format Cells dialog box appears. Use this dialog box to change the formatting just as if you were formatting an individual cell. Click OK to close the Format Cells dialog box when you're finished.

6. **Click OK to close the Style window.**

   Once you've created a style, applying it is just a matter of a few mouse clicks. Select the cell or cells you want to modify, choose Home→Styles→Cell Styles, and then choose your style from the list (Figure 6-3).

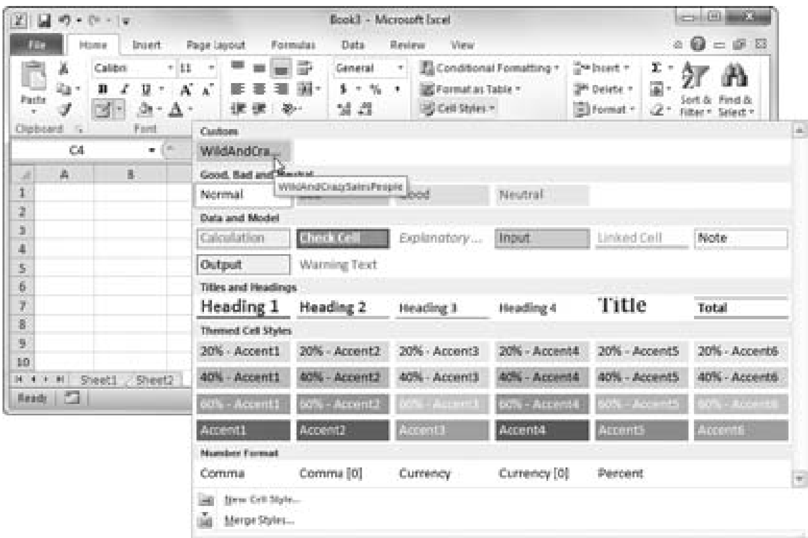

***Figure 6-3:** The styles you create appear in a Custom group at the top of the list. (If you haven't created any styles, you won't see the Custom group at all.)*

Modifying Styles

Keep in mind that you can still modify the formatting of a cell after you've applied a style. But if you do find yourself overriding a style fairly frequently, and always in the same way, it probably indicates that your style isn't quite right. Either create more than one version of the same style, each with the appropriate settings, or clear some of the style checkboxes so that your style doesn't apply formatting settings that you commonly change:

- To modify a style, choose Home→Styles→Cell Styles, find the style you want in the gallery, right-click it, and then choose Modify. You wind up back at the familiar Style dialog box (Figure 6-2), where you can tweak it to your heart's content. You can use this approach to revise your own custom styles, or to change the built-in styles that Excel adds to every workbook.
- To duplicate a style, choose Home→Styles→Cell Styles, find the style you want in the gallery, right-click it, and then choose Duplicate. The Styles dialog box appears with the style you've chosen. Change the formatting as desired, choose a better name, and then click OK to add this style to the Custom category.
- To delete a style you don't want anymore, choose Home→Styles→Cell Styles, find the style, right-click it, and then choose Delete. Excel removes the style formatting from all the cells that used it.

Transferring Styles Between Workbooks

Once you've created a few useful styles, you'll probably want to reuse them in a variety of spreadsheet files. In order to reuse them, you need to copy the style information from one workbook to another. Excel makes this process fairly straightforward:

1. **Open both files in Excel.**

   You need both the source workbook (the one that has the styles you want to copy) and the destination workbook (the one where you want to copy the styles).

2. **Go to the destination workbook.**
3. **Choose Home→Styles→Cell Styles→Merge Styles.**

   The Merge Styles dialog box appears with a list of all files that you currently have open in Excel.

4. **Select the file that has the styles you want to copy into your active workbook, and then click OK.**

   If two or more styles have the same name, Excel prompts you with a warning message, informing you that it will overwrite the current styles with the styles you're importing. Click OK to continue.

   You can now use the styles that you've imported. These styles are now an independent copy of the styles in the source workbook. If you change the styles in one workbook, you won't affect the other workbook unless you merge the changed styles back into it.

Tip: If you want to automatically include your styles in new workbooks, consider creating a template that includes these styles. You can then create new workbooks based on that template. Chapter 16 has the full story.

TIMESAVING TIP

Turning Off Live Preview

Most of the time, live preview is a great way for wishywashy spreadsheet writers to see formatting possibilities without committing. However, if you're using a heavily formatted workbook, you might find the live preview feature slows you down when you're scrolling through a lot of options. In this case, it might make sense to turn off live preview.

To turn it off, choose File→Options, and then pick the General section. Under the "User Interface options," turn off the Enable Live Preview setting, and then click OK. Now you can zip around the ribbon, but you need to actually *apply* a formatting change (by clicking the appropriate button in the ribbon) before you see what it looks like.

Themes: A Package of Styles

As nice as the prebuilt styles are, they don't suit everyone. For example, the standard style colors favor subdued shades of red, gray, brown, and green, which make sense for the company accountant but aren't the most exciting choice for an urban hipster. To jazz things up, you can choose a different *theme* that features livelier colors. When you do, your entire worksheet gets an immediate facelift—you don't need to track down each individual cell and reformat it.

Technically, a theme is a combination of three ingredients:

- **Fonts**. Every theme has one font that's used for headings and another one that's used for everything else. These two fonts might be different sizes of the same typeface, or two complementary typefaces.
- **Colors**. Every theme has a palette of 12 complementary colors. The cell styles that appear under the Themed Cell Styles heading (see Figure 6-1) draw upon these colors for text and background fills. Best of all, these colors don't reflect the preferences of Cheeto-munching programmers. Instead, bona-fide artsy types chose them—in this case, professional designers on the Microsoft payroll.
- **Effects**. Effects are fine alterations that pretty up shapes and other hand-drawn graphics that you can create with Excel's drawing tools (Chapter 19). If you don't have any shapes on your worksheet, the effect settings don't do anything.

To choose a theme, choose Page Layout→Themes→Themes to see a gallery of choices (Figure 6-4).

The secret to understanding themes is realizing how changing the theme affects your worksheet. In other words, how does Excel apply the theme's fonts, colors, and effects to your worksheet? The following sections break it down.

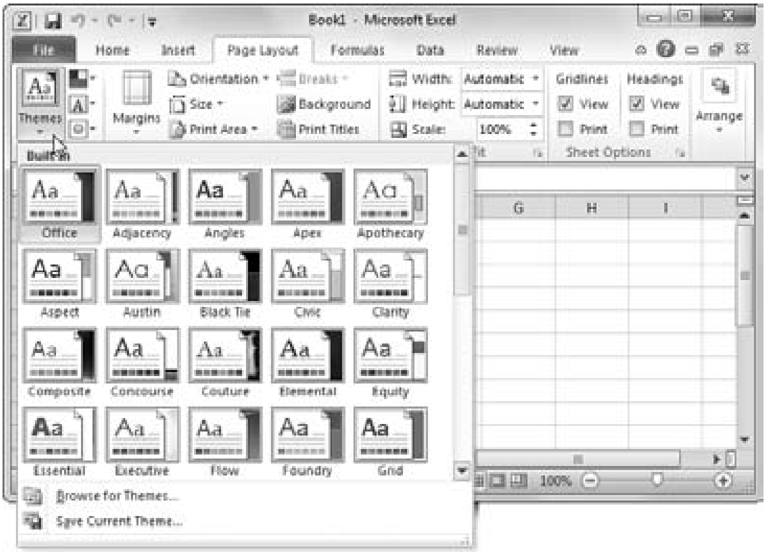

Figure 6-4: *Every workbook begins using the crowd-pleasing Office theme, but you have a long list of other options. You can even search Microsoft's Office Online website for hot new themes. As you hover over a new theme, your workbook adjusts itself automatically, thanks to the magic of Excel's live preview feature.*

Fonts

Every workbook has a standard *body font* that it uses in every cell. Excel uses this standard font unless you explicitly choose a different font using the ribbon's Home→Font section or the Format Cells dialog box.

In a brand new Excel spreadsheet, everything you type starts out in easy-on-the-eyes 11-point Calibri font. If you apply a new theme, you get a new standard font. For example, switch to the traditionally styled Apex theme, and you'll get the elegant Book Antiqua font instead.

Note: All the fonts used in Excel themes are installed as part of Microsoft Office.

The same sort of magic works with the *heading font*, but it's more limited—in fact, the only style that uses the heading font is Title. If you use the Title cell style and switch from one theme to another, Excel updates your cell to use the heading font from that style. In some themes, the heading font is the same as the standard body font (as it is for the standard Office theme). In other themes, these two fonts are different, but complementary. For example, the Apex font uses the stylish Lucida Sans font for all titles.

Note: You might assume that the heading styles (Heading 1, Heading 2, Heading 3, and so on) use the heading font. Oddly enough, that's not how it works. All the heading styles use the body font. Title is the only style that uses the heading font.

If you're feeling a bit reckless, you can override the font that Excel automatically uses for all new workbooks. To override it, select File→Options, and then choose the General section. Under the "When creating new workbooks" heading is a "Use this font" and a "Font size" setting where you can set the standard font and font size. Ordinarily, the automatic font isn't set to a specific font at all—instead, it's set to the special value Body Font. This tells Excel to apply the standard font from the current theme. Usually, this is the choice you want, because it lets you quickly adapt your entire spreadsheet to a theme of your choosing.

Colors

Every theme relies on 12 key colors. When you move from one theme to another, Excel swaps in the new set of colors. Excel alters any place where you've used the 12 theme colors. However, *other* colors aren't affected.

Note: Although each theme has only 12 base colors, Excel varies the saturation of the color to make it bolder or lighter, based on the style you use. For example, the Office theme includes a steel blue color that you can use at full strength (with the style named Accent 1) or lighten to a faint gray-blue mist (with the style named 20% - Accent 1).

To make this system a bit clearer, imagine that a designer runs amok, formatting cells with different background fills. He fills some of the cells with theme colors, and others with custom colors. (Figure 6-5 shows the difference.) When you switch themes, the cells that have the theme colors are changed, while the other cells aren't.

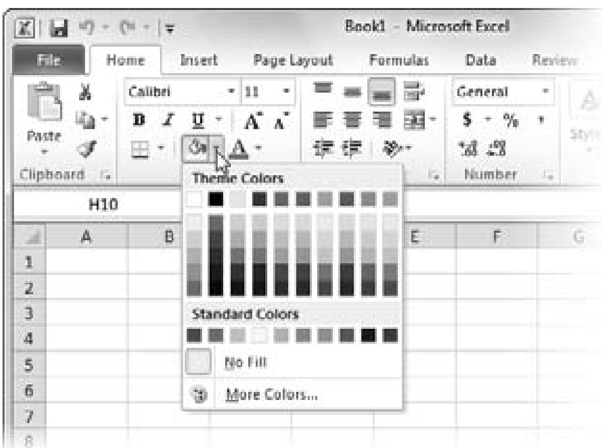

Figure 6-5: *When you set the background fill in a cell, you can pick one of the theme colors (at different saturations), you can use a standard color (which gives you the standard red-green-blue lineup), or you can click More Colors to pick a custom color of your own design. You have the same choices when picking the foreground color for your text.*

Tip: You're always better off using theme colors rather than picking a new custom color. That way you can give your workbook a facelift simply by switching from one theme to another, and the colors will still match. Also, if you choose custom colors that look nice with a specific theme, they're likely to clash horribly when you change to another theme.

Experienced Excel workers rarely waste time picking background and foreground colors out of the ribbon. Instead, they use styles. Any time you use one of the styles from the Themed Cell Styles category (Home→Styles→Cell Styles; see Figure 6-1), you're applying a theme-specific color. As a result, if you pick another theme, all the themed cell styles change to use the new color.

The theme system works well because each color in a theme plays a specific role. In other words, some colors are intended for text, while others are designed to play the role of a complementary background; a few more add eye-catching accents. To see the intended purpose of each color, hover over it in the ribbon (Figure 6-6).

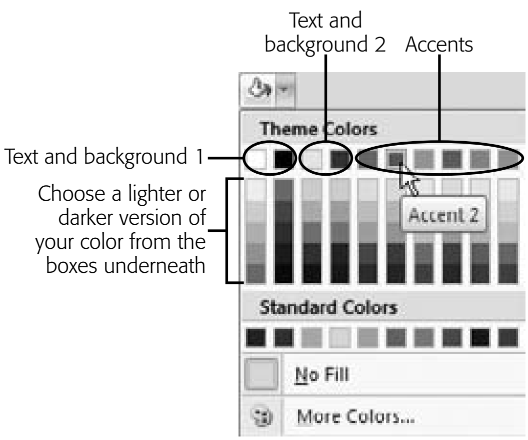

Figure 6-6: *Here, the mouse is over one of the six accent colors that are present in every theme. You can also choose from four colors that are intended for standard text and background fills (on the left).*

Effects

Effects are the simplest part of any theme, because Excel applies them with no work on your part (all you need to do is switch themes, and then the Effects kick in). Excel automatically applies effects to any graphics you've created. You'll learn more about creating shapes and other illustrations in Chapter 19.

Modifying Themes

Excel gives you the ability to use just part of a theme. For example, what if you want the casual fonts from the Office theme but want to use them in conjunction with the rich reds of the Opulent theme? The Excel designers might cringe at your combination, but there's no reason for you to hesitate—from the Page Layout→Themes→Themes gallery, just choose the Office theme, and then, from the Page Layout→Themes→Colors list, pick the Opulent colors (Figure 6-7).

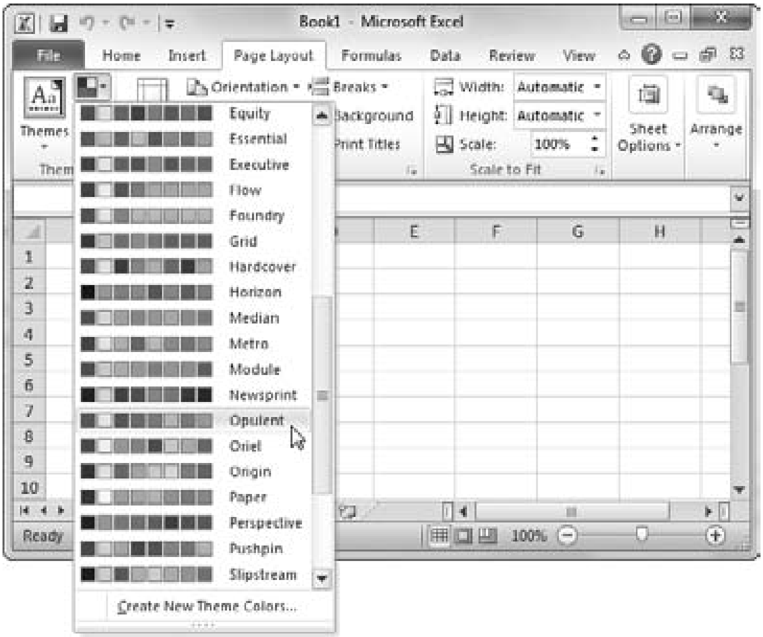

Figure 6-7: *Using the ribbon, you can apply just the colors (as shown here), just the fonts, or just the effects from a theme of your choosing. This technique lets you mix and match different theme parts (which is a bit of a no-no for the artistically minded).*

A more interesting possibility is the way Excel lets you create a brand new custom theme with your own personalized combination of colors and fonts. Here's how:

1. **From the Page Layout→Themes→Themes gallery, choose the theme you want to use as a starting point.**

2. **Choose your favorite body and heading font by going to Page Layout→Themes→Fonts→Create New Theme Fonts.**

   The Create New Theme Fonts dialog box appears (Figure 6-8).

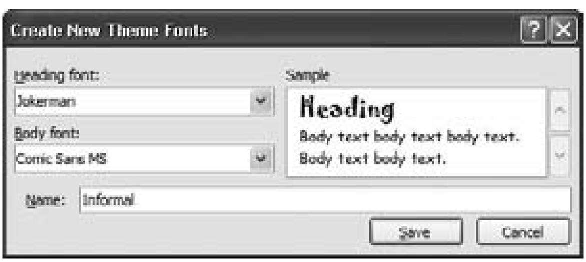

Figure 6-8: *The Create New Theme Fonts dialog box doesn't let you choose font sizes, but you can pick two complementary typefaces: one for body text and one for headings. Before you click Save, enter a descriptive new name.*

3. **Choose your fonts, enter a name for your font combination (like "Informal"), and then click Save.**

   Your font selection appears in the Page Layout→Themes→Font list, so you can use it with any theme.

4. **Choose your favorite colors by going to Page Layout→Themes→Colors→ Create New Theme Colors.**

   The Create New Theme Colors dialog box appears (Figure 6-9).

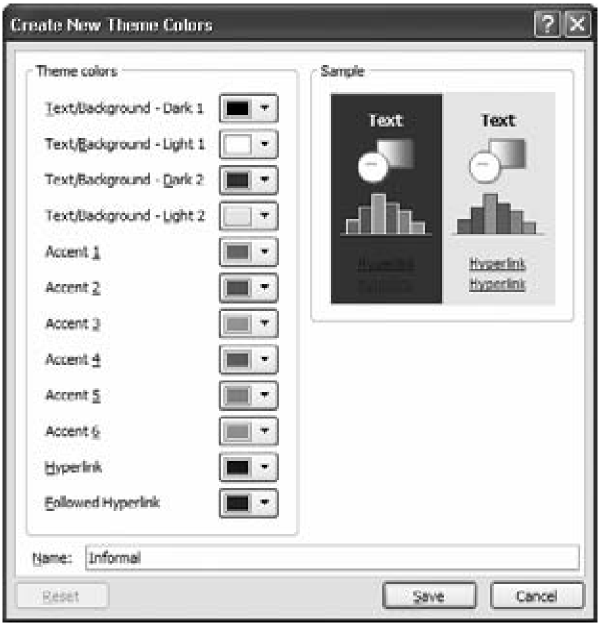

Figure 6-9: *The Create New Theme Colors dialog box shows all 12 theme colors, and lets you adjust each one. Microsoft designed the first two colors for use together as foreground and background colors. The second two colors offer an alternate foreground and background pair. The next six are accented colors for cell backgrounds to highlight important values, and the final two colors are for use with Web-style links (page 90). As you adjust colors, Excel updates the tiny preview pictures.*

5. **Choose your colors, enter a name for your color combination (like "Wacky Office Temp"), and then click Save.**

   Your new color palette appears in the Page Layout→Themes→Colors list, so you can use it with any theme.

6. **Optionally, save your work to a** *.thmx* **file by choosing Page Layout→ Themes→Themes→Save Current Theme.**

   If you want to reuse your theme in other workbooks (or share it with friends), you can save your fonts and colors as a .thmx theme file. To apply your custom theme later on, just choose Page Layout→Themes→Themes→Browse for Themes, browse to your .thmx file, and then choose it.

POWER USERS' CLINIC

Different Office Programs, Same Good Style

Excel shares its theming system with Microsoft Word and PowerPoint. That means you can create a .thmx theme file in Excel, and then use it to set document colors in Microsoft Word, or vice versa. That way, you can create memos in Word, presentations in PowerPoint, and reports in Excel that share the same hot colors.

In fact, you don't even need to create a .thmx file to use this feature (although it's always a good idea). Your theme settings are stored directly in your document file, whether it's an Excel spreadsheet (.xlsx), a Word document (.docx), or a PowerPoint slide deck (.pptx). To pull theme settings directly out of one of these files and place them in your current Excel spreadsheet, choose Page Layout→Themes→Themes→Browse for Themes, and then select the file whose theme you want to use.

Conditional Formatting

A good worksheet highlights the most important information, thereby making it easy to spot. For example, if you look at a worksheet that shows the last year of a company's sales, you want to be able to find underperforming products without having to hunt through hundreds of cells. And even if you're not using Excel in the business world, you still need to home in on key details in a spreadsheet—whether it's a budget-busting dinner in your monthly expense worksheet or a skipped week at the gym in your exercise log. All too often, these essential details are buried in an avalanche of data.

As you learned in Chapter 5, you can use formatting tricks to make important data stand out from the crowd. But the problem with formatting is that it's up to *you* to track down the cells that need to be formatted. Not only is this a time-devouring task, you also run into trouble when you start using formulas (as discussed on page 227). Formulas let you set up elaborate calculations that link cells together, which means that a change to a single cell can cascade through your worksheet, altering data everywhere else. If you're highlighting important information by hand, you just might need to repeat the whole formatting process each time a value changes.

Fortunately, Excel has a feature that's designed to spare you the drudgery. It's called *conditional formatting*, and it lets Excel automatically find and highlight important information. In this section, you'll learn to master conditional formatting to make sure important bits of data stick out for all to see.

Note: In this chapter, you'll learn how to use basic conditional formatting, which applies the formatting settings you learned about in Chapter 5 to make important cells stand out. In Chapter 20, you'll learn about three more advanced types of conditional formatting that build on this system: data bars, color scales, and icon sets. They let you use other graphical tricks (like shaded bars and tiny pictures) to give a *graphical representation* of different values.

The Basics of Conditional Formatting

In Chapter 5, you learned how to create custom format strings. As you saw, Excel lets you create up to three different format strings for the numbers in a single cell (page 187). For example, you can define a format string for positive numbers, a format string for negative numbers, and another format string for zero values. Using this technique, you could create a worksheet that automatically highlights negative numbers in red lettering while leaving nonnegative numbers in black.

This ability to treat negative numbers differently from positive numbers is quite handy, but it's obviously limited. For example, what if you want to flag extravagant expenses that top $100, or you want to flag a monthly sales total if it exceeds the previous month's sales by 50 percent? Custom format strings can't help you there, but conditional formatting fills the gap.

With conditional formatting, you set a condition that, if true, prompts Excel to apply additional formatting to a cell. This new formatting can change the text color, or use some of the other formatting tricks you saw in Chapter 5, including modifying fill colors and fonts. You can also use other graphical tricks, like data bars (shaded bars that grow or shrink based on the number in a cell) and icons.

Highlighting Specific Values

One of the simplest ways to use conditional formatting is to use a little formatting razzle-dazzle to highlight important values. For example, consider the daily calorie intake log shown in Figure 6-10.

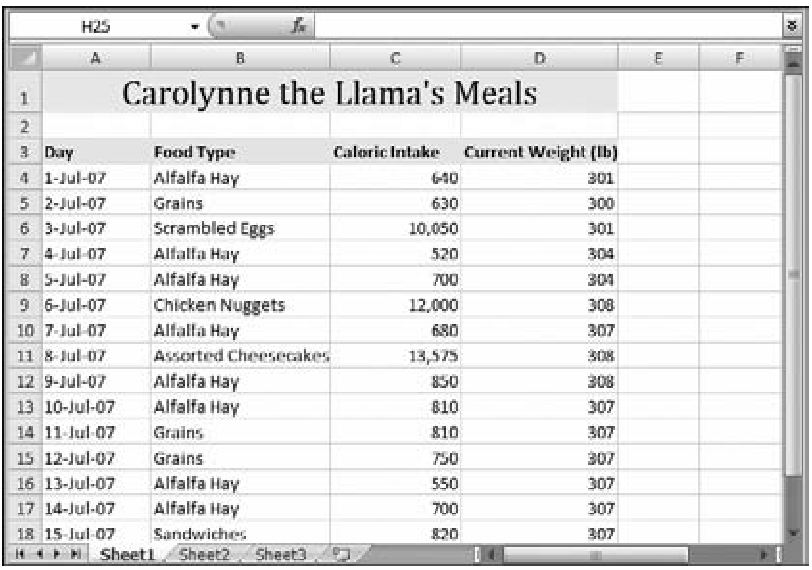

Figure 6-10: *This worksheet tracks the calorie intake of Carolynne, the pet llama, over several weeks. Carolynne's owners have noticed a dramatic weight gain over the same period, but they're at a loss to pinpoint exactly when the overeating took place. Fortunately, conditional formatting can highlight the problem areas.*

To apply conditional formatting, select the cells that you want to examine and format. Next, you need to pick the right conditional formatting rule. A *rule* is an instruction that tells Excel when to apply conditional formatting to a cell and when to ignore it. For example, a typical rule might state, "If the cell value is greater than 10,000, apply bold formatting."

Excel has a wide range of conditional formatting rules, and they fall into two categories (each of which has a separate menu):

- **Highlight specific values**. If your cell contains numbers or dates, you can set a minimum, a maximum, or a range of values that you want Excel to highlight. In the case of text, you can highlight cells that contain certain specific text, start with specific text, and so on. In the case of dates, you can pick out dates that fall within certain ranges (last week, last month, next week, and so on). To see all your choices, choose Home→Styles→Conditional Formatting→Highlight Cells Rules.
- **Highlight values based on where they fall in a series**. These options get Excel to highlight the top values, bottom values, or values that fall above or below average. To see your choices, choose Home→Styles→Conditional Formatting→Top/Bottom Rules.

For example, here's how you can quickly pick out big eating days in the llama food table (Figure 6-10):

1. **In the Caloric Intake column, select the cells.**

   Click the C column's header to select the whole column.

2. **Choose Home→Styles→Conditional Formatting→Highlight Cells Rules→Greater Than.**

   A dialog box appears where you can set the cutoff number and the formatting (Figure 6-11).

Figure 6-11: *Each time Carolynne the llama eats more than 10,000 calories, Excel highlights the cell with a yellow shaded background.*

3. **Set your minimum value in the text box on the left.**

   In this case, use 10000.

Note: Usually, conditional formatting compares a cell value to a fixed number. However, you can also create conditions that compare the cell value to *other* cells in your worksheet. To take this step, select the text box where you'd normally enter the comparison number, and then click the worksheet to select the cell that Excel should use. Excel automatically inserts a cell reference (like $D$2 for cell D2) into the box.

4. **Choose the type of formatting from the list box on the right.**

   You can choose from several presets (like Red Text, Red Border, Red Fill with Dark Red Text, and so on), or you can define your own formatting.

   To define your own format settings, choose Custom Format. An abbreviated version of the Format Cells dialog box appears. A few settings are turned off, because you can't apply them conditionally. For example, you can't conditionally change the font or font size, but you can conditionally set other font characteristics like the use of bold, italic, and underline. Aside from these limitations, the tabs are exactly the same as the ones you're familiar with from the full-blown Format Cells dialog box (page 139). Click OK when you're finished choosing your format options.

Note: Imaginative Excel fans can do a lot with the Format Cells dialog box and conditional formatting. For example, you can highlight specific values by drawing a border around them, adding a different color fill, or changing the number format to add more decimal places.

5. **Click OK.**

   As soon as you click OK, Excel evaluates the conditions and adjusts the formatting as needed. Every time you open your spreadsheet, or change the value in one of the conditional cells, Excel evaluates the condition and adds or removes the formatting as required.

   Figure 6-12 shows the result.

Tip: To remove any type of conditional formatting, select your cells, and then choose Home→Styles→Conditional Formatting→Clear Rules→Clear Rules from Selected Cells. You can also use Home→Styles→Conditional Formatting→Clear Rules→Clear Rules from Entire Sheet to wipe all the conditional formatting from your entire worksheet.

The ribbon is packed with useful conditional formatting choices. However, a few possibilities don't appear in the Highlight Cells Rules and Top/Bottom Rules lists. To see every choice, you can create your conditional formatting rule by hand. Just choose Home→Styles→Conditional Formatting→New Rule. You see the New Formatting Rule dialog box (Figure 6-13).

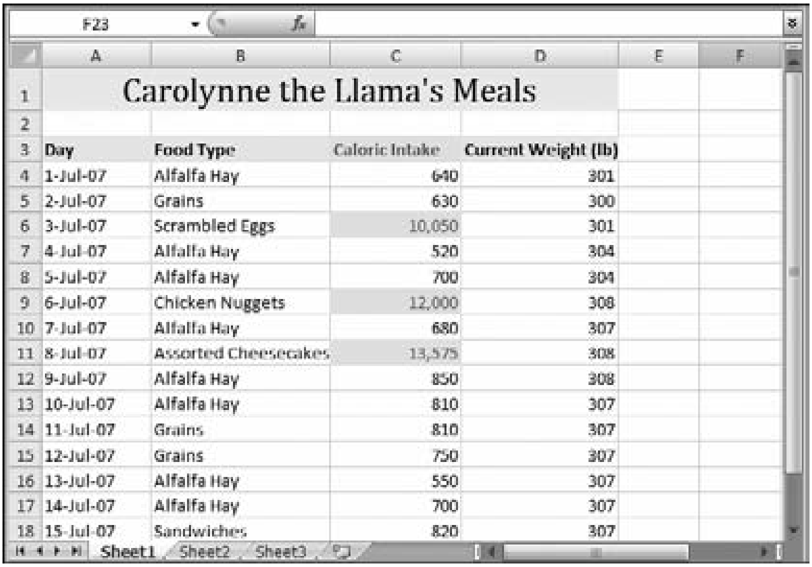

Figure 6-12: *Now Carolynne's days of indulgence stand out. The highlights you see here are a result of the settings applied in "Every time the calorie value tops 10,000, Excel adds a yellow background."*

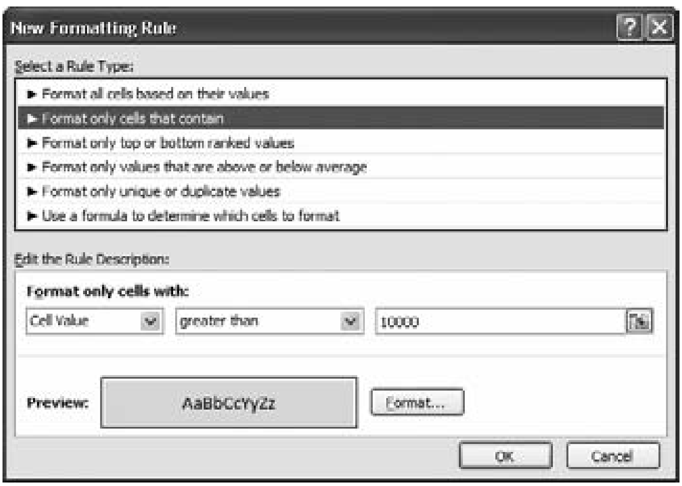

Figure 6-13: *The New Formatting Rule dialog box has two sections. The top portion of the window lets you choose the type of rule. (Ignore the first rule, "Format all cells based on their values," because that's used for data bars, color scales, and icon sets–three features you'll consider a little later in this chapter.) The bottom section lets you define all the rule settings. In this example, a new rule is being created that formats any cell with a value greater than 10,000.*

The New Formatting Rule dialog box is surprisingly intuitive (translation: it's not just for tech jockeys). The "Format only cells that contain" rule is by far the most versatile. It lets you pick out specific numbers, dates, blank cells, cells with errors, and so on. Most people find this formatting rule satisfies most of their conditional formatting needs.

Two rules work well with values that change frequently; these are "Format only top or bottom ranked values" and "Format only values that are above or below average." Both of these rules pick out values that stand out in relationship to the others. For example, if you didn't know that 10,000 calories is the threshold for llama overeating, you might use one of these rules to pick the largest meals, as shown in Figure 6-14.

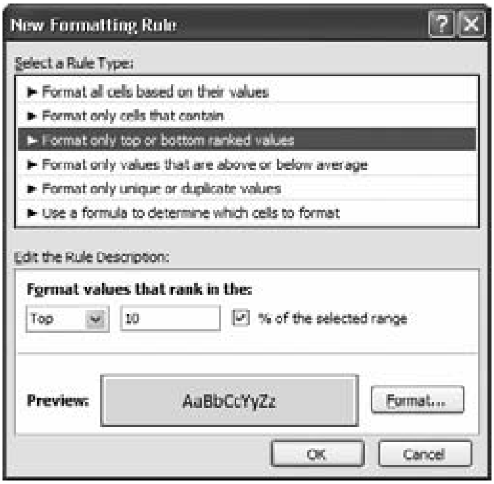

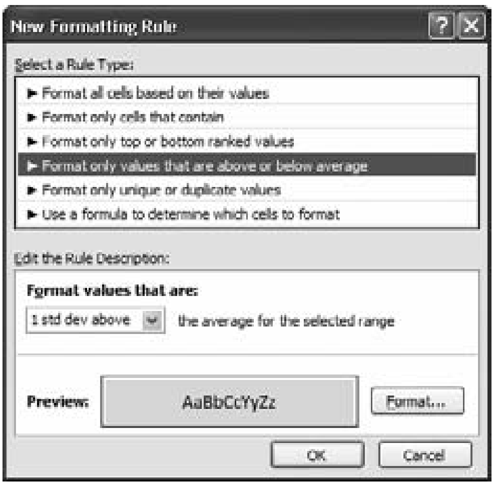

Figure 6-14:
Top: This rule picks out the top 10 percent of all values.

Bottom: This rule picks out values that are one standard deviation above average.

Both rules format the highest values, without your actually needing to know what these values are.

These rules are the foundation of conditional formatting. In the following sections, you'll learn about three more specialized conditional formatting features that use unique formatting to distinguish between different values.

Using Multiple Rules

So far, you've seen examples that use only one conditional formatting rule. However, there's actually no limit to the number of conditional formatting rules you can use at the same time. (Older versions of Excel don't have the same ability—they top out at three conditions.)

Excel gives you two basic ways to use multiple rules:

- You can create rules that format different subsections of data. This lets you apply several different layers of conditional formatting to highlight different values.
- You can create rules that overlap. For example, you can highlight the top five values with red lettering and values above 10,000 with bold. If one of the top five values has a value above 10,000, it gets the combined formatting settings, and Excel displays it in bold red.

If you use conditional rules that overlap, there's always the possibility of conflict. For example, one conditional formatting rule might apply a red background fill while another sets a yellow background fill. If both these rules affect the same cell, only one can win.

In this situation, it all depends on the *order* in which Excel applies conditional formatting rules. If there's a conflict, rules that are applied later override the rules it applied earlier. Ordinarily, Excel applies rules in the same order that you created them, but if this isn't what you want, you can change the order using the Conditional Formatting Rules Manager, which is shown in Figure 6-15. To get to the Conditional Formatting Rules Manager, select one of the cells that uses the conditional formatting, and then choose Home→Styles→Conditional Formatting→Manage Rules.

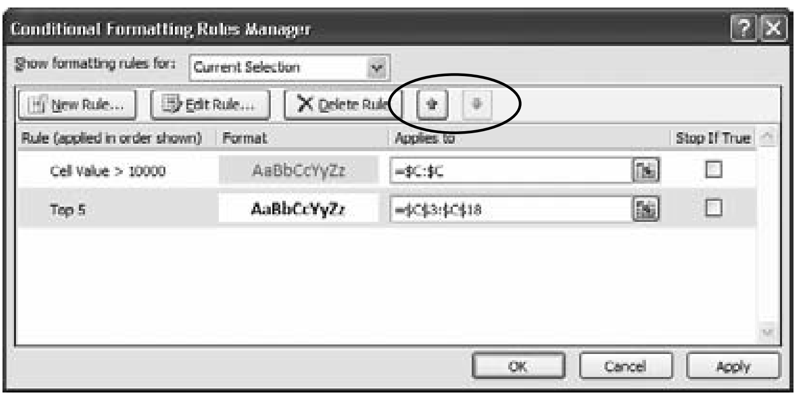

Figure 6-15: *To reorder a rule, select it, and then click the up or down arrow button (circled). Excel applies the rules at the top of the list first. In this example, Excel applies the Cell Value > 10000 rule before the Top 5 rule. As a result, the Top 5 rule formatting may override the Cell Value > 10000 formatting, if they conflict.*

Note: Ordinarily, the Conditional Formatting Rules Manager shows only the conditional formatting rules that apply to the currently selected cell (or cells). However, you can choose a specific worksheet in the "Show formatting rules for" list to see everything that's defined on that worksheet. Now you have a nice way to review all the conditional formatting rules that you've created, but it can be a little confusing. When showing all the formatting on a worksheet, it's important to remember that the order of rules isn't important if the rules apply to different sets of cells.

The Conditional Formatting Rules Manager isn't just for reordering your rules. It also lets you:

- Create rules (click New Rule)
- Modify rules (select a rule in the list, and then click Edit Rule)
- Delete rules (select a rule, and then click Delete Rule)

Finally, there's one easily overlooked gem: the Stop If True column. You can use this setting to tell Excel to stop evaluating conditional formatting rules for a cell. This setting is a handy way to limit where Excel shows data bars, color scales, and icon sets. For example, imagine you create two rules: a Top 5 rule that gives cells a black background and white bold text, and a Cell Value > 10000 rule that gives cells red text. If you put the Top 5 rule first and switch on Stop If True setting, you ensure that the rules will never overlap. If a cell value is both in the Top 5 *and* greater than 10000, it will get only the Top 5 formatting—which is good, because red text is difficult to read on a black background.

CHAPTER

7

Viewing and Printing Worksheets

The previous chapters have given you all the tools you need to create nicely formatted worksheets. While this is all well and good, these features can quickly bury you in an avalanche of data. If you want to see more than one part of your workbook at once, or if you want an overview of the entire worksheet, then you must seize control of Excel's viewing features.

These features include zooming (which lets you magnify cells or just fit more information into your Excel window), panes (which let you see more than one part of a worksheet at once), and freezing (which lets you keep certain cells visible at all times). This chapter teaches you how to use these tools, store a custom view, and even save a *workspace* (a configuration that lets you edit multiple files in one window).

No matter what your worksheets look like on a screen, sometimes the best way to review them is in print. The second half of this chapter tackles printing your worksheets. You'll learn Excel's basic printing options and a few tricks that can help you preview page breaks and make sure large amounts of data get divided the way you want.

Controlling Your View

So far, most of the worksheets in this book have included only a small amount of data. But as you cram your worksheets with dozens of columns, and hundreds or even thousands of rows, editing becomes much trickier. The most challenging problems are keeping track of where you are in an ocean of information and making sure the data you want stays visible. Double that if you have multiple large worksheets in one workbook.

The following sections introduce the basic tools you can use to view your data, along with a few tips for managing large worksheets.

Zooming

Excel's zoom feature lets you control how much data you'll see in the window. When you *reduce* the zoom percentage—say, from 100 percent to 10 percent—Excel shrinks your cells, letting you see more of them at once, which also makes it harder to read the data. Very small zoom percentages are ideal for looking at the overall layout of a worksheet. When you *increase* the zoom percentage—say, from 100 percent to 200 percent—Excel magnifies your cells, letting you see more detail but fewer cells. Larger zoom percentages are good for editing.

Note: Excel lets you zoom in to 500 percent and out all the way to 10 percent.

You can most easily adjust the zoom percent by using the zoom slider in the bottom-right part of the Status bar. The zoom slide also displays the current zoom percentage. But if you want to specify the exact zoom percentage by hand (say, 142 percent), then you can choose View→Zoom→Zoom (or click the zoom percentage next to the zoom slider). A Zoom dialog box appears (shown in Figure 7-1).

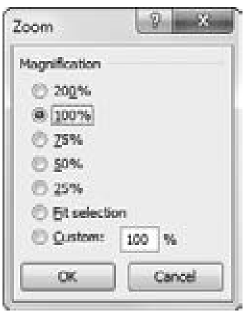

Figure 7-1:
Using the Zoom dialog box, you can select a preset zoom percentage or, in the Custom box, type in your own percentage. However, using the Zoom slider (on the Status bar's right side) is almost always faster than making frequent trips to the Zoom dialog box.

The standard zoom setting is 100 percent, although other factors like the size of the font you're using and the size and resolution of your computer screen help determine how many cells fit into Excel's window. As a rule of thumb, every time you double the zoom, Excel cuts in half the number of rows you can see. Thus, if you can see 20 rows at 100 percent, then you'll see 10 rows at 200 percent.

Note: Changing the zoom affects how your data appears in the Excel window, but it won't have any effect on how your data is printed or calculated.

You can also zoom in on a range of cells. When your data extends beyond the edges of your monitor, this handy option lets you shrink a portion to fit your screen. Conversely, if you've zoomed out to get a bird's eye view of all your data, and you want to swoop in on a particular section, Excel lets you expand a portion to fit your screen. To zoom in on a group of cells, first select some cells (Figure 7-2), and then choose View→Zoom→Zoom to Selection (Figure 7-3). (You can perform this same trick by highlighting some cells, opening the Zoom dialog box, and then choosing "Fit selection.") Make sure you select a large section of the worksheet—if you select a small group, you'll end up with a truly jumbo-sized zoom.

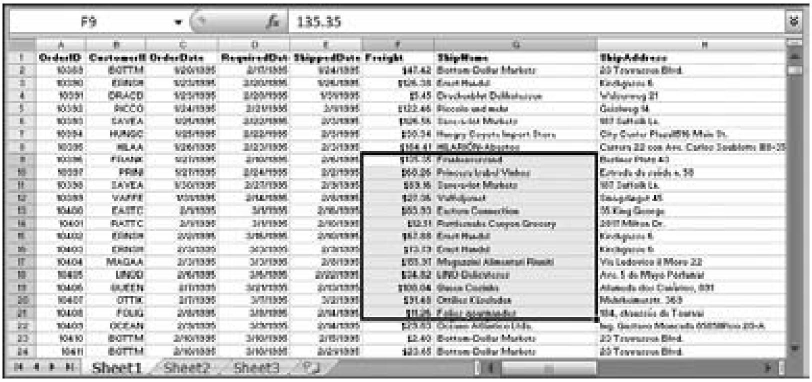

Figure 7-2: *To magnify a range of cells, select them, as shown here, and then choose View→Zoom→Zoom to Selection to have Excel expand the range to fill the entire window, as shown in Figure 7-3.*

Figure 7-3: *Here, Excel increased the zoom to 97 percent (from 57 percent in Figure 7-2).*

Tip: If you're using a mouse with a scroll wheel, then you can zoom with the wheel. Just place the mouse pointer over the worksheet (anywhere will do), hold down the Ctrl key, and roll the scroll wheel up (to zoom in) or down (to zoom out).

GEM IN THE ROUGH

Filling the Screen with Cells

If you really want to see the maximum number of cells at once, Excel provides a little-known feature that strips away the ribbon, the formula bar, and all other extraneous screen elements, making more room for cells. To make the switch, choose View→Workbook Views→Full Screen. To return things to the way they were, right-click anywhere on the worksheet grid, and then choose Close Full Screen (or just press Esc).

Most people find that Full Screen mode is just a little too drastic. Another good option is to collapse the ribbon, which reclaims a significant portion of screen real estate. To do so, in the ribbon, double-click any tab title. Excel hides the ribbon surface, but leaves just the tab titles above your worksheet. Even when the ribbon is collapsed, you can still use it—just click a tab title (which pops that tab back into view), and then click the command you want. The ribbon disappears again as soon as you're done. If you're an unredeemed keyboard lover, then you can use the ribbon in the same way whether it's collapsed or expanded. Just press Alt, and then follow the KeyTips (page 5). And if you get tired of the collapsed ribbon, you can double-click any tab title or press Ctrl+F1 to show the full ribbon once again.

Viewing Distant Parts of a Spreadsheet at Once

Zooming is an excellent way to survey a large expanse of data or focus on just the important cells, but it won't help if you want to simultaneously view cells that aren't near each other. For example, if you want to focus on both row 1 and row 138 at the same time, then zooming won't help. Instead, try splitting your Excel window into multiple *panes*—separate frames that each provide a different view of the same worksheet. You can split a worksheet into two or four panes, depending on how many different parts you want to see at once. When you split a worksheet, each pane contains an identical replica of the entire worksheet. When you make a change to the worksheet in one pane, Excel automatically applies the same change in the other panes. The beauty of panes is that you can look at different parts of the same worksheet at once.

You can split a window horizontally or vertically (or both). When you want to compare different *rows* in the same worksheet, use a horizontal split. To compare different *columns* in the same worksheet, use a vertical split. And if you want to be completely crazy and see four different parts of your worksheet at once, then you can use a horizontal and a vertical split—but that's usually too confusing to be much help.

Excel gives you two ways to split the windows. Here's the easy way:

1. **Find the splitter controls on the right side of the screen.**

   Figure 7-4 shows you where to find them.

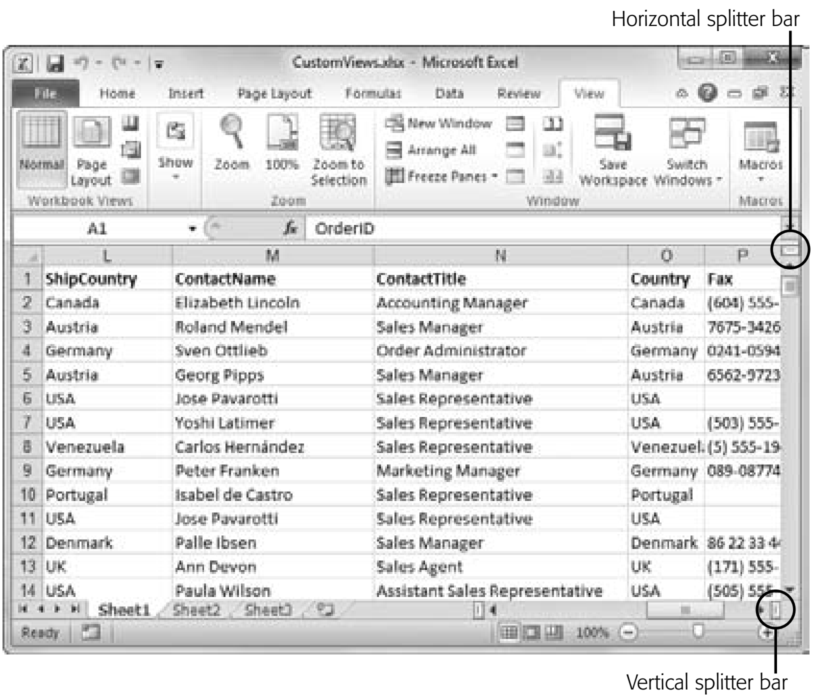

Figure 7-4: *Every Excel window contains both horizontal and vertical splitter controls.*

2. **Drag either control to split the window into two panes. As you drag, Excel displays a gray bar showing where it'll divide the window. Release the splitter control when you're happy with the layout. (At this point, you don't need to worry about whether you can actually view the data you want to compare; you're simply splitting up the window.)**

   If you want to split the window into an upper and lower portion, drag the vertical control down to the location where you want to split the window.

   If you want to split the window into a left and right portion, drag the horizontal control leftwards—to the location where you want to split the window.

Note: If for any reason you *do* want to split the window into four panes, use both controls. The order you follow isn't important.

If you don't like the layout you've created, simply move the splitter bars again by dragging them, just as you did before.

3. **Within each pane, scroll to the cells you want to see.**

   For example, if you have a 100-row table that you split horizontally in order to compare the top five rows and the bottom five rows, scroll to the top of the upper pane, and then scroll to the bottom of the lower pane. (Again, the two panes are replicas of each other; Excel is just showing you different parts of the same worksheet.)

Using the scroll bars in panes can take some getting used to. When the window is split in two panes, Excel synchronizes scrolling between both panes in *one direction*. For example, if you split the window into top and bottom halves, Excel gives you just one *horizontal* scroll bar (at the bottom of the screen), which controls both panes (Figure 7-5). Thus, when you scroll to the left or right, Excel moves both panes horizontally. On the other hand, Excel gives you separate *vertical* scroll bars for each pane, letting you independently move up and down within each pane.

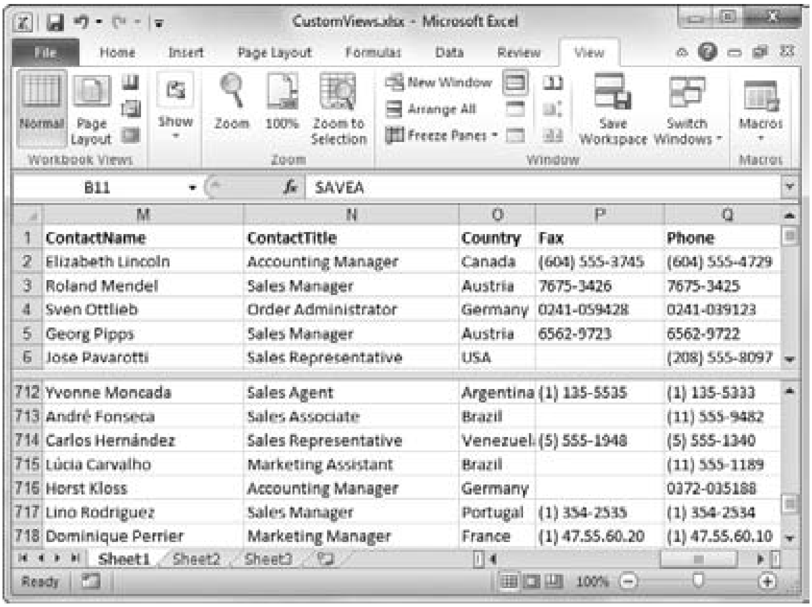

Figure 7-5: *Here you can see the data in rows 1 through 6 and rows 709 through 715 at the same time. As you move from column to column, both panes move in sync, letting you see, for instance, the phone number information in both panes at once. (You can scroll up or down separately in each pane.)*

Tip: If you want the data in one pane—for example, column titles—to remain in place, you can freeze that pane. The next section tells you how.

The reverse is true with a vertical split; in this case, you get one vertical scroll bar and two horizontal bars, and Excel synchronizes both panes when you move up or down. With four panes, life gets a little more complicated. In this case, when you scroll left or right, the frame that's just above or just below the current frame moves, too. When you scroll up or down, the frame that's to the left or to the right moves with you. Try it out.

Note: If you want to remove your panes, just drag the splitter bars back to the edges of the window, or double-click the splitter bar.

You can also create panes by using the ribbon command View→Window→Split. When you do, Excel carves the window into four equal panes. You can change the pane sizes as described above, or use View→Window→Split again to return to normal.

Note: If you use Excel's worksheet navigation tools—like the Go To and Find commands—*all* your panes move to the newly found spot. For example, if you use the Find command in one pane to scroll to a new cell, the other panes display the same cell.

Freezing Columns or Rows

Excel has another neat trick up its sleeve to help you manage large worksheets: *freezing*. Freezing is a simpler way to make sure a specific set of rows or columns remains visible at all times. When you freeze data, it remains fixed in place in the Excel window, even as you move to another location in the worksheet in a different pane. For example, say you want to keep visible the first row that contains column titles. When you freeze that row, you can always tell what's in each column—even when you've scrolled down several screenfuls. Similarly, if your first column holds identifying labels, you may want to freeze it so that when you scroll off to the right, you don't lose track of what you're looking at.

Tip: Excel lets you print out worksheets with a particular row or column fixed in place. Page 217 tells you how.

You can freeze rows at the top of your worksheet, or columns at the left of your worksheet, but Excel does limit your freezing options in a few ways:

- You can freeze rows or columns only in groups. That means you can't freeze column A and C without freezing column B. (You can, of course, freeze just one row or column.)
- Freezing always starts at column A (if you're freezing columns) or row 1 (if you're freezing rows). That means that if you freeze row 13, Excel also freezes all the rows above it (1 through 12) at the top of your worksheet.

- If a row or column isn't visible and you freeze it, you can't see it until you unfreeze it. For example, if you scroll down so that row 100 appears at the top of the worksheet grid, and then freeze the top 100 rows, you can't see rows 1 to 99 anymore. This may be the effect you want, or it may be a major annoyance.

Note: As far as Excel is concerned, frozen rows and columns are a variation on panes (described earlier). When you freeze data, Excel creates a vertical pane for columns or a horizontal pane for rows. It then fixes that pane so you can't scroll through it.

To freeze a row or set of rows at the top of your worksheet, just follow these steps:

1. **Make sure the row or rows you want to freeze are visible and at the top of your worksheet.**

   For example, if you want to freeze rows 2 and 3 in place, make sure they're visible at the top of your worksheet. Remember, rows are frozen starting at row 1. That means that if you scroll down so that row 1 isn't visible, and you freeze row 2 and row 3 at the top of your worksheet, then Excel also freezes row 1—and keeps it hidden so you can't scroll up to see it.

2. **Move to the first row you want *unfrozen*, and then move left to column A.**

   At this point, you're getting into position so that Excel knows where to create the freeze.

3. **Select View→Window→Freeze Panes→Freeze Panes.**

   Excel splits the worksheet, but instead of displaying a gray bar (as it does when you create panes), it uses a solid black line to divide the frozen rows from the rest of the worksheet. As you scroll down the worksheet, the frozen rows remain in place.

   To unfreeze the rows, just select View→Freeze Panes→Unfreeze Panes.

Freezing columns works the same way:

1. **Make sure the column or columns you want to freeze are visible and at the left of your worksheet.**

   For example, if you want to freeze columns B and C in place, make sure they're visible at the edge of your worksheet. Remember, columns are frozen starting at column A. That means that if you scroll over so that column A isn't visible, and you freeze columns B and C on the left side of your worksheet, Excel also freezes column A—and keeps it hidden so you can't scroll over to see it.

2. **Move to the first column you want *unfrozen*, and then move up to row 1.**

   At this point, you're getting into position so that Excel knows where to create the freeze.

3. **Select View→Window→Freeze Panes→Freeze Panes.**

   Excel splits the worksheet, but instead of displaying a gray bar (as it does when you create panes), Excel uses a solid black line to divide the frozen columns from the rest of the worksheet. As you scroll across the worksheet, the frozen columns remain in place.

   To unfreeze the columns, select View→Window→Freeze Panes→Unfreeze Panes.

Tip: If you're freezing just the first row or the leftmost column, then there's no need to go through this whole process. Instead, you can use the handy View→Window→Freeze Panes→Freeze Top Row or View→Window→Freeze Panes→Freeze First Column.

You can also freeze columns and rows *at the same time*, which is useful when you have identifying information that you need to keep visible both on the left and the top of your worksheet. Figure 7-6 shows an example.

Figure 7-6: *Here, both column A and row 1 are frozen, and thus always remain visible. The easiest way to create these frozen regions is to scroll to the top of the worksheet, position the active cell B2, and choose View→Window→Freeze Panes→Freeze Panes. Excel then automatically freezes rows above and the columns to the left in separate panes.*

Tip: You can also create a horizontal or vertical pane by moving one of the splitter bars, and then freezing that pane. Just drag the splitter bar to the appropriate position, and select View→Window→Freeze Panes→Freeze Panes.

Hiding Data

In some cases, your problem isn't that you need to keep data visible, but that you need to *hide* it. For example, say you have a column of numbers that you need only for a calculation but don't want to see when you edit or print the sheet. Excel provides the perfect solution: *hiding* rows and columns. Hiding doesn't delete information; it just temporarily tucks it out of view. You can restore hidden information any time you need it.

Technically, hiding a row or column is just a special type of resizing. When you instruct Excel to hide a column, it simply shrinks the column down to a width of 0. Similarly, when you hide a row, Excel compresses the row height.

Note: You can also hide an entire worksheet of data. See Chapter 4 for details.

You can hide data in a few ways:

- To hide a column, right-click the column header (the letter button on the top of the column), and then choose Hide. Or, put your cursor in any row in that column, and then select Home→Cells→Format→Hide & Unhide→Hide Columns.
- To hide a row, right-click the row header (the number button at the left of the row), and then choose Hide. Or, put your cursor in any column in that row, and then select Home→Cells→Format→Hide & Unhide→Hide Rows.
- To hide multiple rows or columns, just select all the ones you want to disappear before choosing Hide.

To unhide a column or row, select the *range* that includes the hidden cells. For example, if you hid column B, select columns A and C by dragging over the numeric row headers. Then choose Home→Cells→Format→Hide & Unhide→Unhide Columns (or Unhide Rows). Or just right-click the selection, and then choose Unhide. Either way, Excel makes the missing columns or rows visible, and then highlights them so you can see which information you've restored.

Tip: To unhide all columns (or rows) in a worksheet, select the entire worksheet (by clicking the square in the top-left corner of the grid), and then select Home→Cells→Format→Hide & Unhide→Unhide Columns (or Unhide Rows).

Forgetting that you've hidden data is as easy as forgetting where you put your keys. While Excel doesn't include a hand-clapper to help you locate your cells, it does offer a few clues that some of your row numbers or column letters are missing, as shown in Figure 7-7.

Figure 7-7: *This worksheet jumps directly from column A to column O, which tells you that B through N are hidden.*

Tip: Excel doesn't let you hide individual cells. However, Excel gurus use workarounds. The first one is to format the cell so that the text's white (because white lettering on a white background is invisible). Another solution is to format the cell with the custom number format ***;;;*** (which doesn't show anything for positive, negative, or text values; see page 150 for more on custom formatting). When you use either of these tricks, you can still see the cell content by moving to the cell and looking in the formula bar.

Saving View Settings

If you regularly tweak things like the zoom, visible columns, and the number of panes, you can easily spend more time adjusting your worksheet than editing it. Fortunately, Excel lets you save your view settings with *custom views*. Custom views let you save a combination of view settings in a workbook. You can store as many custom views as you want. When you want to use a particular view you've created, simply select it from a list and Excel applies your settings.

Custom views are particularly useful when you frequently switch views for different tasks, like editing and printing. For example, if you like to *edit* with several panes open and all your data visible, but you like to *print* your data in one pane with some columns hidden, custom views let you quickly switch between the two layouts.

Note: You can't save a custom view for one worksheet, and then apply it to another.

Custom views can save the following settings:

- The location of the active cell. (In other words, your position in the worksheet. For example, if you've scrolled to the bottom of a 65,000-row spreadsheet, then the custom view returns you to the active cell in a hurry.)
- The currently selected cell (or cells).
- Column widths and row heights, including hidden columns and rows.
- Split panes (page 198).

- View settings, like the zoom percentage, which you set using the ribbon's View tab.
- Print settings (page 213), like the page margins.
- Filter settings, which affect what information Excel shows in a data list (see Chapter 14).

Note: Excel does have one restriction on custom views. You can't use views with Excel's tables feature, which is described in Chapter 14. In fact, once you add a table to your worksheet, Excel turns the Custom Views button off.

To create a custom view, follow these steps:

1. **Adjust an open worksheet for your viewing pleasure.**

   Set the zoom, hide or freeze columns and rows, and move to the place in the worksheet where you want to edit.

2. **Choose View→Workbook Views→Custom View.**

   The Custom Views dialog box appears, showing you a list of all the views defined for this workbook. If you haven't created any yet, this list is empty.

3. **Click the Add button.**

   The Add View dialog box appears.

4. **Type in a name for your custom view.**

   You can use any name, but consider something that'll remind you of your view settings (like "50 percent Zoom"), or the task that this view is designed for (like "All Data at a Glance"). A poor choice is one that won't mean anything to you later ("View One" or "Zoom with a View") or something obscure like "'57 Chevy."

   The Add View dialog box also gives you the chance to specify print settings or hidden rows and columns that Excel *shouldn't* save as part of the view. Turn off the appropriate checkboxes if you don't want to retain this information. Say you hide column A, but you clear the "Hidden rows, columns, and filter settings" checkbox because you don't want to save this as part of the view. The next time you restore the view, Excel won't make any changes to the visibility of column A. If it's hidden, it stays hidden; if it's visible, it stays visible. On the other hand, if you want column A to always be hidden when you apply your new custom view, then keep the "Hidden rows, columns, and filter settings" checkbox turned on when you save it.

   After you've typed your view name and dealt with the inclusion settings, click OK to create your new view. Excel adds your view to the list.

5. **Click Close.**

   You're now ready to use your shiny new view or add another (readjust your settings and follow this procedure again).

Applying your views is a snap. Simply select View→Workbook Views→Custom Views to return to the Custom Views dialog box (Figure 7-8), select your view from the list, and then click Show. Because Excel stores views with the workbook, they'll always be available when you open the file, even if you take that file to another computer.

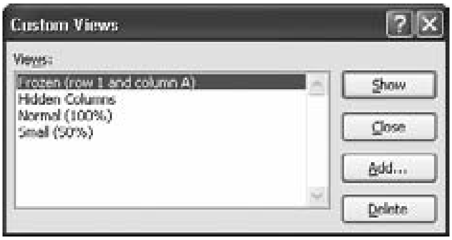

Figure 7-8:
You can use the Custom Views dialog box to show or delete existing views or to create new ones (click Add, and then follow the procedure from step 4).

Tip: For some examples of custom views in action, visit this book's Missing CD page at *www.missing-manuals.com/cds* and download CustomViews.xls, a sample spreadsheet with an array of custom views already set up.

Viewing Multiple Workbooks at Once

In its usual state, Excel lets you view only one open workbook file at a time. If you want to compare two or more workbooks, you must switch between them using the Windows taskbar. But this action is a pain, especially if you want to compare two worksheets side-by-side.

But Excel does provide a handy tool that lets you place several open workbooks inside one large Excel window, and save your arrangement. This setup is called a *custom workspace*. With custom workspaces, you can arrange all the workbooks you need for a particular task the way you like them, and then save that arrangement of windows in a special workspace file. Then when it's time to get back to work on your project, you simply open the workspace file, and Excel restores all the windows exactly the way you left them.

Note: You can use a custom workspace to work on different parts of a single workbook at once. However, custom workspaces really come in handy if you need to work on multiples files simultaneously.

To create and save a workspace, follow these steps:

1. **Open all the spreadsheet files you want to make part of your workspace. Close all other Excel files.**

   Should you want different worksheets from the *same* workbook to be part of your workspace, you must open duplicate versions of the workbook. To do this, go to the workbook and select View→Window→New Window. Excel opens a second (or third, or fourth…) window that shows the same workbook. Don't worry though—any change you make in one window automatically appears in the others, because there's still just one open workbook. The only way you can tell that you have more than one window open for the same workbook is to look at the title bar of the window, which adds a colon and a number. For example, when you open a second view on MyBeanieBabies.xlsx, you'll see the window title MyBeanieBabies.xlsx:2.

   At this point, all your windows appear in the taskbar at the bottom of the screen. You can also check which spreadsheets are open (and jump from one to another) using the View→Window→Switch Windows list, which shows all the currently open windows.

2. **Choose View→Window→Arrange All.**

   The Arrange Windows dialog box appears.

3. **Choose an Arrange option, and then click OK.**

   Excel gives you four choices for window arranging:

   - **Horizontal**, as shown in Figure 7-9, stacks the windows from top to bottom. Excel arranges the windows one above the other, each occupying the full width of the Excel window (similar to when you split a worksheet with the horizontal splitter bar).
   - **Vertical** instructs Excel to tile the windows from left to right.
   - **Tiled** arranges the windows in a grid pattern whose composition changes depending on the number of files you're arranging. This option is great if you have a huge monitor.
   - **Cascade** layers the windows on top of each other with just a smidge of each window showing.

If you open multiple windows on the same workbook, you can select the "Windows of active workbook" option to tell Excel to ignore any other open workbooks.

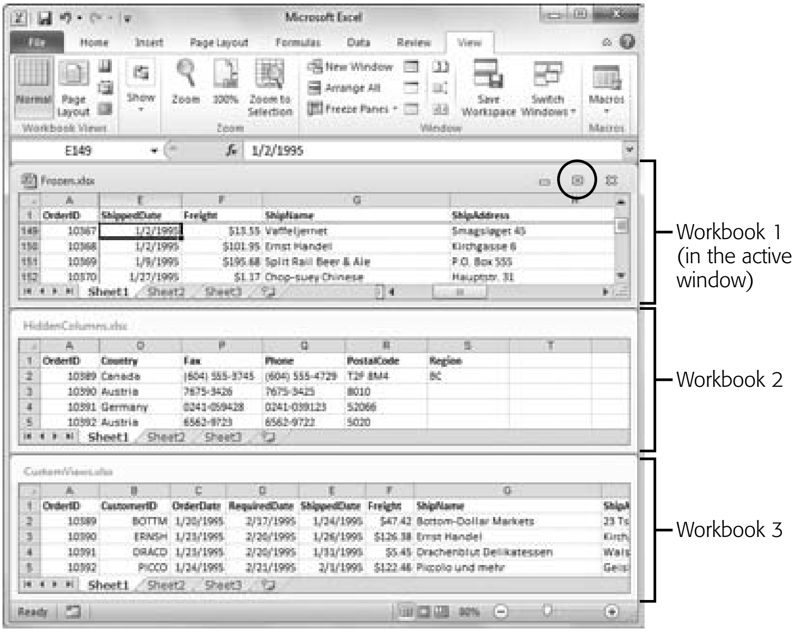

Figure 7-9: *These three spreadsheets have been arranged horizontally. HiddenColumns.xlsx is the active window. To return to the standard, one-file view, just double-click the blue title bar on any window, or click its Maximize button at the right (circled).*

If you'd like to save a particular set of arranged windows, Excel lets you do so by creating a custom workspace. However, unlike a custom view, you can't save custom workspaces inside an individual spreadsheet file. Therefore, you must save custom workspaces as separate *workspace files*, which specify those files you've included and the position of their windows. (These files, which use the .xlw file extension, don't actually contain any of the individual spreadsheet file data.) When you open the custom workspace file, Excel automatically loads all the files you were using and returns them to their original locations.

Once you've completed the above sequence of steps to create your workspace, you can save it at any time. Just select View→Window→Save Workspace, and then choose a file name. You can open a workspace file in the same way you open a spreadsheet: by choosing File→Open or by double-clicking the file on your desktop or in Windows Explorer.

Note: If you ever decide to delete a workspace file, bear in mind that you're deleting only the information about the customized window arrangement. The individual Excel files themselves are stored separately.

Workspaces have two minor quirks you should be aware of:

- The workspace file stores the *location* of the Excel files it uses. If you move one of these files somewhere else, the workspace file can't find it and load it.
- If you open a workspace file and then change the window arrangement or open new worksheets, Excel won't prompt you to save the new workspace settings. Instead, you need to explicitly choose View→Window→Save Workspace again.

Tip: You can use custom workspaces as a shortcut to open multiple files you want to work on at the same time—even if you don't want to use Excel's window-arranging features. To do so, just open all the files into separate windows, and then save the workspace.

Printing

Printing in Excel is pretty straightforward—as long as your spreadsheet fits on a normal 8.5×11-inch piece of paper. If you're one of the millions of spreadsheet owners who don't belong to that club, welcome to the world of Multiple Page Disorder: the phenomenon in which pages and pages of apparently unrelated and noncontiguous columns start spewing from your printer. Fortunately, Excel comes with a slew of print-tweaking tools designed to help you control what you're printing. First off, though, it helps to understand the standard settings that Excel uses.

Note: You can change most of the settings listed; this is just a list of what happens if you *don't* adjust any settings before printing a spreadsheet.

- In the printout, Excel uses all the formatting characteristics you've applied to the cells, including fonts, fills, and borders. However, Excel's gridlines, row headers, and column headers *don't* appear in the printout.
- If your data is too long (all the rows won't fit on one page) or too wide (all the columns won't fit), Excel prints the data on multiple pages. If your data is both too long *and* too wide, Excel prints in the following order: all the rows for the first set of columns that fit on a printed page, then all the rows for the next set of columns that fit, and so on (this is known as "down, then over"). When printing on multiple pages, Excel never prints *part* of an individual column or row.
- Excel prints your file in color if you use colors and you have a color printer.
- Excel sets margins to 0.75 inches at the top and bottom of the page, and 0.7 inches on the left and right sides of the page. Ordinarily, Excel doesn't include headers and footers (so you don't see any page numbers).
- Excel doesn't include hidden rows and columns in the printout.

How to Print an Excel File

Excel 2010 uses its backstage view to make printing a whole lot less confusing. Its key feature is a built-in preview that shows you what the on-paper printout will look like before you click Print.

If you're in a tremendous hurry to get your printout and you're not interested in playing with print settings, just choose File→Print, and then click the big Print button shown in Figure 7-10.

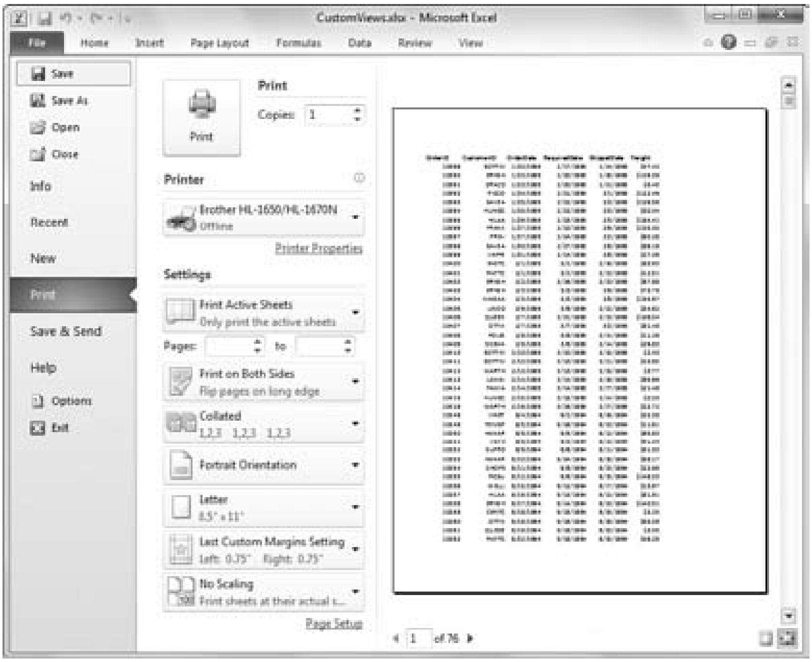

Figure 7-10: *When it comes to printing, Excel's backstage view is a small miracle of efficiency. Not only does it let you tweak the most common print settings, but it also shows you the effect in an instantly updated preview.*

If this no-fuss printing approach doesn't give you the results you want, you need to take a closer look at the print settings you can tweak. Here's a walkthrough that takes you through all your options:

1. **Choose File→Print (or press Ctrl+P).**

   Excel switches into backstage view, and shows a two-part display that puts printing options on the left and a print preview on the right (Figure 7-10).

2. **If you want to print multiple copies of your data, use the Copies box.**

   If you want to print more than one identical copy of your data, change the Copies box accordingly. (This is automatically set at 1 to print a single copy of your work.)

If you're printing more than one copy and your worksheet has multiple pages, you should also review the collating setting, which appears further down. This setting determines whether Excel duplicates each page separately. For example, if you print 10 pages and your printout is set to Uncollated, Excel prints 10 copies of page 1, 10 copies of page 2, and so on. If your printout it set to Collated, Excel prints the entire 10-page document, then prints out another copy, and so on. You'll still end up with 10 copies of each page, but they'll be grouped together for added convenience.

3. **Select a printer from the drop-down list under the Printer heading.**

   Excel automatically selects your regular printer. If you have more than one printer installed and you want to use a different printer, you need to select this printer yourself. You can also adjust printer settings by clicking the Printer Properties link. Every printer has its own set of options here, but common Properties settings include print quality and paper handling (like double-sided printing for those lucky enough to have a printer that supports it).

4. **Choose what you want to print from the first list under the Settings heading.**

   Print Active Sheets prints the current worksheet. If you've grouped more than one worksheet together using the techniques described on page 125, Excel prints all the selected worksheets, one after the other.

   Print Entire Workbook prints all the worksheets in your file.

   Print Selection prints out just a portion of a worksheet. To make this feature work, you need to start by selecting a range of cells, columns, or rows before you start your print out, and *then* choose File→Print.

5. **If you want to print just some pages, use the two Pages boxes.**

   By default, Excel prints as many pages as it needs to output all the data you've chosen to print. Alternately, you can choose a range of pages using the Pages option. For example, you can choose to print only the first three pages by typing 1 into the first box and 3 instead the second. You can also print just the fourth page by printing from 4 to 4.

Note: In order to use the "Print range" box effectively, you need to know how many pages you need to print your worksheet and what data will appear on each page. You can step through all the pages in your printout using the handy print preview shown in Figure 7-10.

6. **Set the orientation and paper size.**

   Orientation is one of the all-time most useful print settings. It lets you control whether you're printing on pages that are upright (choose Portrait Orientation) or turned horizontally on their sides (choose Landscape Orientation). If Excel is splitting your rows across multiple pages when you print your worksheet, it makes good sense to switch to landscape orientation. That way, Excel prints your columns across a page's long edge, which accommodates more columns (but fewer rows per page).

If you're fed up with trying to fit all your data on an ordinary sheet no matter which way you turn it, you may be tempted to try using a longer sheet of paper. You can then tell Excel what paper you've decided to use by choosing it from the list just under the orientation setting. (Of course, the paper needs to fit into your printer.) Letter is the standard 8.5×11-inch sheet size, while Legal is another common choice—it's just as wide but comes in a bit longer at 8.5×14 inches.

7. **Adjust your margins.**

   Beneath the options for page orientation and paper size is the margin setting, which determines the amount of space between your worksheet content and the edges of the page.

   You can set the margin setting in two ways. The easiest approach is to pick one of the presets (Normal, Wide, or Narrow), as shown in Figure 7-11.

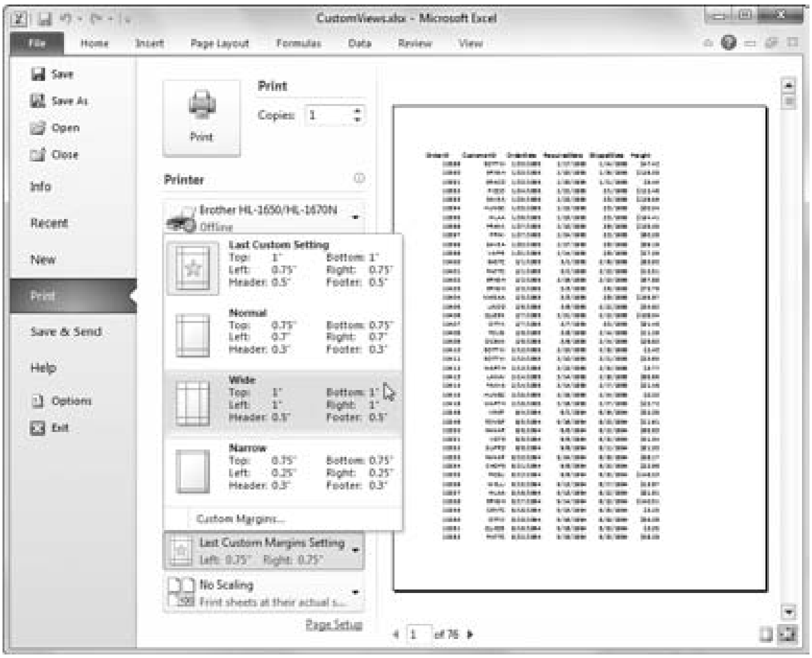

***Figure 7-11:** Each margin preset includes several numbers. The Top, Bottom, Left, and Right values measure the amount of space between your worksheet content and the top, bottom, left, and right edges of the page, respectively.*

For more control, you can choose Custom Margins and fill in your own margin values (Figure 7-12). Logically enough, when you reduce the size of your margins, you can accommodate more information. However, you can't *completely* eliminate your margins. Most printers require at least a little space (usually no less than 0.25 inches) to grip onto the page, and you won't be able to print on this part (the very edge of the page). If you try to make the margins too small, Excel won't inform you of the problem; instead, it'll just stick with the smallest margin your current printer allows.

If you have only a few rows or columns of information, you may want to use one of the "Center on page" options. Select Horizontally to center your columns between the left and right margins. Select Vertically to center your data between the top and bottom of the page.

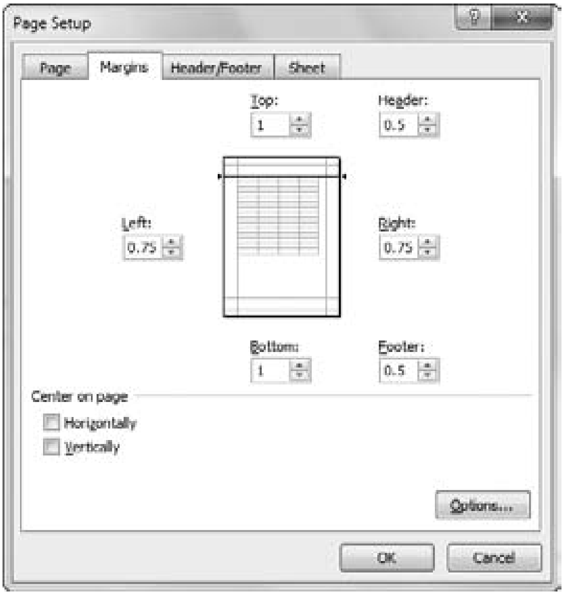

Figure 7-12:
Excel allocates space at the top and bottom of your printout for a header or footer (page 220). In this example, the header margin is set to 0.5, which means that any header information will appear half an inch below the top of the page. The top margin is set to 1, meaning the worksheet data will appear one inch below the top of the page. When adjusting either of these settings, be careful to make sure the top margin is always larger than the header margin; otherwise, your worksheet's data will print on top of your header. The same holds true with footers when changing the bottom margin. (But if you aren't using headers or footers, their margin settings don't matter.)

Tip: A good rule of thumb is to adjust margins symmetrically (printouts tend to look nicest that way). Thus, if you shrink the left margin to 0.5, make the same change to the right margin. Generally, if you want to fit more data and you don't need any header or footer space, then you can safely reduce all your margins to 0.5. If you really want to cram in the maximum amount of data you can try 0.25, but that's the minimum margin that most printers allow.

8. **If you need to shrink your printout and cram more information into each page, pick a scaling option.**

   No matter how drastically you reduce your margins, you'll only be able to fit a few extra rows and columns in a page. A more powerful approach for fitting mass amounts of data into a smaller number of pages is to use *scaling*. Page 223 gives more detail, but for now, you can try one of the handy scaling presets:

- **Fit All Columns on One Page** squashes your page width-wise, making it narrower. This way, the columns won't leak off the edge into a separate page.
- **Fit All Rows on One Page** squashes your page height-wise, making it shorter. This way, all your rows will appear on the same page.
- **Fit Sheet on One Page** squashes your page both ways, making sure all your data fits on a single sheet.

Note: Excel performs scaling by reducing the font size in the printout. If you try to cram too much data into too small a space, your text might shrink into near-oblivion. It can be hard to judge just how small your text is from the print preview, so you might need to print your worksheet to see how much scaling is too much.

9. **If you still want more options, click the Page Setup link.**

   This shows the Page Setup dialog box (Figure 7-13), which holds a few of Excel's more specialized print settings. The Page Setup dialog box is organized into several tabs. The Page and Margins tabs duplicate settings that are provided elsewhere in backstage view. The Header/Footer tab isn't the most convenient way to add a header or footer (instead, see page 220). However, the Sheet tab has a number of options you won't find anywhere else:

   - **Print area** lets you specify the range of cells you want to print. While this tool definitely gets the job done, it's easier to use the Print Area tool (described in the box on page 217). And some people find that the Print dialog box's Selection setting (step 4) is an easier approach to printing small groups of cells.
   - **Print titles** lets you print specific rows at the top of every page, or specific columns on the left side of every page. For example, you could use this setting to print column titles on the top of every page.

Note: Due to a strange Excel quirk, you can't modify the "Print area" or "Print titles" settings while Excel is showing a preview. Instead, you need to close backstage view (just press Esc), head to the Page Layout→Sheet Options section of the ribbon, and click the dialog box launcher. This shows the same Page Setup dialog box, but with all its options enabled.

   - **Gridlines** prints the grid of lines separating columns and rows that you see on your worksheet.

- **Row and column headings** prints the column headers (which contain the column letters) at the top of each page and the row headers (with the row numbers) on the left side of each page.
- **Black and white** tells Excel to render all colors as a shade of gray, regardless of your printer settings.
- **Draft quality** tells Excel to use lower-quality printer settings to save toner and speed up printing—assuming your printer has these features, of course.
- **Comments** lets you print the comments that you've added to a worksheet. Excel can either append them to the cells in the printout or add them at the end of the printout, depending on the option you select. For the lowdown on Comments, see page 724.
- **Cell errors** lets you configure how Excel should print a cell if it contains a formula with an error. You can choose to print the error that's shown (the standard option), or replace the error with a blank value, two dashes (--), or the error code #N/A (meaning not available). You'll learn much more about formulas in Chapter 8.
- **Page order** sets the way Excel handles a large worksheet that's too wide and too long for the printed page's boundaries. When you choose "Down, then over" (the standard option), Excel starts by printing all the rows in the first batch of columns. Once it's finished this batch, Excel then moves on to the next set of columns, and prints those columns for all the rows in your worksheet, and so on. When you chose "Over, then down," Excel moves across your worksheet first. That means it prints all the columns in the first set of rows. After it's printed these pages, it moves to the next set of rows, and so on.

10. **Now that all your configuration is complete, click the Print button to send the spreadsheet to the printer.**

    Excel prints your document using the settings you've selected.

    If you're printing a very large worksheet, Excel shows a Printing dialog box for a few seconds as it sends the pages to the printer. If you decide to cancel the printing process—and you're quick enough—you can click the Cancel button in this Printing dialog box to stop the operation. If you don't possess the cat-like reflexes you once did, you can also open your printer queue to cancel the process. Look for your printer icon in the notification area at the bottom-right of your screen, and double-click that icon to open a print window. Select the offending print job in the list, and then press Delete (or choose Document→Cancel from the print window's menu). Some printers also provide their own cancel button that lets you stop a print job even after it's left your computer.

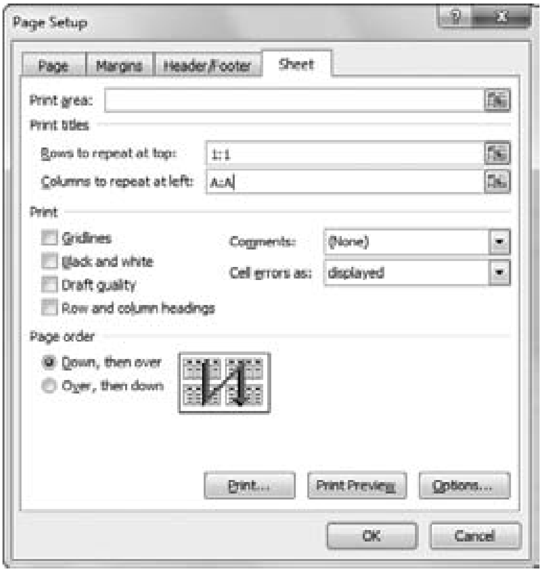

Figure 7-13:
In this example, Excel uses the "Print titles" section to ensure that every page in this printout will display the first row of the spreadsheet as well as the first column.

GEM IN THE ROUGH

Printing Parts of a Spreadsheet

When working with large worksheets, you'll often want to print only a small portion of your total data. Excel gives you several ways to limit your printout. You can hide the rows or columns you aren't interested in, or you can select the cells you want to print, and, in the Print dialog box's "Print what" box, choose Selection. But if you frequently need to print the same area, you're better off defining and using a *print area.*

A print area designates a portion of your worksheet as the only region that Excel will print. (The one exception is if you choose Selection from the "Print what" box, in which case Excel prints the selected cells, not the print area.) Once you define a print area, Excel retains it until you remove it. That means you can make changes, save, close, and open your spreadsheet, and the same print area remains in place.

To set a print area, select the rows, columns, or group of cells, and then choose Page Layout→Page Setup→Print Area→Set Print Area. The portion of the worksheet that you've highlighted now has a thin dashed outline, indicating that this is the only region Excel will print. You can only have one print area at a time, and setting a new one always clears the previous one. To remove your print area so that you can print the entire worksheet, choose Page Layout→Page Setup→Print Area→Clear Print Area.

Page Layout View: A Better Print Preview

When you're preparing to print that 142-page company budget monstrosity, there's no reason to go in blind. Instead, prudent Excel fans use Page Layout view to check out what their printouts look like *before* they appear on paper.

Page Layout view is a bit like the print preview that you saw backstage (Figure 7-11), but it's more powerful. First, it's bigger and easier to navigate. More importantly, it lets you do a few things that aren't possible in backstage view, like setting headers and footers, editing cell values, and tweaking other page layout settings from the ribbon.

To see the Page Layout view for a worksheet, choose View→Workbook Views→Page Layout View. Or, for an even quicker alternative, use the tiny Page Layout View button in the Status bar, which appears immediately to the left of the zoom slider. Either way, you see a nicely formatted preview (Figure 7-14).

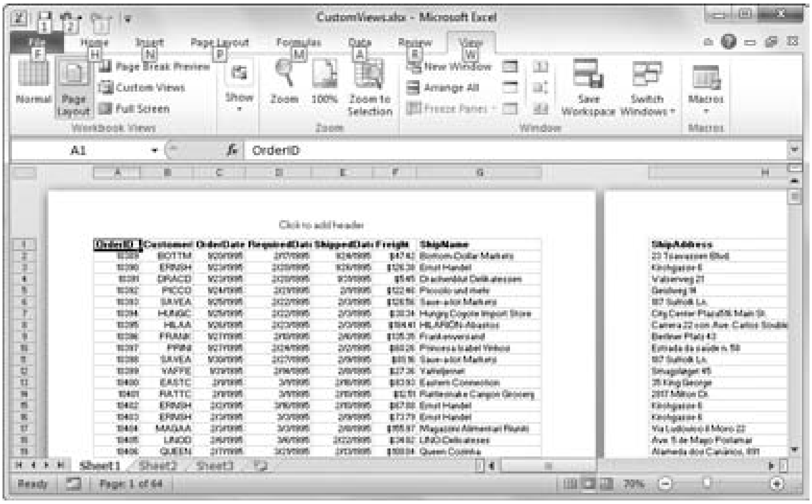

Figure 7-14: *The Page Layout view shows the first (and part of the second) page of this worksheet's 76 printed pages. This worksheet has 19 columns, but since they're wider than the width of a single printed page, the first page includes only the leftmost seven columns, as shown here. You can scroll to the right to see the additional columns that'll turn up on other pages, or scroll down to see more rows.*

How does Page Layout view differ from Normal view? For starters, Page Layout view:

- Paginates your data. You see exactly what fits on each page, and how many pages your printout requires.
- Reveals any headers and footers you've set as part of the page setup. These details don't appear in the Normal worksheet view.
- Shows the margins that Excel will use for your pages.
- Doesn't show anything that Excel won't print (like the letters at the top of each column). The only exception is the cell gridlines, which are shown to help you move around your worksheet.
- Includes a bit of text in the Status bar that tells you where you are, page-wise, in a large spreadsheet. For example, you might see the text "Page: 5 of 26."

Note: Don't confuse Page Layout view with an ordinary print preview (like the one you see when you choose File→Print). A print preview provides a fixed "snapshot" of your printout. You can look, but you can't touch. Page Layout view is vastly better because it shows what your printout will look like *and* it lets you edit data, change margins, set headers and footers, create charts, draw pictures–you get the idea. In fact, you can do everything you do in Normal view mode in Page Layout view. The only difference is you can't squeeze quite as much data into the view at once.

If you aren't particularly concerned with your margin settings, you can hide your margins in Page Layout view so you can fit more information into the Excel window. Figure 7-15 shows you how.

Figure 7-15: *Move your mouse between the pages and your mouse pointer changes into this strange two-arrow beast. You can then click to hide the margins in between pages (as shown here), and click again to show them (as shown in Figure 7-14). Either way, you see an exact replica of your printout. The only difference is whether you see the empty margin space.*

Here are some of the tasks you may want to perform in Page Layout view:

- To tweak print settings and see the effect, choose the Page Layout tab in the ribbon. The most important print-related sections are Page Setup (which lets you change orientation and margin settings), Scale to Fit (which lets you cram more information into your printed pages), and Sheet Options (which lets you control whether gridlines and column headers appear on the printout).
- To move from page to page, you can use the scroll bar at the side of the window, or you can use the keyboard (like Page Up, Page Down, and the arrow keys). When you reach the edge of your data, you see shaded pages with the text "Click to add data" superimposed. If you want to add information further down the worksheet, just click one of these pages, and then start typing.
- To adjust the page margins, first make sure the ruler is visible by turning on the View→Show→Ruler checkbox. Then, drag one of the margin lines on the ruler, as shown in Figure 7-16. If you want to set page margins by typing in the exact margin width, use the Page Layout tab of the ribbon instead.

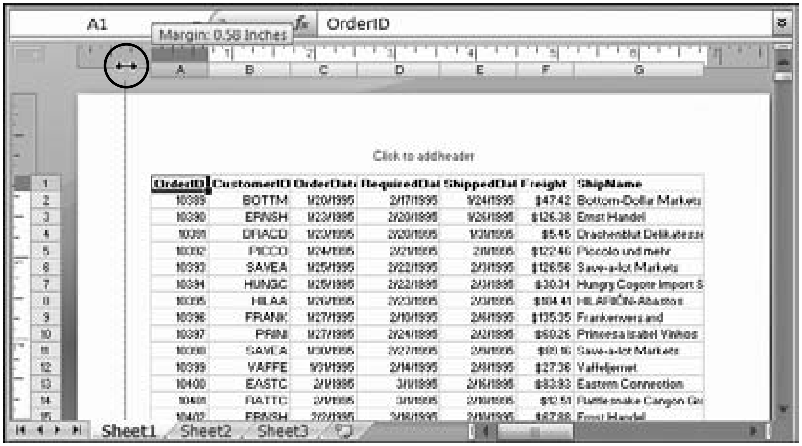

Figure 7-16: *The Page Layout view lets you set margins by dragging the margin edge with your mouse. Here, the left margin (circled) is about to be narrowed down to 0.58 inches. If you're also using a header or footer (below), make sure you don't drag the page margin above the header or below the footer. If you do, then your header or footer will overlap your worksheet's data.*

When you're ready to return to the Normal worksheet view, choose View→Workbook Views→Normal (or just click the Status bar's tiny Normal View button).

Creating Headers and Footers

A *header* is a bit of text that's printed at the top of every page in your printout. A *footer* is a bit of text that's printed at the bottom of every page. You can use one, both, or neither in a printout.

Ordinarily, every new workbook starts out without a header or footer. However, Page Layout view gives you an easy way to add either one (or both). Just scroll up to the top of any page to create a header (or the bottom to create a footer), and then look for the box with the text "Click to add header" or "Click to add footer". Click inside this box, and you can type the header or footer text you want.

Note: You won't see the header or footer boxes if you've drastically compressed your margins. That's because the header and footer don't fit. To get them back, resize the margins so that they're larger. When you're finished adding the header or footer, you can try adjusting the margins again to see just how small you can get them.

Of course, a good header or footer isn't just an ordinary piece of text. Instead, it contains information that changes dynamically, like the file name, current page, or the date you printed it. You can get these pieces of information using specialized header and footer *codes*, which are distinguished by their use of square brackets. For example, if you enter the code *[Page]* into a footer, Excel replaces it with the current page number. If you use the code *[Date]*, Excel substitutes the current date (when you fire off your printout). Of course, no one wants to memorize a long list of cryptic header and footer codes. To help you get these important details right, Excel adds a new tab to the ribbon named Header & Footer Tools | Design (Figure 7-17) when you edit a header or footer.

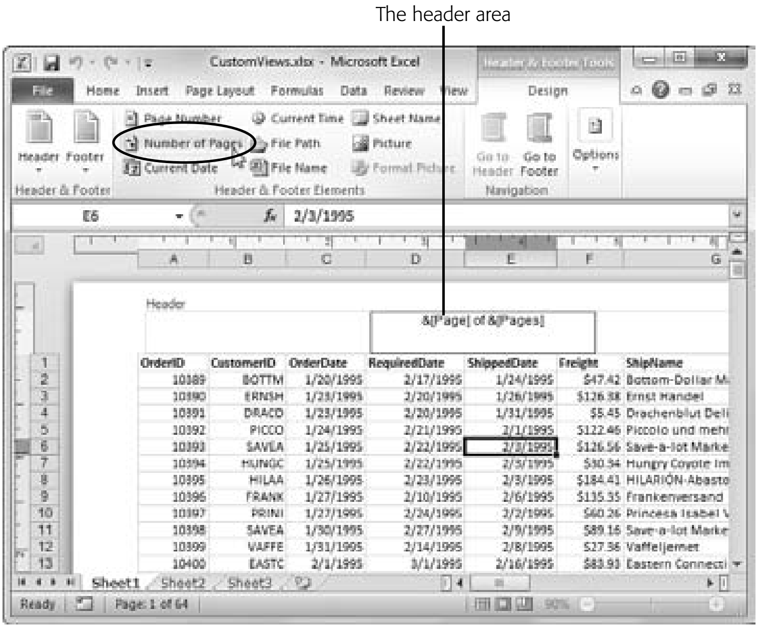

Figure 7-17: *The Header & Footer Tools | Design tab is chock-full of useful ingredients you can add to a header or footer. Click a button in the Header & Footer Elements section to insert a special Excel code that represents a dynamic value, like the number of pages (circled).*

The quickest way to get a header or footer is to go to the Header & Footer Tools | Design→Header & Footer section (shown in Figure 7-17), and then choose one of the Header or Footer list's ready-made options. Some of the options you can use for a header or footer include:

- Page numbering (for example, Page 1 or Page 1 of 10).
- Worksheet name (for example, Sheet 1).
- File name (for example, *myfile.xlsx* or *C:\MyDocuments\myfile.xlsx*).
- The person who created the document, and the date it was created.
- A combination of this information.

Oddly enough, the header and footer options are the same. It's up to you to decide whether you want page numbering at the bottom and a title at the top, or vice versa.

If none of the standard options match what you need, you can edit the automatic header or footer, or you can create your own from scratch. Start typing in the header or footer box, and use the buttons in the Header & Footer Elements section to paste in the code you need for a dynamic value. And if you want to get more creative, switch to the Home tab of the ribbon, and then use the formatting buttons to change the font, size, alignment, and color of your header or footer.

Finally, Excel gives you a few high-powered options in the Header & Footer Tools | Design→Options section. These include:

- **Different First Page**. This option lets you create one header and footer for the first page, and use a different pair for all subsequent pages. Once you've checked this option, fill in the first page header and footer on the first page, and then head to the second page to create a new header and footer that Excel can use for all subsequent pages.
- **Different Odd & Even pages**. This option lets you create two different headers (and footers)—one for all even-numbered pages and one for all odd-numbered pages. (If you're printing a bunch of double-sided pages, you can use this option to make sure the page number appears in the correct corner.) Use the first page to fill in the odd-numbered header and footer, and then use the second page to fill in the even-numbered header and footer.
- **Scale with Document**. If you select this option, then when you change the print scale to fit in more or less information on your printout (page 223), Excel adjusts the headers and footers proportionately.
- **Align with Page Margins**. If you select this option, Excel moves the header and footer so that they're centered in relation to the margins. If you don't select this option, Excel centers them in relation to the whole page. The only time you'll notice a difference is if your left and right margins are significantly different sizes.

All these settings affect both headers and footers.

Controlling Pagination

Sooner or later it will happen to you—you'll face an intimidatingly large worksheet that, when printed, is hacked into dozens of apparently unconnected pages. You could spend a lot of time assembling this jigsaw printout (using a bulletin board and lots of tape), or you could take control of the printing process and tell Excel exactly where to split your data into pages. In the following sections, you'll learn several techniques to do just that.

Page Breaks

One of Excel's often overlooked but surprisingly handy features is *manual page breaks*. The idea is that you tell Excel explicitly where to start a new page. For example, you can tell Excel to start a new page between subsequent tables on a worksheet (rather than print a page that has the end of the first one and the beginning of the next).

To insert a page break, move to the leftmost column (column A), and then scroll down to the first cell that you want to appear on the new page. Then, choose Page Layout→Page Setup→Breaks→Insert Page Break. You see a dotted line that indicates the dividing lines in between pages (Figure 7-18).

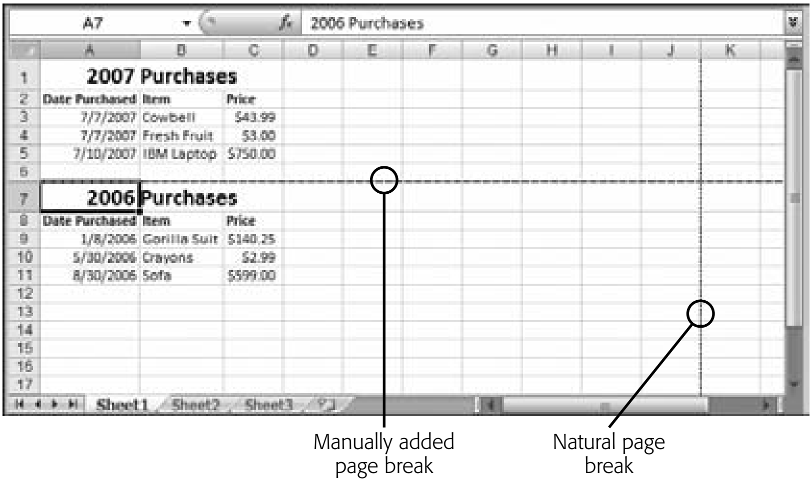

Figure 7-18: *Using a page break, you can make sure the second table (2010 Purchases) always begins on a new page. The dotted line shows where one page ends and the new page starts. When you add a page break, you see a dotted line for it, and you see a dotted line that shows you where additional page breaks naturally fall, based on your margins, page orientation, and paper size settings.*

Note: There's no limit to how many page breaks you can add to a worksheet–if you have a dozen tables that appear one after the other, you can place a page break after each one to make sure they all start on a new page.

You can also insert page breaks to split your worksheet vertically into pages. This is useful if your worksheet is too wide to fit on one page, but you want to control exactly where the page break will fall. To do so, move to the first row, scroll to the column where the new page should begin, and then choose Page Layout→Page Setup→Breaks→Insert Page Break.

You can remove page breaks one at a time by moving to an adjacent cell and choosing Page Layout→Page Setup→Breaks→Remove Page Break. Or you can clear them all using Page Layout→Page Setup→Breaks→Reset All Page Breaks.

Scaling

Page breaks are a nifty feature for making sure you paginate your printouts just the way you want them. However, they can't help you fit more information on a page. They simply let you place page breaks earlier than they would ordinarily appear, so they fall in a more appropriate place.

If you want to fit more on a page, you need to shrink your information down to a smaller size. Excel includes a scaling feature that lets you take this step easily without forcing you to reformat your worksheet.

Scaling lets you fit more rows and columns on a page, by shrinking everything proportionally. For example, if you reduce scaling to 50 percent, you fit twice as many columns and rows on a page. Conversely, you can use scaling to enlarge your data.

To change the scaling percentage, just type a new percentage into the Page Layout→Scale to Fit→Scale box. The data still appears just as big on your worksheet, but Excel shrinks or expands it in the printout. To gauge the effect, you can use the Page Layout view (page 217) to preview your printout.

Rather than fiddling with the scaling percentage (and then seeing what its effect is on your worksheet by trial and error), you may want to force your data to fit into a fixed number of pages. To do this, you set the values in the Page Layout→Scale to Fit→Width box and the Page Layout→Scale to Fit→Height box. Excel performs a few behind-the-scenes calculations and adjusts the scaling percentage accordingly. For example, if you choose one page tall and one page wide, then Excel shrinks your entire worksheet so that everything fits into one page. This scaling is tricky to get right (and can lead to hopelessly small text), so make sure you review your worksheet in the Page Layout view before you print it.

Tip: Page Break Preview mode, described next, gives you yet another way to squeeze more data onto a single page.

Page Break Preview: A Bird's-Eye View of Your Worksheet

You don't must be a tree-hugging environmentalist to want to minimize the number of pages you print out. Enter the Page Break Preview, which gives you a bird's-eye view of how an entire worksheet's going to print. Page Break Preview is particularly useful if your worksheet has lots of columns. That's because Page Break Preview zooms out so you can see a large amount of data at once, and it uses thick blue dashed lines to show you where page breaks will happen, as shown in Figure 7-19. In addition, the Page Break Preview numbers every page, placing the label "Page X" (where "X" is the page number) in large gray lettering in the middle of each page.

To preview the page breaks in your data, select View→Workbook Views→Page Break Preview, or use the tiny Page Break Preview button in the Status bar. A window appears, informing you that you can use Page Break Preview mode to move page breaks. You can choose whether you want to see this message again; if not, turn on the "Do not show this dialog again" checkbox before clicking OK.

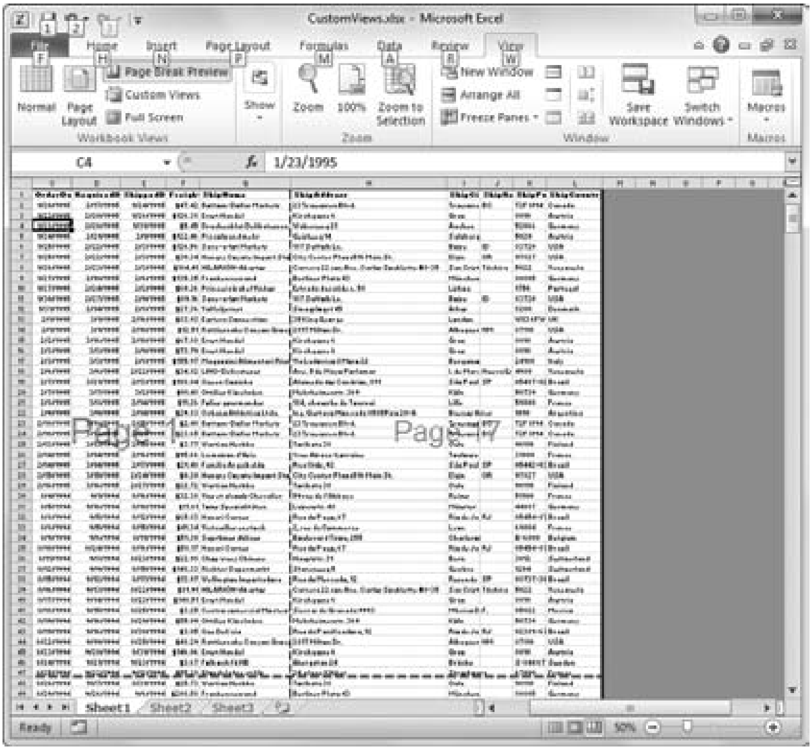

Figure 7-19: *This example shows a large worksheet in Page Break Preview mode. The worksheet is too wide to fit on one page (at least in portrait orientation), and the thick dotted line clearly indicates that the page breaks after column G and after row 47. (Excel never breaks a printout in the middle of a column or row.)*

Once you're in Page Break Preview mode, you can do all of the things you do in Normal view mode, including editing data, formatting cells, and changing the zoom percentage to reveal more or fewer pages. You can also click the blue dashed lines that represent page breaks, and drag them to include more or less rows and columns in your page.

Excel lets you make two types of changes using page breaks:

- **You can make *less* data fit onto a page**. To do so, drag the bottom page break up or the right-side page break to the left. Usually, you'll perform these steps if you notice that a page break is in an awkward place, like just before a row with some kind of summary or subtotal.
- **You can make *more* data fit onto a page**. To do so, drag the bottom page break down or the right-side page break to the right.

Of course, everyone wants to fit more information onto their printouts, but there's only so much space on the page. So what does Excel do when you expand a page by dragging the page break? It simply adjusts the scaling setting you learned about earlier (page 223). The larger you make the page, the smaller the Scaling percentage setting becomes. That means your printed text may end up too tiny for you to read. (The text on your computer's display doesn't change, however, so you don't have any indication of just how small your text has become until you print out your data, or take a look at it in Page Layout view.)

Note: Scaling affects all the pages in your printout. That means when you drag one page break to expand a page, you actually end up compressing *all* the pages in your workbook. However, the page *breaks* don't change for other pages, which means you may end up with empty, unused space on some of the pages.

The best advice: If your goal is merely to fit more information into an entire printout, change the scaling percentage manually instead of using the Page Break Preview. On the other hand, if you need to squeeze just a little bit more data onto a specific page, use the Page Break Preview.

CHAPTER 8

Building Basic Formulas

Most Excel fans don't turn to the world's leading spreadsheet software just to create nicely formatted tables. Instead, they rely on Excel's industrial-strength computing muscle, which lets you reduce reams of numbers to neat subtotals and averages. Performing these calculations is the first step to extracting meaningful information out of raw data.

Excel provides a number of different ways to build formulas, letting you craft them by hand or point-and-click them into existence. In this chapter, you'll learn about all of these techniques. You'll start by examining the basic ingredients that make up any formula, and then take a close look at the rules Excel uses when evaluating a formula.

Creating a Basic Formula

First things first: what exactly do formulas do in Excel? A *formula* is a series of instructions that you place in a cell in order to perform some kind of calculation. These instructions may be as simple as telling Excel to sum up a column of numbers, or they may incorporate advanced statistical functions to spot trends and make predictions. But in all cases, all formulas share the same basic characteristics:

- You enter each formula into a single cell.
- Excel calculates the result of a formula every time you open a spreadsheet or change the data a formula uses.
- Formula results are usually numbers, although you can create formulas that have text or Boolean (true or false) results.

- To view any formula (for example, to gain some insight into how Excel produced a displayed result), you must move to the cell containing the formula, and then look in the *formula bar* (see Figure 8-1). The formula bar also doubles as a handy tool for editing your formulas.
- Formulas can evaluate a combination of numbers you input (useful when you want to use Excel as a handy calculator) or, more powerfully, the contents of other cells. Formulas can even process an entire group of cells when using certain functions.

One of the simplest formulas you can create is this one:

```
=1+1
```

The equal sign is how you tell Excel that you're entering a formula (as opposed to a string of text or numbers). The formula that follows is what you want Excel to calculate. Note that the formula doesn't include the *result*. When creating a formula in Excel, you write the question, and then Excel coughs up the answer, as shown in Figure 8-1.

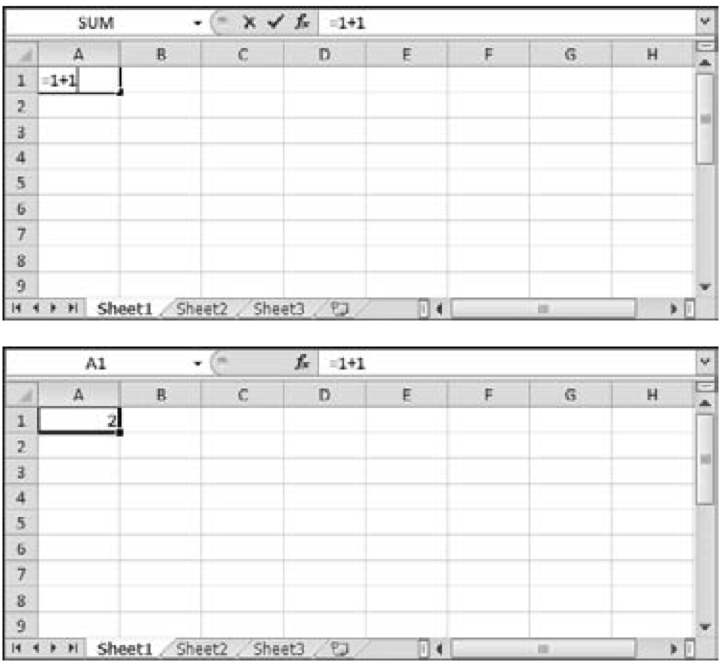

Figure 8-1:
Top: This simple formula begins its life when you enter it into a cell. The checkmark and X buttons to the left of the formula bar let you quickly complete or cancel, respectively, your formula.

Bottom: Or you can press Enter, and Excel displays the result in the cell. The formula bar always displays the complete formula (=1+1). In formula lingo, this particular example consists of two literal values (1 and 1) and one arithmetic operator (+).

All formulas use some combination of the following ingredients:

- **The equal sign** (=). Every formula must begin with the equal sign. It signals to Excel that the cell contains a formula, not just ordinary text.

- **The simple operators**. These ingredients include everything you fondly remember from high-school math class, including addition (+), subtraction (–), multiplication (*), division (/), and exponentiation (^). Table 8-1 lists these ingredients, also known as *arithmetic operators*.
- **Numbers**. These ingredients are known as constants or *literal values*, because they never change (unless you edit the formula).
- **Cell references**. These references point to another cell, or a range of cells, that you need data from in order to perform a calculation. For example, say you have a list of 10 numbers. A formula in the cell beneath this list may refer to all 10 of the cells above it in order to calculate their average.
- **Functions**. Functions are specialized formulas built into Excel that let you perform a wide range of calculations. For example, Excel provides dedicated functions that calculate sums and averages, standard deviations, yields, cosines and tangents, and much more. The next four chapters describe these functions, which span every field from financial accounting to trigonometry.
- **Spaces**. Excel ignores these. However, you can use them to make a formula easier to read. For example, you can write the formula *=3*5 + 6*2* instead of *=3*5+6*2*. (The only exception to this rule applies to cell ranges, where spaces have a special meaning. You'll see this described on page 237.)

Table 8-1. *Excel's arithmetic operators*

Operator	Name	Example	Result
+	Addition	=1+1	2
–	Subtraction	=1–1	0
*	Multiplication	=2*2	4
/	Division	=4/2	2
^	Exponentiation	=2^3	8
%	Percent	=20%	0.20

Note: The percentage (%) operator divides a number by 100.

Excel's Order of Operations

For computer programs and human beings alike, one of the basic challenges when it comes to reading and calculating formula results is figuring out the *order of operations*—mathematician-speak for deciding which calculations to perform first when there's more than one calculation in a formula. For example, given the formula:

```
=10 - 8 * 7
```

the result, depending on your order of operations, is either 14 or –46. Fortunately, Excel abides by the standard rules for order of operations, meaning it doesn't necessarily process your formulas from left to right. Instead, it evaluates complex formulas piece-by-piece in this order:

1. **Parentheses (any calculations within parentheses are always performed first)**
2. **Percent**
3. **Exponents**
4. **Division and Multiplication**
5. **Addition and Subtraction**

Note: When Excel encounters formulas that contain operators of equal *precedence* (that is, the same order of operation priority level), it evaluates these operators from left to right. However, in basic mathematical formulas, this has no effect on the result.

For example, consider the following formula:

```
=5 + 2 * 2 ^ 3 - 1
```

To arrive at the answer of 20, Excel first performs the exponentiation (2 to the power of 3):

```
=5 + 2 * 8 - 1
```

And then the multiplication:

```
=5 + 16 - 1
```

And then the addition and subtraction:

```
=20
```

To control this order, you can add parentheses. For example, notice how adding parentheses affects the result in the following formulas:

```
5 + 2 * 2 ^ (3 - 1) = 13
(5 + 2) * 2 ^ 3 - 1 = 55
(5 + 2) * 2 ^ (3 - 1) = 28
5 + (2 * (2 ^ 3)) - 1 = 20
```

You must always use parentheses in pairs (one open parenthesis for every closing parenthesis). If you don't, then Excel gets confused and lets you know you need to fix things, as shown in Figure 8-2.

Tip: Remember, when you're working with a lengthy formula, you can expand the formula bar to see several lines at a time. To do so, click the down arrow at the far right of the formula bar (to make it three lines tall), or drag the bottom edge of the formula bar to make it as many lines large as you'd like. Page 32 shows an example.

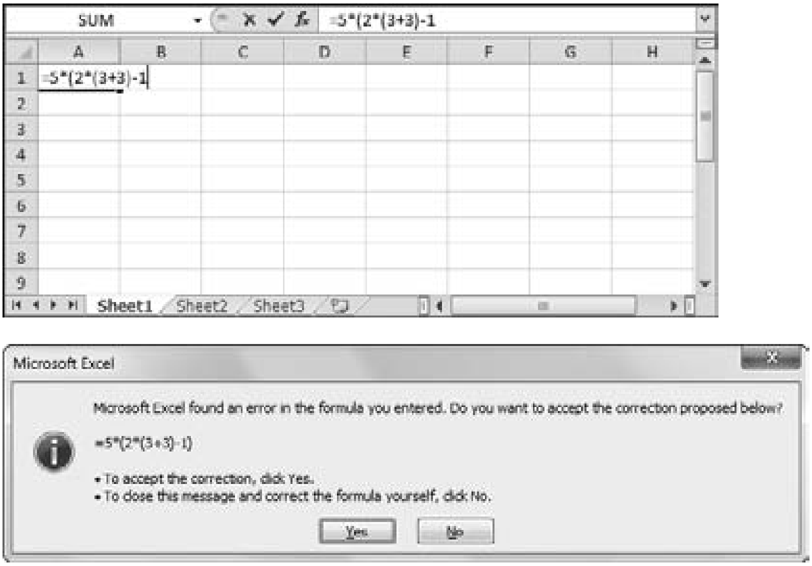

Figure 8-2:
Top: If you create a formula with a mismatched number of opening and closing parentheses (like this one), Excel won't accept it.

Bottom: Excel offers to correct the formula by adding the missing parentheses at the end. You may not want this addition, though. If not, cancel the suggestion, and then edit your formula by hand. Excel helps a bit by highlighting matched sets of parentheses. For example, as you move to the opening parenthesis, Excel automatically bolds both the opening and closing parentheses in the formula bar.

Cell References

Excel's formulas are handy when you want to perform a quick calculation. But if you want to take full advantage of Excel's power, then you're going to want to use formulas to perform calculations on the information that's already in your worksheet. To do that you need to use *cell references*—Excel's way of pointing to one or more cells in a worksheet.

For example, say you want to calculate the cost of your Amazonian adventure holiday, based on information like the number of days your trip will last, the price of food and lodging, and the cost of vaccination shots at a travel clinic. If you use cell references, you can enter all this information into different cells, and then write a formula that calculates a grand total. This approach buys you unlimited flexibility because you can change the cell data whenever you want (for example, turning your three-day getaway into a month-long odyssey), and Excel automatically refreshes the formula results.

Cell references are a great way to save a *ton* of time. They come in handy when you want to create a formula that involves a bunch of widely scattered cells whose values frequently change. For example, rather than manually adding up a bunch of subtotals to create a grand total, you can create a grand total formula that uses cell references to point to a handful of subtotal cells. They also let you refer to large groups of cells by specifying a *range*. For example, using the cell reference lingo you'll learn on page 237, you can specify all the cells in the first column between the second and 100th rows.

Every cell reference points to another cell. For example, if you want a reference that points to cell A1 (the cell in column A, row 1), then use this cell reference:

```
=A1
```

In Excel-speak, this reference translates to "get the value from cell A1, and insert it in the current cell." So if you put this formula in cell B1, then it displays whatever value's currently in cell A1. In other words, these two cells are now linked.

Cell references work within formulas just as regular numbers do. For example, the following formula calculates the sum of two cells, A1 and A2:

```
=A1+A2
```

Note: In Excel lingo, A1 and A2 are *precedents*, which means another cell needs them to perform a calculation. Cell B1, which contains the formula, is called the *dependent*, because it depends on A1 and A2 to do its work. These terms become important when you need to hunt for errors in a complex calculation using Excel's error-checking tools (page 392).

Provided both cells contain numbers, you'll see the total appear in the cell that contains the formula. If one of the cells contains text, then you'll see an error code instead that starts with a # symbol. Errors are described in more detail on page 242.

Note: This chapter focuses on how to perform calculations using cells that contain ordinary numbers. Excel also lets you manipulate other types of content in a formula, like text and dates. You'll learn more about these topics in Chapter 11.

How Excel Formats Cells That Contain Cell References

As you learned in Chapter 5, the way you format a cell affects how Excel displays the cell's value. When you create a formula that references other cells, Excel attempts to simplify your life by applying automatic formatting. It reads the number format that the *source cells* (that is, the cells being referred *to*) use, and applies that to the cell with the formula. If you add two numbers and you've formatted both with the Currency number format, then your result also has the Currency number format. Of course, you're always free to change the formatting of the cell after you've entered the formula.

GEM IN THE ROUGH

Excel As a Pocket Calculator

Sometimes you need to calculate a value before you enter it into your worksheet. Before you reach for your pocket calculator, you may like to know that Excel lets you enter a formula in a cell, and then use the result in that same cell. This way, the formula disappears and you're left with the result of the calculated value.

Start by typing your formula into the cell (for example *=65*88*). Next, press F9 to perform the calculation. Finally, just hit Enter to insert this value into the cell.

Remember, when you use this technique, you replace your formula with the calculated value. If your calculation is based on the values of other cells, then Excel won't update the result if you change those other cells' values. That's the difference between a cell that has a value, and a cell that has a formula.

Excel has a similar trick that's helpful if you want to take a whole batch of formulas (in different cells), and replace them all with values. It's the Paste Values command. To try it out, select the cells that have the formulas you want to change, copy them (Home→Clipboard→Copy), and then paste them somewhere in your worksheet using the Home→Clipboard→Paste→Paste Values command. The pasted cells now have the numbers, not the formulas.

Usually, Excel's automatic formatting is quite handy. Like all automatic features, however, it's a little annoying if you don't understand how it works when it springs into action. Here are a few points to consider:

- Excel copies only the number format to the formula cell. It ignores other details, like fonts, fill colors, alignment, and so on. (Of course, you can manually copy formats using the Format Painter, as discussed on page 173.)
- If your formula uses more than one cell reference, and the different cells use different number formats, Excel uses its own rules of precedence to decide which number format to use. For example, if you add a cell that uses the Currency number format with one that uses the Scientific number format, then the destination cell has the Scientific number format. Sadly, these rules aren't spelled out anywhere, so if you don't see the result you want, it's best to just set your own formatting.
- If you change the formatting of the source cells *after* you enter the formula, it won't have any effect on the formula cell.
- Excel copies source cell formatting only if the cell that contains the formula uses the General number format (which is the format that all cells begin with). If you apply another number format to the cell *before* you enter the formula, then Excel doesn't copy any formatting from the source cells. Similarly, if you change a formula to refer to new source cells, then Excel doesn't copy the format information from the new source cells.

Functions

A good deal of Excel's popularity is due to the collection of *functions* it provides. Functions are built-in, specialized algorithms that you can incorporate into your own formulas to perform powerful calculations. Functions work like miniature computer programs—you supply the data, and the function performs a calculation and gives you the result.

In some cases, functions just simplify calculations that you could probably perform on your own. For example, most people know how to calculate the average of several values, but when you're feeling a bit lazy, Excel's built-in AVERAGE() function automatically gives you the average of any cell range. Even more usefully, Excel functions perform feats that you probably wouldn't have a hope of coding on your own, including complex mathematical and statistical calculations that predict *trends*—hidden relationships in your data that you can use to make guesses or predict the future.

Tip: You can create your own Excel functions by writing a series of instructions using VBA (Visual Basic for Applications) code. Chapter 29 shows you how.

Every function provides a slightly different service. For example, one of Excel's statistical functions is named COMBIN(). It's a specialized tool used by probability mathematicians to calculate the number of ways a set of items can be combined. Although this sounds technical, even ordinary folks can use COMBIN() to get some interesting information. You can use the COMBIN() function, for example, to count the number of possible combinations there are in certain games of chance.

The following formula uses COMBIN() to calculate how many different five-card combinations there are in a standard deck of 52 playing cards:

```
=COMBIN(52,5)
```

Functions are always written in all capitals. (More in a moment on what those numbers inside the parentheses are doing.) However, you don't need to worry about the capitalization of function names, because Excel automatically capitalizes the function names that you type in.

Using a Function in a Formula

Functions alone don't actually *do* anything in Excel. Functions need to be part of a formula to produce a result. For example, COMBIN() is a function name. But it actually *does* something—that is, give you a result—only when you've inserted it into a formula, like so: *=COMBIN(52,5)*.

Whether you're using the simplest or the most complicated function, the *syntax*—or, rules for including a function within a formula—is always similar. To use a function, start by entering the function name. Excel helps you out by showing a pop-up list with possible candidates as you type, as shown in Figure 8-3. This handy feature is called Formula AutoComplete.

UP TO SPEED

Learning New Functions

This book will introduce you to dozens of new functions. Sometimes you'll start off by looking at a sample formula that uses the function, but for more complex functions, start by considering the *function description*.

The function description assigns a name to each argument. You can learn about the type of data the function requires before you start wading into an example with real numbers. For example, here's the function description for the COMBIN() function:

```
COMBIN(number_in_set, number_chosen)
```

You can tell the difference between a sample formula and the function description by the fact that the function description doesn't include the initial equal sign (=) that you need in all formulas.

Sometimes a function takes an *optional argument*. The argument isn't required, but it may be important depending on the behavior you want. Optional arguments are always shown in square brackets. (Excel uses the same convention in its help and formula tooltips.)

You'll see plenty of function descriptions in this book. You can look up function descriptions in Excel. Page 252 tells you where to look.

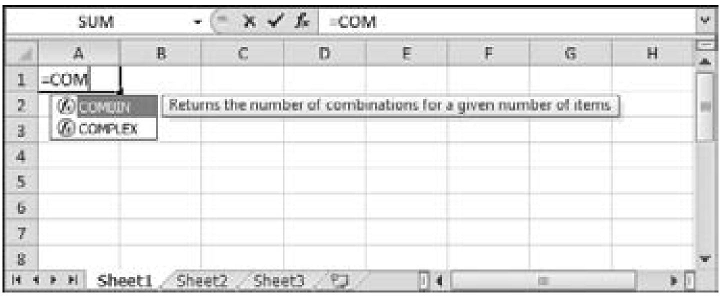

Figure 8-3: *After you type =COM, Excel helpfully points out that it knows only two functions that start that way: COMBIN() and COMPLEX(). If your fingers are getting tired, use the arrow keys to pick the right one out of the list, and then click Tab to pop it into your formula. (Or, you can just double-click it.)*

After you type the function name, add a pair of parentheses. Then, inside the parentheses, put all the information the function needs to perform its calculations.

In the case of the COMBIN() function, Excel needs two pieces of information, or *arguments*. The first is the number of items in the set (the 52-card deck), and the second's the number of items you're randomly selecting (in this case, 5). Most functions, like COMBIN(), require two or three arguments. However, some functions can accept many more, while a few don't need any arguments at all. Once again, Formula AutoComplete guides you by telling you what arguments you need, as shown in Figure 8-4.

Once you type this formula into a cell, the result (2598960) appears in your worksheet. In other words, there are 2,598,960 different possible five-card combinations in any deck of cards. Rather than having to calculate this fact using probability theory—or, heaven forbid, trying to count out the possibilities manually—the COMBIN() function handled it for you.

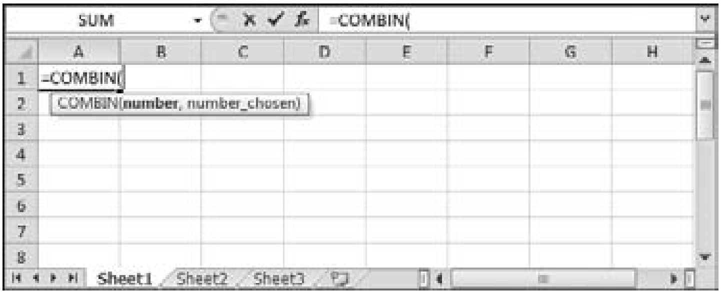

Figure 8-4: *When you type the opening parenthesis after a function name, Excel automatically displays a tooltip indicating what arguments the function requires. The argument you're currently entering is bolded in the tooltip. The argument names aren't crystal clear, but if you already know how the function works, they're usually enough to jog your memory.*

Note: Even if a function doesn't take any arguments, you still need to supply an empty set of parentheses after the function name. One example is the RAND() function, which generates a random fractional number. The formula =RAND() works fine, but if you forget the parentheses and merely enter *=RAND*, then Excel displays an error message (#*NAME*?) that's Excelian for: "Hey! You got the function's name wrong." See Table 8-2 for more information about Excel's error messages.

UP TO SPEED

Understanding Functions

Even though it's relatively easy to understand the basics behind how functions work and how to combine them in a formula, that doesn't mean you'll understand what all functions do and *why* you should use a particular function. If you don't already know a little probability theory, for instance, then the COMBIN() function may not be very useful. Excel's packed full of advanced formulas like COMBIN() that are tailored for statisticians, accountants, and mathematicians. Most of these you'll never need to use.

In this book, functions that'll likely be useful to a nonspecialist reader are explained completely. For example, you may not know the financial term *net present value*, but you'll probably still be interested in using Excel's NPV() function to calculate the value of your investments. On the other hand, if you don't know the meaning of a *complex conjugate*—an abstract concept used in some engineering calculations—you won't be interested in the IMCONJUGATE() function.

This book won't explain the math behind these more specialized functions. (In fact, properly explaining some of these concepts would require at least a whole other book!) Instead, these functions will be briefly pointed out in a note or table in the relevant chapter. This way, you can easily find these functions if they're relevant to your work and you already know the underlying math or statistical concepts that power them.

Using Cell References with a Function

One of the particularly powerful things about functions is that they don't necessarily need to use literal values in their arguments. They can also use cell references. For example, you could rewrite the five-card combination formula (mentioned earlier) so that it specifies the number of cards that'll be drawn from the deck based on a number that you've typed in somewhere else in the spreadsheet. Assuming this information's entered into cell B2, the formula would become:

```
=COMBIN(52,B2)
```

Building on this formula, you can calculate the probability (albeit astronomically low) of getting the exact hand you want in one draw:

```
=1/COMBIN(52,B2)
```

You could even multiply this number by 100 or use the Percent number style to see your percentage chance of getting the cards you want.

Tip: Excel gives you a detailed function reference to find functions and learn about them. Excel's information doesn't make for light reading, though; for the most part, it's in IRS-speak. You'll learn more about using this reference on page 252.

Using Cell Ranges with a Function

In many cases, you don't want to refer to just a single cell, but rather a *range* of cells. A range is simply a grouping of multiple cells. These cells may be next to each other (say, a range that includes all the cells in a single column), or they could be scattered across your worksheet. Ranges are useful for computing averages, totals, and many other calculations.

To group together a series of cells, use one of the three following reference operators:

- **The comma (,) separates more than one cell**. For example, the series *A1, B7, H9* is a cell range that contains three cells. The comma's known as the *union operator*. You can add spaces before or after a comma, but Excel just ignores or removes them (depending on its mood).
- **The colon (:) separates the top-left and bottom-right corners of a block of cells**. You're telling Excel: "Hey, use *this* block of cells in my formula." For example, *A1:A5* is a range that includes cells A1, A2, A3, A4, and A5. The range *A2:B3* is a grid that contains cells A2, A3, B2, and B3. The colon is the *range operator*—by far the most powerful way to select multiple cells.
- **The space can find cells that are common to two or more different cell ranges**. For example, the expression *A1:A3 A1:B10* is a range that consists of only three cells: A1, A2, and A3 (because those three cells are the only ones found in both ranges). The space is technically the *intersection operator*, and it's not used terribly often.

Tip: As you might expect, Excel lets you specify ranges by selecting cells with your mouse, instead of typing in the range manually. You'll see this trick later in this chapter on page 245.

You can't enter ranges directly into formulas that just use the simple operators. For example, the formula *=A1:B1+5* doesn't work, because Excel doesn't know what to do with the range A1:B1. (Should it sum up the range? Or average it? Excel has no way of knowing.) Instead, you need to use ranges with functions that know how to use them. For instance, one of Excel's most basic functions is named SUM(); it calculates the total for a group of cells. To use the SUM() function, you enter its name, an open parenthesis, the cell range you want to add up, and then a closing parenthesis.

Here's how you can use the SUM() function to add together three cells, A1, A2, and A3:

```
=SUM(A1,A2,A3)
```

And here's a more compact syntax that performs the same calculation using the range operator:

```
=SUM(A1:A3)
```

A similar SUM() calculation's shown in Figure 8-5. Clearly, if you want to total a column with hundreds of values, it's far easier to specify the first and last cell using the range operator rather than including each cell reference in your formula!

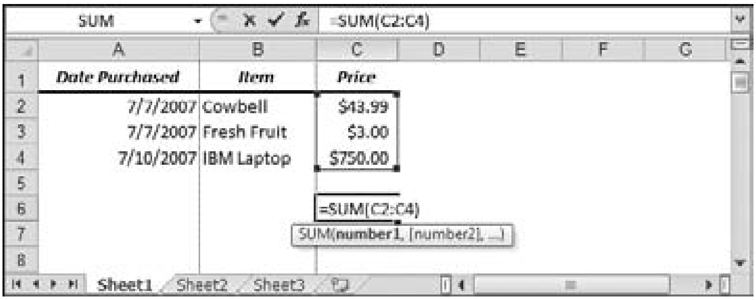

Figure 8-5:
Using a cell range as the argument in the SUM() function is a quick way to add up a series of numbers in a column. Note that when you enter or edit a formula, Excel highlights all the cells that formula uses with different colored borders. In this example, you see the range of cells C2, C3, and C4 in a blue box.

Sometimes your worksheet may have a list with unlimited growth potential, like a list of expenses or a catalog of products. In this case, you can code your formulas to include an *entire* column by leaving out the row number. For example, the range A:A includes all the cells in column A (and, similarly, the range 2:2 includes all the cells in row 2).

The range A:A also includes any heading cells, which isn't a problem for the SUM() function (because it ignores text cells), but could cause problems for other functions. If you don't want to include the top cell, then you need to think carefully about what you want to do. You could create a normal range that stretches from the second cell to the last cell using the mind-blowingly big range A2:A1048576. However, this could cause a problem with older versions of Excel, which don't support as many rows. You're better off creating a table (described in Chapter 14). Tables expand automatically, updating any linked formulas.

Excel Compatibility Functions

Some of Excel's functions use extraordinarily complex logic behind the scenes. Over the years, Excel experts have found minor flaws and quirks in some functions, like cases where the functions deviate from mathematical standards.

Correcting these problems is a bit messy. If different versions of Excel use subtly different calculation logic, you could find that your numbers change unpredictably—for example, when you upgrade your software or when you pass your spreadsheet to a colleague with a different version of Excel. In this case, consistency is more important than absolute, theoretical accuracy.

To avoid this sort of situation, the designers of Excel rarely change an existing function. Instead, in Excel 2010 they've added new functions that have similar names but incorporate a period. For example, in Excel 2010 the RANK.AVG() and RANK.EQ() functions replace the old-school RANK() function. Although RANK() still works, Microsoft recommends you use one of its replacements in new worksheets. (If you're curious, the replacements change how Excel ranks tied values, as you'll learn on page 273.) Because RANK() is kicking around only to ensure that old worksheets keep working, it's called a *compatibility function*.

So how do you recognize compatibility functions, to make sure you don't accidentally use one when you actually want the more modern replacement? The trick is to read the function tooltip, which clearly identifies compatibility functions, as shown in Figure 8-6.

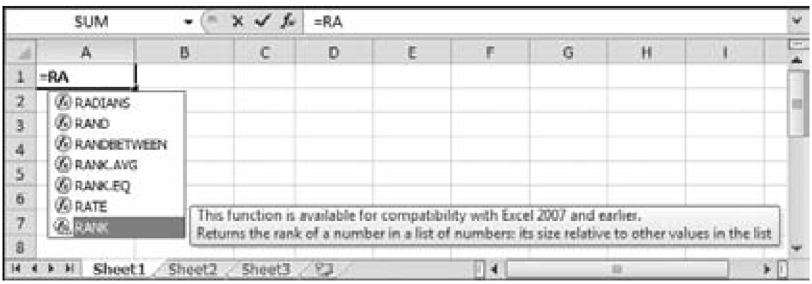

Figure 8-6: *The RANK() function is included for compatibility with old worksheets. Compatibility functions always appear at the bottom of the Formula AutoComplete list, and they're further distinguished by an icon with a yellow exclamation mark.*

Unfortunately, the case for ditching compatibility functions isn't as clear-cut as it seems. The problem is that the new functions won't work in older versions of Excel. For example, imagine you use a function like RANK.EQ() and send your spreadsheet to a colleague who's using Excel 2007. Because Excel doesn't know anything about this function, it can't evaluate the formula. Instead, it shows the infamous #NAME? error (page 242) in the cell.

So the bottom line is this: For now, it's safer to keep using compatibility functions instead of their newer replacements. But if a replacement has the behavior you really want, and you don't expect to share your work with other people, you can use the new functions.

Note: Almost all of the functions you'll learn about in the following chapters are traditional functions that have been with Excel for generations. You'll get a clear warning when we discuss new Excel 2010 functions (those with a period in their names), so you don't use them unknowingly. And if you're still paranoid, you can use the Compatibility Checker (page 41) to scan your worksheet for potential issues, including new functions that won't work in old versions of Excel.

Formula Errors

If you make a syntax mistake when entering a formula (like leaving out a function argument or including a mismatched number of parentheses), Excel lets you know right away. Moreover, like a stubborn schoolteacher, Excel won't accept the formula until you've corrected it. It's also possible, though, to write a perfectly legitimate formula that doesn't return a valid answer. Here's an example:

```
=A1/A2
```

If both A1 and A2 have numbers, this formula works without a hitch. However, if you leave A2 blank, or if you enter text instead of numbers, then Excel can't evaluate the formula, and it reminds you with an error message.

Excel lets you know about formula errors by using an *error code* that begins with the number sign (#) and ends with an exclamation point (!), as shown in Figure 8-7. In order to remove this error, you need to track down the problem and resolve it, which may mean correcting the formula or changing the cells it references.

When you click the exclamation mark icon next to an error, you see a menu of choices (as shown in Figure 8-7):

- **Help On This Error** pops open Excel's online help, with a (sometimes cryptic) description of the problem and what could have caused it.
- **Show Calculation Steps** pops open the Evaluate Formula dialog box, where you can work your way through a complex formula one step at a time. Page 388 describes how this advanced feature works.
- **Ignore Error** tells Excel to stop bothering you about this problem, in any worksheet you create. You won't see the green triangle for this error again (although you'll still see the error code in the cell).
- **Edit in Formula Bar** brings you to the formula bar, where you can change the formula to fix a mistake.

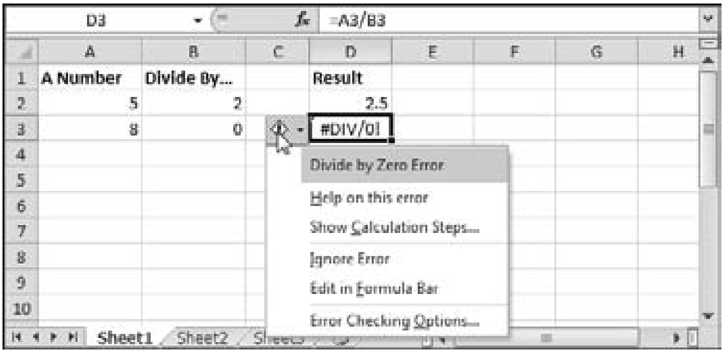

Figure 8-7: *When Excel spots an error, it inserts a tiny green triangle into the cell's top-left corner. When you move to the offending cell, Excel displays an exclamation mark icon next to it (a smart tag). Hover over the exclamation mark to view a description of the error (which appears in a tooltip), or click the exclamation icon to see a list of menu options.*

- **Error Checking Options** opens up the Excel Options dialog box, and brings you to the section where you can configure the settings Excel uses for alerting you about errors. You can turn off *background error checking*, or change the color of the tiny error triangles using the settings under the Error Checking heading. (Background error checking is the feature that flags cells with tiny green triangles when the cells contain a problem.) You can also tell Excel to start paying attention to errors you previously told it ignore by clicking the Reset Ignored Errors button. Underneath that button is a section named "Error checking rules" that has a number of options that focus on specific types of errors. For example, you can choose to have Excel ignore numbers stored as text, formulas that ignore part of a range, and other situations that technically aren't errors, but usually indicate that you've done something you didn't mean to. Excel *always* reports genuine errors, like #VALUE! and #NAME? regardless of what choices you make in this dialog box.

Note: Sometimes a problem isn't an error, but simply the result of data that hasn't yet been entered. In this case, you can solve the problem by using a conditional *error-trapping formula*. This conditional formula checks whether the data's present, and it performs the calculation only if it is. The next section shows one way to use an error-trapping formula.

Table 8-2 lists the error codes that Excel uses.

Table 8-2. *Excel's error codes*

Error code	Description
#VALUE!	You used the wrong type of data. Maybe your function expects a single value and you submitted a whole range. Or, more commonly, you might have used a function or created a simple arithmetic formula with a cell that contains text instead of numbers.
#NAME?	Excel can't find the name of the function you used. This error code usually means you misspelled a function's name, although it can indicate you used text without quotation marks or left out the empty parentheses after the function name. (You'll see how you can use text in a formula in Chapter 11.) This error will also happen if you try to use a new Excel 2010 function in older version of Excel.
#NUM!	There's a problem with one of the numbers you're using. For example, this error code appears when a calculation produces a number that's too large or too small for Excel to deal with.
#DIV/0	You tried to divide by zero. This error code also appears if you try to divide by a cell that's blank, because Excel treats a blank cell as though it contains the number 0 for the purpose of simple calculations with the arithmetic operators. (Some functions, like AVERAGE(), are a little more intelligent and ignore blank cells.)
#REF!	Your cell reference is invalid. This error most often happens if you delete or paste over the cells you were using, or if you try to copy a cell from one worksheet to another. (For information about creating formulas that span worksheets and workbooks, refer to Chapter 13.)
#N/A	The value isn't available. This error can happen if you try to perform certain types of lookup or statistical functions that work with cell ranges. For example, if you use a function to search a range and it can't find what you need, you may get this result. (You'll learn about lookup functions in Chapter 12.) Sometimes people enter a #N/A value manually in order to tell Excel to ignore a particular cell when creating charts and graphs. The easiest way to do this is to use the NA() function (rather than entering the text #N/A).
#NULL!	You used the intersection operator incorrectly. Remember, the intersection operator finds cells that two ranges share in common. This error results if there are no cells in common. Oftentimes, people use the intersection operator by accident, as the operator's just a single space character.
########	This code isn't actually an error condition—in all likelihood, Excel has successfully calculated your formula. However, the formula can't be displayed in the cell using the current number format. To solve this problem, you can widen the column, or possibly change the number format (page 141) if you require a certain number of fixed decimal places.

Tip: Chapter 13 describes a collection of Excel tools designed to help you track down the source of an error in a complex formula—especially one where the problem isn't immediately obvious.

TROUBLESHOOTING MOMENT

Circular References

One of the more aggravating problems that can happen with formulas is the infamous *circular reference*. A circular reference happens when you create a formula that depends, indirectly or directly, on its own value. For example, consider what happens if you enter the following formula in cell B1.

```
=B1+10
```

In order for this formula to work, Excel would need to take the current B1 value, and add 10. However, this operation *changes* the value of B1, which means Excel needs to recalculate the formula. If unchecked, this process would continue in an endless loop without ever producing a value.

More subtle forms of circular references are possible. For example, you can create a formula in one cell that refers to a cell in another cell that refers back to the original cell. This is what's known as an *indirect circular reference*, but the problem is the same.

Ordinarily, Excel doesn't allow circular references. When you enter a formula that contains a circular reference, Excel displays an error message and forces you to edit the formula until you've removed the circular reference. However, you can configure Excel to allow circular references by modifying the calculation settings, as described in the Formulas section of the Excel Options dialog box. In this case, Excel repeats the loop a fixed number of times, or until the value seems to settle down and stop changing.

Occasionally, this technique is useful for calculating certain types of approximations in advanced formulas. In most cases, though, this approach is rather dangerous because it means you don't catch accidental circular references, which can lead to invalid data. A better approach is to write a custom function that performs a calculation in a loop using VBA (Visual Basic for Applications), as described in Chapter 28.

Logical Operators

So far, you've seen the basic arithmetic operators (which are used for addition, subtraction, division, and so on) and the cell reference operators (used to specify one or more cells). There's one final category of operators that's useful when creating formulas: *logical operators*.

Logical operators let you build conditions into your formulas so the formulas produce different values depending on the value of the data they encounter. You can use a condition with cell references or literal values.

For example, the condition A2=A4 is true if cell A2 contains the same value as cell A4. On the other hand, if these cells contain different values (say 2 and 3), then the formula generates a false value. Using conditions is a stepping-stone to using conditional logic. Conditional logic lets you perform different calculations based on different scenarios.

For example, you can use conditional logic to see how large an order is, and provide a discount if the total order cost's over $5,000. Excel *evaluates* the condition, meaning it determines if the condition's true or false. You can then tell Excel what to do, based on that evaluation.

Table 8-3 lists all the logical operators you can use to build formulas.

Table 8-3. Logical operators

Operator	Name	Example	Result
=	Equal to	1=2	FALSE
>	Greater than	1>2	FALSE
<	Less than	1<2	TRUE
>=	Greater than or equal to	1>=1	TRUE
<=	Less than or equal to	1<=1	TRUE
<>	Not equal to	1<>1	FALSE

You can use logical operators to build standalone formulas, but that's not particularly useful. For example, here's a formula that tests whether cell A1 contains the number 3:

```
=(A2=3)
```

The parentheses aren't actually required, but they make the formula a little bit clearer, emphasizing the fact that Excel evaluates the condition first, and then displays the result in the cell. If you type this formula into the cell, then you see either the uppercase word TRUE or FALSE, depending on the content in cell A2.

On their own, logical operators don't accomplish much. However, they really shine when you start combining them with other functions to build conditional logic. For example, you can use the SUMIF() function, which totals the value of certain rows, depending on whether the row matches a set condition. Or you can use the IF() function to determine what calculation you should perform.

The IF() function has the following function description:

```
IF(condition, [value_if_true], [value_if_false])
```

In English, this line of code translates to: If the condition is true, display the second argument in the cell; if the condition is false, display the third argument.

Consider this formula:

```
=IF(A1=B2, "These numbers are equal", "These numbers are not equal")
```

This formula tests if the value in cell A1 equals the value in cell B2. If this is true, you'll see the message "These numbers are equal" displayed in the cell. Otherwise, you'll see "These numbers are not equal".

Note: If you see a quotation mark in a formula, it's because that formula uses text. You must surround all literal text values with quotation marks. (Numbers are different: You can enter them directly into a formula.)

People often use the IF() function to prevent Excel from performing a calculation if some of the data is missing. Consider the following formula:

```
=A1/A2
```

This formula causes a divide-by-zero error if A2 contains a 0 value. Excel then displays an error code in the cell. To prevent this from happening, you can replace this formula with the conditional formula shown here:

```
=IF(A2=0, 0, A1/A2)
```

This formula checks if cell A2 is empty or contains a 0. If so, the condition's true, and the formula simply gives you a 0. If it isn't, then the condition's false, and Excel performs the calculation A1/A2.

Practical examples of conditional logic abound in Chapters 12 and 13.

FREQUENTLY ASKED QUESTION

Showing and Printing Formulas

How in the world do I print out formulas that appear in my cells?

When you print a worksheet, Excel prints the calculated value in each cell rather than any formula that happens to be inside a cell. Usually, that's what you want to happen. But in some cases, rather than a printout of the formula's results, you want a record of the calculations used to generate the results.

Excel gives you a view setting so you can get this record. Just choose Formulas→Formula Auditing→Show Formulas. Now, Excel displays the cells' formula contents instead of the results—but on the current worksheet only. Excel also widens the columns so they can show more information (as formulas tend to be longer than their results). Repeat this process, and then uncheck the setting to return to normal life.

Formula Shortcuts

So far, you've learned how to build a formula by entering it manually. That's a good way to start out because it forces you to understand the basics of formula writing. But writing formulas by hand is a drag; plus, it's easy to type in the wrong cell address. For example, if you type A2 instead of A3, you can end up with incorrect data, and you won't necessarily notice your mistake.

As you become more comfortable with formulas, you'll find that Excel gives you a few tools—like point-and-click formula creation and the Insert Function button—to speed up your formula writing and reduce your mistakes. You'll learn about these features in the following sections.

Point-and-Click Formula Creation

Instead of entering a formula by typing it out letter-by-letter, Excel lets you create formulas by clicking the cells you want to use. For example, consider this simple formula that totals the numbers in two cells:

```
=A1+A2
```

To build this formula by clicking, just follow these steps:

1. **Move to the cell where you want to enter the formula.**

   This cell's where the result of your formula's calculation will appear. While you can pick any cell on the worksheet, A3 works nicely because it's directly below the two cells you're adding.

2. **Press the equal sign (=) key.**

   The equal sign tells Excel you're going to enter a formula.

3. **Move to the first cell you want to use in your formula (in this case, A1).**

   You can move to this first cell by pressing the up arrow key twice, or by clicking it with the mouse. You'll notice that moving to another cell doesn't cancel your edit, as it would normally, because Excel recognizes that you're building a formula. When you move to the new cell, the cell reference appears automatically in the formula (which Excel displays in cell A3, as well as in the formula bar just above your worksheet). If you move to another cell, Excel changes the cell reference accordingly.

4. **Press the + key.**

   Excel adds the + sign to your formula so that it now reads *=A1+*.

5. **Finish the formula by moving to cell A2, and then pressing Enter.**

   Again, you can move to A2 either by pressing the up arrow key or by clicking the cell directly. Remember, you can't just finish the formula by moving somewhere else; you need to press Enter to tell Excel you're finished writing the formula. Another way to complete your edit is to click the checkmark that appears on the formula bar, to the left of the current formula. Even experienced Excel fans get frustrated with this step. If you click another cell before you press Enter, then you won't move to the cell—instead, Excel inserts the cell into your formula.

Tip: You can use this technique with any formula. Just type in the operators, function names, and so on, and use the mouse to select the cell references. If you need to select a range of cells, then just drag your mouse to highlight the whole group of cells. You can practice this technique with the SUM() function. Start by typing *=SUM(* into the cell, and then selecting the range of cells you want to add. Finish by adding a final closing parenthesis, and then press Enter.

Point-and-Click Formula Editing

You can use a similar approach to edit formulas, although it's slightly trickier.

1. **Move to the cell that contains the formula you want to edit, and then put it in edit mode by double-clicking it or pressing F2.**

Excel highlights all the cells that this formula uses with a colored outline. Excel's even clever enough to use a helpful color-coding system. Each cell reference uses the same color as the outline surrounding the cell it's referring to. This can help you pick out where each reference is.

2. **Click the outline of the cell you want to change. (Your pointer changes from a fat plus sign to a four-headed arrow when you're over the outline.) With the mouse button still held down, drag this outline over to the new cell (or cells) you want to use.**

   Excel updates the formula automatically. You can also expand and shrink cell range references. To do so, put the formula-holding cell into edit mode, and then click any corner of the border that surrounds the range you want to change. Next, drag the border to change the size of the range. If you want to move the range, click any part of the range border, and then drag the outline in the same way as you would with a cell reference.

3. **Press Enter or click the formula bar checkmark to accept your changes.**

   That's it.

The Formulas Tab

The ribbon is stocked with a few buttons that make formula writing easier. To take a look, click the Formulas tab.

The most important part of the Formulas tab is the Function Library section at the left. It includes the indispensable Insert Function button, which you'll take for a spin in the next section. It also includes many more buttons that arrange Excel's vast catalog of functions into related categories for easier access. Figure 8-8 shows how it works.

The Function Library divides its functions into the following categories:

- **AutoSum** has a few shortcuts that let you quickly add, average, or otherwise deal with a list of numbers. You'll see how they work on page 269.
- **Recently Used** has exactly what you'd expect—functions that you've recently chosen from the Function Library. If you're just starting out with functions, you'll see that Excel fills the Recently Used list with a small set of commonly used functions, like SUM().
- **Financial** functions let you track your car loan payments and calculate how many more years until you can retire rich. You'll tackle them in Chapter 10.
- **Logical** functions let you create conditional logic for even smarter spreadsheets that make calculation decisions. You already had a quick introduction to conditional logic earlier in this chapter, but you'll get more details in Chapter 13.
- **Text** functions manipulate words, sentences, and other non-numeric information. Chapter 11 has the scoop.

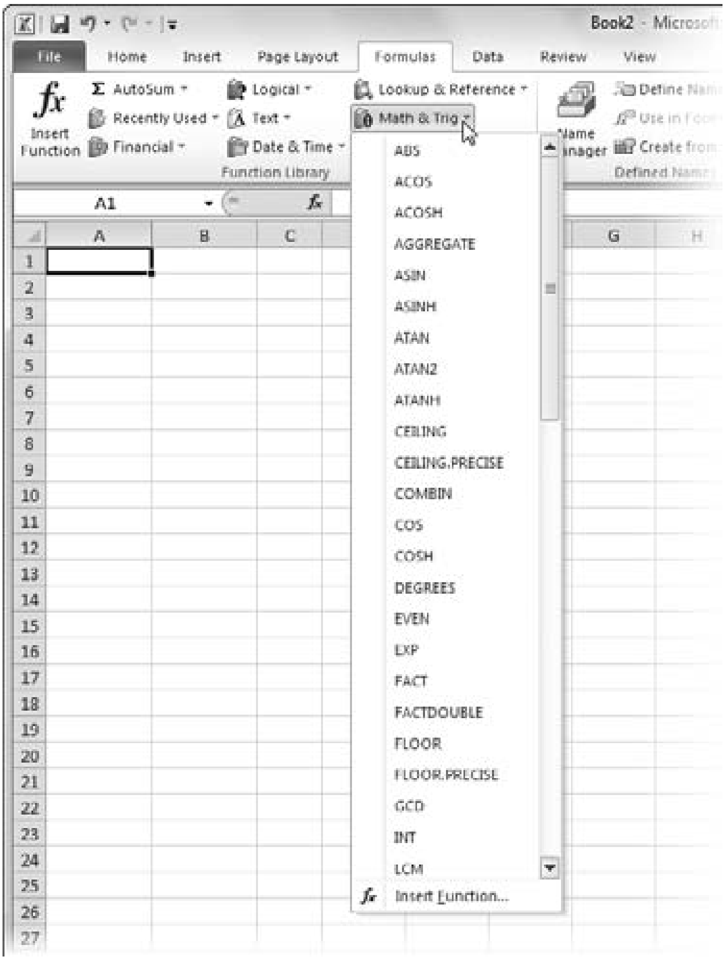

Figure 8-8:
Each button in the Function Library section (other than Insert Function) pops up a mini menu of function choices. Choose one, and Excel inserts that function into the current formula. You can use this technique to find functions that you've used recently, or to browse the main function categories. This example shows some of the functions in the Math & Trig section, which you'll explore in the next chapter.

- **Date & Time** functions perform calendar math, and can help you sort out ages, due dates, and more. You'll try them out in Chapter 11.
- **Lookup & Reference** functions perform the slightly mind-bending feat of searching for information in other cells. You'll see why they're so useful in Chapter 12.
- **Math & Trig** functions are the mathematic basics, including sums, rounding, and all the other high-school trigonometry you're trying to forget. Chapter 9 explores this section.
- **More Functions** groups together some heavy-duty Excel functions intended for specialized purposes. This category includes high-powered statistical and

engineering functions. It also includes *cube* functions, which are designed for working in highly technical OLAP (*online analytical processing*) databases, and information functions, which get "information" about the other cells in your worksheet. You won't see all of these functions in this book, because some are reserved for true specialists, but you'll dip your toe into a few statistical functions in Chapter 9, and you'll explore information functions in Chapter 12.

The Function Library section isn't the only part of the Formulas tab that you'll use. To the right are buttons for using named cells, tracking down errors, and changing calculation settings (three techniques you'll pick up in Chapter 13).

Using the Insert Function Button

Excel provides more than 300 built-in functions. In order to use a function, however, you need to type its name *exactly*. That means that every time you want to employ a function, you'll need to refer to this book, call on your own incredible powers of recollection, or click over to the convenient Insert Function button.

To use the Insert Function feature, choose Formulas→Function Library→Insert Function. However, formula pros skip straight to the action by clicking the *fx* button that appears just to the left of the formula bar. (Or, they press the Shift+F3 shortcut key.)

No matter which approach you use, Excel displays the Insert Function dialog box (shown in Figure 8-9), which offers three ways to search for and insert any of Excel's functions.

- If you're looking for a function, the easiest way to find one is to choose a category from the "Or select a category" drop-down list. For example, when you select the Math & Trig category, you see a list of functions with names like SIN() and COS(), which perform basic trigonometric calculations.
- If you choose the Most Recently Used category, you'll see a list of functions you've recently picked from the ribbon or the Insert Function dialog box.
- If you're really ambitious, you can type a couple of keywords into the "Search for a function" text box. Next, click Go to perform the search. Excel gives you a list of functions that match your keywords.

When you spot a function that looks promising, click it once to highlight its name. Excel then displays a brief description of the function at the bottom of the window. For more information, you can click the "Help on this function" link in the bottom-left corner of the window. To build a formula using this function, click OK.

Excel then inserts the function into the currently active cell, followed by a set of parentheses. Next, it closes the Insert Function dialog box, and then opens the Function Arguments dialog box (Figure 8-10).

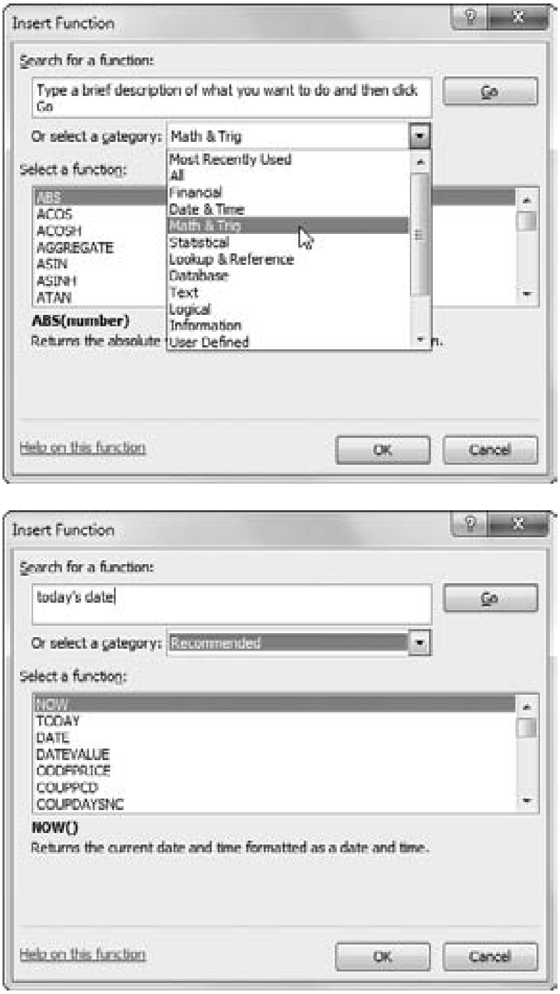

Figure 8-9:
Top: The Insert Function dialog box lets you quickly find the function you need. You can choose a category that seems likely to have the functions you're interested in.

Bottom: You can also try to search by entering keywords in the "Search for a function" box. Either way, when you click one of the functions in the list, Excel presents you with a description of the function at the bottom of the window.

Note: Depending on the function you're using, Excel may make a (somewhat wild) guess about which arguments you want to supply. For example, if you use the Insert Function dialog box to add a SUM() function, you'll see that Excel picks a nearby range of cells. If this isn't what you want, just replace the range with the correct values.

Now you can finish creating your formula by using the Function Arguments dialog box, which includes a text box for every argument in the function. It also includes a help link for detailed information about the function, as shown in Figure 8-11.

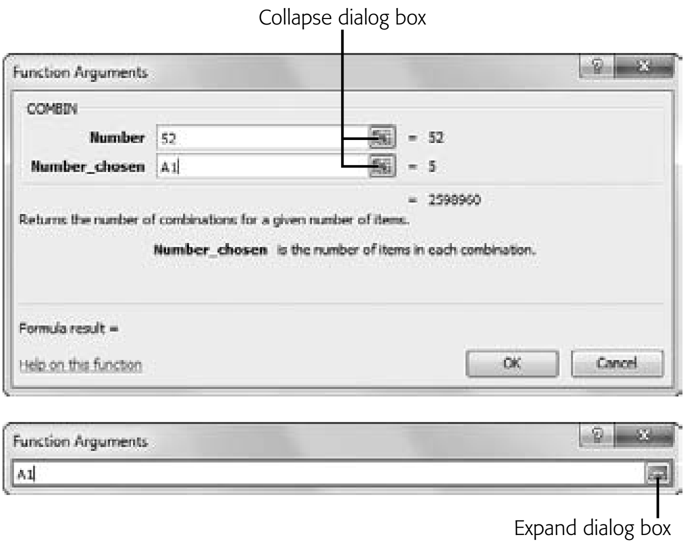

Figure 8-10:
Top: Here, the COMBIN() function has just been inserted via the Insert Function dialog box. Because the COMBIN() function requires two arguments (Number and Number_chosen), the Function Arguments dialog box shows two text boxes. The first argument uses a literal value (52), while the second argument uses a cell reference (A1). (You can use literal values or a cell reference for either argument–it's up to you.) As you enter the arguments, Excel updates the formula in the worksheet's active cell, and displays the result the calculation at the bottom of the Function Arguments dialog box.

Bottom: If you need more room to see the worksheet and select cells, you can click the Collapse Dialog Box icon to reduce the window to a single text box. Clicking the Expand Dialog Box icon restores the window to its normal size.

To complete your formula, follow these steps:

1. **Click the text box for the first argument.**

   A brief sentence describing the argument appears in the Function Arguments dialog box.

   Some functions don't require any arguments. In this case, you don't see any text boxes, although you still see some basic information about the function. Skip directly to step 4.

2. **Enter the value for the argument.**

   If you want to enter a literal value (like the number 52), type it in now. To enter a cell reference, you can type it in manually, or click the appropriate cell on the worksheet. To enter a range, drag the cursor to select a group of cells.

   You may need to move the Function Arguments dialog box to the side to expose the part of the worksheet you want to click. The Collapse Dialog Box icon (located to the immediate right of each text box) is helpful since clicking it shrinks the window's size. This way, you'll have an easier time selecting cells from your worksheet. To return the window to normal, click the Expand Dialog Box icon, which is to the right of the text box.

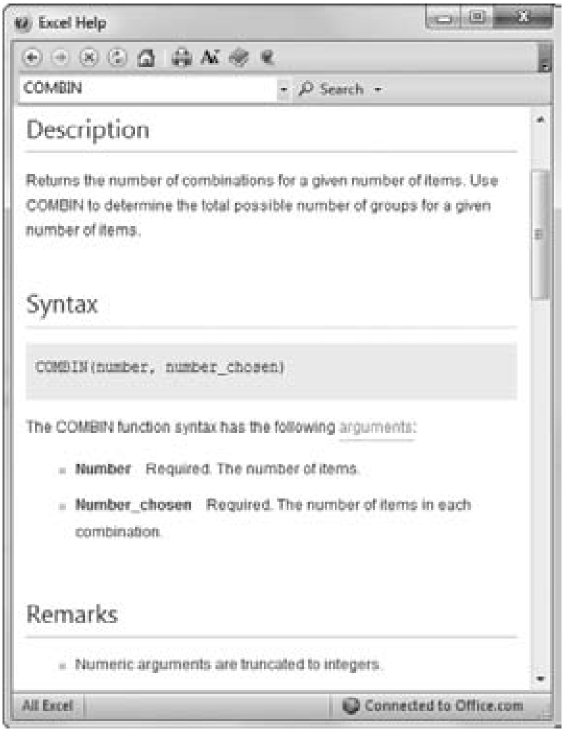

Figure 8-11:
Both the Insert Function and Function Arguments dialog boxes make it easy to get detailed reference information about any function by clicking the "Help on this function" link at the bottom left of the window. The help page shown here shows the reference for the trigonometric SIN() function, which calculates the sine of an angle. Excel's help page includes a brief description, important notes, and a couple of sample formulas that use the function, complete with results.

3. **Repeat step 2 for each argument in the function.**

   As you enter the arguments, Excel updates the formula automatically.

4. **Once you've specified a value for every required argument, click OK.**

   Excel closes the window and returns you to your worksheet.

Copying Formulas

Sometimes you need to perform similar calculations in different cells throughout a worksheet. For example, say you want to calculate sales tax on each item in a product catalog, the monthly sales in each store of a company, or the final grade for each student in a class. In this section, you'll learn how Excel makes it easy with *relative cell references*. Relative cell references are cell references that Excel updates automatically when you copy them from one cell into another. They're the standard kind of references that Excel uses (as opposed to absolute cell references, which are covered in the next section). In fact, all the references you've used so far have been relative references, but you haven't yet seen how they work with copy-and-paste operations.

POWER USERS' CLINIC

Functions That Return Arrays

A few exotic functions actually give you *multiple* results, so rather than simply generating a single value, as a function like SUM() would, these functions generate more than one value. To use these functions properly, you need to create a special type of formula called an *array formula*. (An array is a group of numbers.) You can create an array formula quite easily. Simply select all the cells that Excel will use to display the results, type in the formula, and then end by pressing Ctrl+Shift+Enter. It's this final keystroke that actually creates the array formula and links the cells together. Usually, you'll you see a different result appear in each cell.

Technically, when you create an array formula, you have a single formula that's shared between several cells. You know that you have an array formula because the formula appears in the formula bar with curly braces { } around it whenever you move into any one of the linked cells. However, you can't edit the individual cells in the array formula. (If you try, Excel warns you that you can't change part of an array.) To make a change to an array formula, you need to select all the cells in the array, edit the formula, and then press Ctrl+Shift+Enter to recreate the array formula.

Some functions that require array formulas include FREQUENCY() and TREND(), which are shown on pages 277 and 289, respectively; and TRANSPOSE(), discussed on page 353.

Consider the worksheet shown in Figure 8-12, which contains a teacher's grade book. In this example, each student has three grades: two tests and one assignment. A student's final grade is based on the following percentages: 25 percent for each of the two tests, and 50 percent for the assignment.

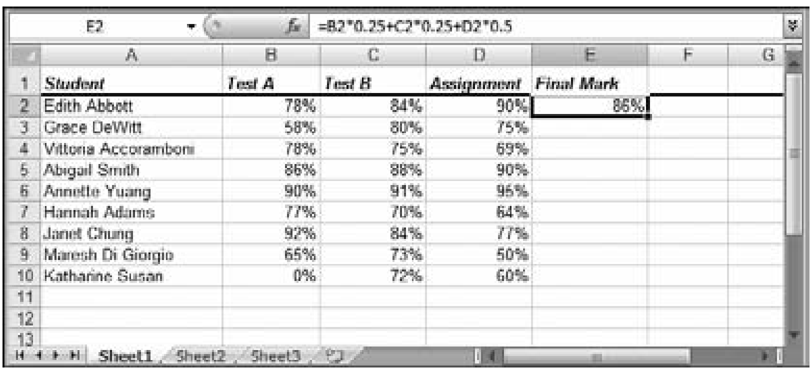

Figure 8-12: *This worksheet shows a list of students in a class, and calculates the final grade for each student using two test scores and an assignment score. So far, the only formula that's been added is for the first student (in cell E2).*

The following formula calculates the final grade for the first student (Edith Abbott):

```
=B2*25% + C2*25% + D2*50%
```

The formula that calculates the final mark for the second student (Grace DeWitt) is almost identical. The only change is that all the cell references are offset by one row, so that B2 becomes B3, C2 becomes C3, and D2 becomes D3:

```
=B3*25% + C3*25% + D3*50%
```

You may get fed-up entering all these formulas by hand. A far easier approach is to copy the formula from one cell to another. Here's how:

1. **Move to the cell containing the formula you want to copy.**

   In this example, you'd move to cell E2.

2. **Copy the formula to the clipboard by pressing Ctrl+C.**

   You can also copy the formula by choosing Home→Clipboard→Copy.

3. **Select the range of cells you want to copy the formula into.**

   Select cells E3 to E10.

4. **Paste in the new formulas by pressing Ctrl+V.**

   You can also paste the formula by choosing Home→Clipboard→Paste.

   When you paste a formula, Excel magically copies an appropriate version of the formula into each of the cells from E3 to E10. These automatic formula adjustments happen for any formula, whether it uses functions or just simple operators. Excel then automatically calculates and displays the results, as shown in Figure 8-13.

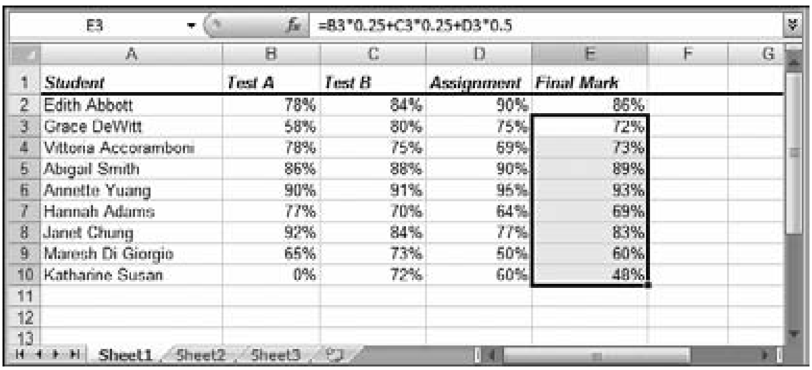

	A	B	C	D	E	F	G
1	*Student*	*Test A*	*Test B*	*Assignment*	*Final Mark*		
2	Edith Abbott	78%	84%	90%	86%		
3	Grace DeWitt	58%	80%	75%	72%		
4	Vittoria Accoramboni	78%	75%	69%	73%		
5	Abigail Smith	86%	88%	90%	89%		
6	Annette Yuang	90%	91%	95%	93%		
7	Hannah Adams	77%	70%	64%	69%		
8	Janet Chung	92%	84%	77%	83%		
9	Maresh Di Giorgio	65%	73%	50%	60%		
10	Katharine Susan	0%	72%	60%	48%		
11							
12							
13							

Figure 8-13: *When you paste the formula into one or more new cells, each Final Grade formula operates on the data in its own row. This means that you don't must tweak the formula for each student. The formula bar shows the formula contained in cell E3.*

Tip: There's an even quicker way to copy a formula to multiple cells by using the AutoFill feature introduced in Chapter 2 (page 75). In the student grade example, you start by moving to cell E2, which contains the original formula. Click and hold the small square at the bottom-right corner of the cell outline, and then drag the outline down until it covers all cells from E3 to E10. When you release the mouse button, Excel inserts the formula copies in the AutoFill region.

Absolute Cell References

Relative references are a true convenience since they let you create formula copies that don't need the slightest bit of editing. But you've probably already realized that relative references don't always work. For example, what if you have a value in a specific cell that you want to use in multiple calculations? You may have a currency conversion ratio that you want to use in a list of expenses. Each item in the list needs to use the same cell to perform the conversion correctly. But if you make copies of the formula using relative cell references, then you'll find that Excel adjusts this reference automatically and the formula ends up referring to the wrong cell (and therefore the wrong conversion value).

Figure 8-14 shows the problem with the worksheet of student grades. In this example, the test and assignment scores aren't all graded out of 100 possible points; each item has a different total score available (listed in row 12). In order to calculate the percentage a student earned on a test, you need to divide the test score by the total score available. This formula, for example, calculates the percentage for Edith Abbott's performance on Test B:

```
=B2/B12*100%
```

POWER USERS' CLINIC

Cell Formulas "Under the Hood"

To understand how Excel adjusts copied formulas, you need to know a little more about how Excel stores formulas. Internally, Excel formulas are actually stored using a system called R1C1 reference (page 27). With R1C1 referencing, when you create a formula, it doesn't contain cell references; instead, it contains cell *offsets*. Offsets tell Excel how to find a cell based on its position relative to the current cell, as you'll see below.

For example, in the student grade calculation shown in Figure 8-11, the formula:

```
=B2*25% + C2*25% + D2*50%
```

looks like this in R1C1 representation:

```
=RC[-3]*25% + RC[-2]*25% + RC[-1]*50%
```

Notice that Excel has translated the cell reference B2 into the notation RC[-3], which means "get the number in the same row, but three columns to the left." If the formula were using a number from the row above, then you'd see an R1C1 reference like R[-1]C[-3], which would tell Excel to go one row up and three columns to the left. Negative numbers in relative cell referencing indicate movement to the left (for columns) or up (for rows); positive numbers indicate movement to the right (columns) or down (rows).

When you copy a formula from one cell to another, Excel actually copies the R1C1 formula, rather than the formula you've entered (unless you've instructed Excel to use absolute cell references, as explained on page 254). To view your formulas in R1C1 style, you can temporarily change the type of cell addressing used in your spreadsheet. (See page 27 for instructions.) You'll probably want to turn this mode off when you're about to create a formula, as it's almost always easier to write formulas using Excel's standard cell addressing.

To calculate Edith's final grade for the class, you use the following formula:

```
=B2/B12*25% + C2/C12*25% + D2/D12*50%
```

Like many formulas, this one contains a mix of cells that should be relative (the individual scores in cells B2, C2, and D2) and those that should be absolute (the possible totals in cell B12, C12, and D12). As you copy this formula to subsequent rows, Excel incorrectly changes all the cell references, causing a calculation error.

Fortunately, Excel provides a perfect solution. It lets you use *absolute cell references*—cell references that always refer to the same cell. When you create a copy of a formula that contains an absolute cell reference, Excel doesn't change the reference (as it does when you use *relative* cell references; see the previous section). To indicate that a cell reference is absolute, use the dollar sign ($) character. For example, to change B12 into an absolute reference, you add the $ character twice, once in front of the column and once in front of the row, which changes it to $B$12.

Here's the corrected class grade formula (for Edith) using absolute cell references:

```
=B2/$B$12*25% + C2/$C$12*25% + D2/$D$12*50%
```

This formula still produces the same result for the first student. However, you can now copy it correctly for use with the other students. To copy this formula into all the cells in column E, use the same procedure described in the previous section on relative cell references.

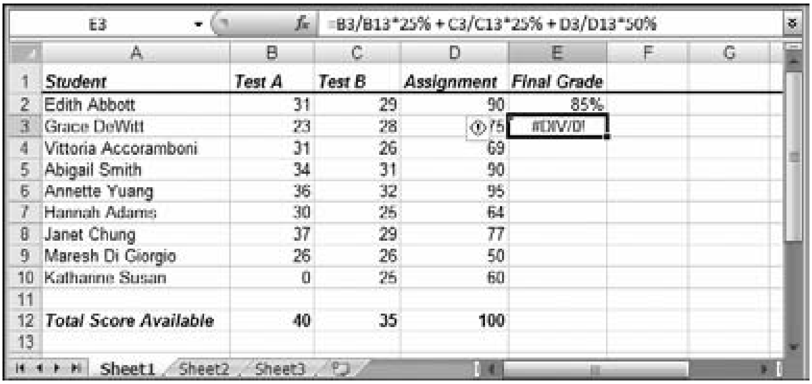
E3 =B3/B13*25% + C3/C13*25% + D3/D13*50%

	A	B	C	D	E	F	G
1	*Student*	*Test A*	*Test B*	*Assignment*	*Final Grade*		
2	Edith Abbott	31	29	90	85%		
3	Grace DeWitt	23	28	75	#DIV/0!		
4	Vittoria Accoramboni	31	26	69			
5	Abigail Smith	34	31	90			
6	Annette Yuang	36	32	95			
7	Hannah Adams	30	25	64			
8	Janet Chung	37	29	77			
9	Maresh Di Giorgio	26	26	50			
10	Katharine Susan	0	25	60			
11							
12	*Total Score Available*	40	35	100			
13							

Figure 8-14: *In this version of the student grade book, both the tests and the assignment are graded on different scales (as listed in row 12). Thus, the formula for calculating the final class grade uses the values in cells B12, C12, and D12. When you copy the Final Grade formula from the first row (cell E2) to the rows below it, Excel offsets the formula to use B13, C13, and D13—none of which exist. Thus a problem happens—shown here as a divide-by-zero error. To fix this, you need to use absolute cell references.*

Partially Fixed References

You might wonder why you need to use the $ character twice in an absolute reference (before the column letter *and* the row number). The reason is that Excel lets you create *partially* fixed references. To understand partially fixed references, it helps to remember that every cell reference consists of a column letter and a row number. With a partial fixed reference, Excel updates one component (say, the column part) but not the other (the row) when you copy the formula. If this sounds complex (or a little bizarre), consider a few examples:

- You have a loan rate in cell A1, and you want all loans on an entire worksheet to use that rate in calculations. If you refer to the cell as $A$1, its column and row always stay the same when you copy the formula to another cell.
- You have several rows of loan information. The first column of a row always contains the loan rate for all loans on that row. In your formula cell, if you refer to cell $A1, then when you copy the formula across columns and rows, the row changes (2, 3, 4, etc.) but the column doesn't (A2, A3, A4, etc.).
- You have a table of loan rates organized by the length of the loan (10-year, 15-year, 20-year, etc.) along the top of a worksheet. Loans in each column are calculated using the rate specified at the top of that column. If you refer to the rate cell as A$1 in your first column's formula, then the row stays constant (1), but the column changes (B1, C1, D1, etc.) as you copy the formula across columns and down rows.

Tip: You can quickly change formula references into absolute or partially fixed references. Just put the cell into edit mode (by double-clicking it or pressing F2). Then, move through the formula until you've highlighted the appropriate cell reference. Now, press F4 to change the cell reference. Each time you press F4, the reference changes. If the reference is A1, for instance, it becomes $A$1, then A$1, then $A1, and then A1 again.

UP TO SPEED

Creating an Exact Formula Copy

There's another way to copy a formula that prevents Excel from automatically changing the formula's cell references. The trick is to copy the formula itself rather than copy the whole cell (which is what you do when performing a basic copy-and-paste operation on a formula).

The process takes a few more steps, and it lets you paste only one copy at a time, but it can still come in handy if you don't want Excel to use relative references. Here's how it works:

1. First, move to the cell that contains the formula you want to copy.
2. Place this cell in edit mode by double-clicking it or pressing F2.
3. Select all the text in the cell. You can use the mouse, or you can use the arrow keys (just hold down Shift as you scroll from the beginning to the end of the cell).
4. Once you've selected the complete formula, press Ctrl+C to copy it.
5. Press Esc to leave edit mode once you're finished.
6. Move to the new cell, and then press Ctrl+V to paste it.

Keep in mind that when you use this approach, you create an exact copy of the formula. That means this technique doesn't help in situations where some cell references need to be absolute, and others need to be relative.

Referring to Other Worksheets

To reference a cell in another worksheet, you simply need to preface the cell with the worksheet name, followed by an exclamation mark. For example, say you've created a formula to double the value of the number in cell A1 on a worksheet named Sheet1. You'd use this formula:

```
=A1*2
```

If you want to use the same formula in another worksheet (in the same workbook), you'd insert this formula in the new worksheet:

```
=Sheet1!A1*2
```

Note: If you use the point-and-click method to build formulas, you'll find that you don't need to worry about the syntax for referring to cells on other worksheets. If you switch to another worksheet, and then click a cell in it, while you're building a formula, Excel automatically inserts the correct reference with the worksheet name.

Referring to Other Workbooks

It's fairly common for one worksheet to reference another within the same workbook file. (For a refresher on the difference between worksheets and workbooks see page 118.) On page 359, you'll see an example where one worksheet includes a product catalog, and a second worksheet builds an invoice. In that example, the second worksheet uses formulas that refer to data on the first sheet.

It's less common for a worksheet to refer to data in another *file*. The potential problem with this type of link is that there's no way to guarantee that referenced files will always be available. If you rename the file or move it to a different folder, then the link breaks. Fortunately, Excel doesn't leave you completely stranded in this situation—instead, it continues with the most recent version of the data it was able to retrieve. You can also tweak the link to point to the new location of the file.

To create a link to a cell in another workbook, you need to put the file name at the beginning of the reference, and then enclose it in square brackets. You follow the file name with the sheet name, an exclamation mark, and the cell address. Here's an example:

```
=[SimpleExpenses.xlsx]Sheet1!B3
```

If the file name or the sheet name contains any spaces, you need to enclose the whole initial portion inside apostrophes, like so:

```
='[Simple Expenses.xlsx]Sheet1'!B3
```

When you enter this formula, Excel checks to see if this file is already open in another Excel window. If not, then Excel attempts to find the file in the current folder (that is, the same folder that holds the workbook you're editing). If it can't find the file, it shows a standard Open dialog box. You can use this dialog box to browse to the file you want to use. Once you select the file, Excel updates the formula accordingly.

Note: Once again, if you use point-and-click formula creation (for example, start entering a formula in one workbook, and then click a cell in another workbook), Excel will add the right cell reference, complete with the appropriate file name.

Excel performs a little sleight of hand with formulas that reference other files. If the referenced workbook's currently open, the formula displays only the file name. However, once you close the referenced file, the formula changes so that it includes the full folder path. Excel makes this change automatically, whether you type a file name in by hand or build the formula by pointing and clicking cells in another open workbook.

Here's an example of what the full formula looks like when you close the file that contains the Sheet1 worksheet:

```
='[C:\Users\Matthew\Documents\Simple Expenses.xlsx]Sheet1'!B3
```

Updating links

You use a formula that *refers* to another workbook (rather than just copying the information you need) when you know the information in that workbook might change. So what does Excel do to keep your linked information up to date? It all depends.

If the linked workbook is currently open, Excel spots any changes right away. For example, imagine you have a workbook named Summary.xlsx that refers to cell B3 in a workbook named SimpleExpenses.xlsx (as shown in the previous example). When you modify the number in B3 in SimpleExpenses.xlsx, Excel recalculates the linked formula in Summary.xlsx immediately.

It's a different story if you work with both files *separately*. For example, imagine you close Summary.xlsx, change the total in SimpleExpenses.xlsx, and then close SimpleExpenses.xlsx. The next time you open Summary.xlsx, Excel is in a difficult position. It needs to get the updated information in SimpleExpenses.xlsx, but it doesn't want to open the file to do so (at least not without your OK). So instead, the Summary.xlsx file uses the most recent information—the value it pulled out of cell B3 the last time SimpleExpenses.xlsx was open.

Excel also notices that your workbook uses linked formulas. The first time you open a workbook with linked formulas, Excel shows a message bar just above the grid of cells, warning you that it hasn't retrieved the latest information (Figure 8-15).

If you want to update your linked formulas, click the Enable Content button in the message bar. Excel opens the linked workbook (in this case, SimpleExpenses.xlsx) behind the scenes, and then uses the latest information to recalculate your formula.

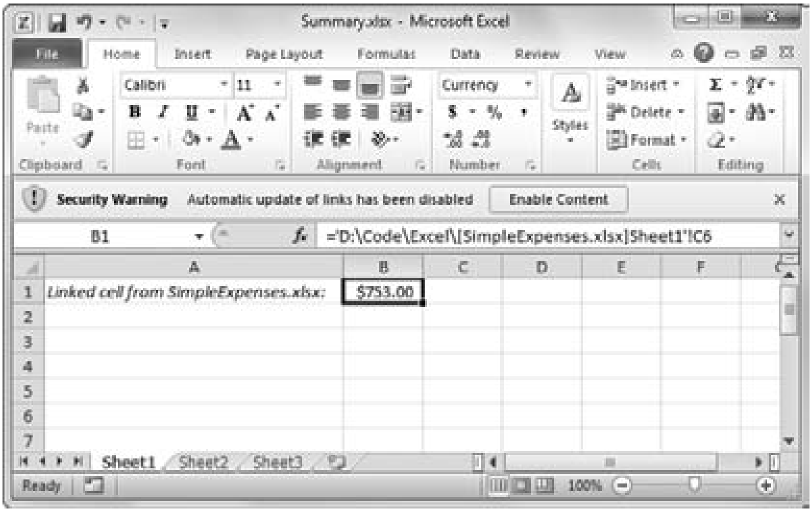

Figure 8-15: *The security warning tells you that the workbook you've opened has linked cells, but Excel hasn't tried to update them. That's because updating them involves opening another file, which ultra-paranoid Excel isn't ready to do.*

In Excel 2007, you'd see this security warning every time you opened a workbook with linked formulas. But Excel 2010 introduces a feature called *trusted documents* (which is explained in detail on page 804). There bottom line is this: once you click Enable Content, your workbook is trusted, and you won't see the warning message the next time you open it. That's because Excel stores a list of trusted workbooks in your Windows user settings, and it checks this list whenever you open a workbook. (The only exception to this system is if you move your workbook, rename it, or log on as a different user. In any of these cases, Excel won't trust the workbook anymore, you'll see the warning message the next time you open it, and need to click Enable Content again.)

Now, here's the tricky bit. Even if Excel trusts your workbook, it won't automatically update your linked formulas when you open the workbook, because it doesn't know if that's what you want to do. Instead, it shows you a dialog box that asks if you want to update the links (Figure 8-16).

If you don't want to update links right away (and you click Don't Update), you can still update your links later. The first approach is to open the workbook you're linking *to*, while the other workbook is still open. In this case Excel updates the links automatically. The second approach is to use the Edit Links dialog box, described on page 261.

Note: Confused? Here's the executive summary. The first time you open a workbook with linked formulas, you see the security message (Figure 8-16) and you must click Enable Content to update the links. Any time you open the workbook afterwards, you see the link message (Figure 8-17) and you must click Update to update the links. Keep reading to learn how you can tweak this behavior.

Figure 8-16:
If you click Update, Excel silently opens the linked workbook, updates the links, and then closes the workbook. If you click Don't Update, Excel does nothing, but you can update them later.

Troubleshooting links

Sometimes, despite the best of intentions, your formula points to a workbook that Excel can't find. For example, if you rename SimpleExpenses.xlsx or move it to another folder, then you'll run into this problem. In this situation, Excel gives you an error message when you attempt to update your links (Figure 8-17).

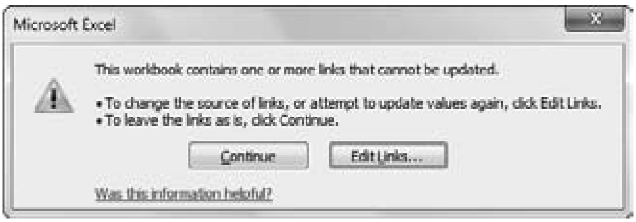

Figure 8-17:
If Excel can't find the workbook you referenced, you see an error message that gives you the choice of ignoring it and working with what you have (click Continue), or relinking the workbook to the correct file (click Edit Links and see Figure 8-18).

If you know that a link's gone bad, you've got a problem—depending on how many linked formulas your workbook has, you could be stuck updating dozens of cells. Fortunately, the Edit Links dialog box (shown in Figure 8-18) gives you an easier alternative. Rather than changing each formula by hand, you can "relink" everything by pointing Excel to the right file.

The Edit Links dialog box

The Edit Links dialog box is invaluable for fixing broken links. In addition, some of the buttons in the Edit Links dialog box work with links that aren't broken. They include:

- **Update Values** refreshes the formulas that use linked cells. This is useful if you chose not to update your links when you first opened the workbook file.
- **Open Source** opens the linked workbook in another Excel window, so you can really see what's going on. This also has the side effect of updating all the linked formulas in your current workbook.

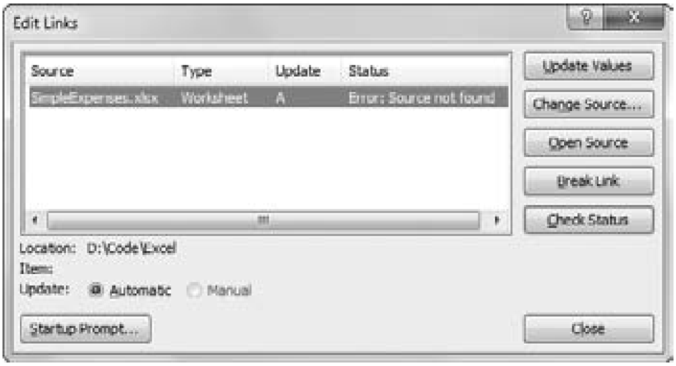

Figure 8-18:
In this example, Excel can't find the linked file SimpleExpenses.xlsx. You can select the link, and then click Change Source to browse for the linked workbook (click OK when you find it); or click Break Link to make your workbook independent, and just use whatever data you extracted last time permanently.

- **Break Link** severs your ties with the linked workbook. Excel replaces any linked formulas with the most recent data.
- **Check Status** rechecks the file, and updates the Status column to let you know whether it exists, whether it's currently open, and so on.
- **Startup Prompt** opens a Startup Prompt dialog box where you can choose the behavior Excel uses when it opens this workbook. You have three choices: refresh the linked cells automatically, never refresh, or ask you (the standard setting). This is a useful option is you're tired of being harassed by the dialog box shown in Figure 8-17 every time you open the workbook.

FREQUENTLY ASKED QUESTION

How Changing the Location of Cells Affects Formulas

OK, I know how Excel adjusts a formula when I copy it to another location. But what happens if I move cells around after I've created a formula?

No worries. It turns out that Excel is surprisingly intelligent. Consider the following simple formula:

```
=B1+A2
```

If you cut and paste the contents of A2 to A3, Excel automatically updates your formula to point to the new cell, without complaining once. It also performs the same automatic cleanup if you drag the contents of a cell to another location (although if you simply make a duplicate copy of the cell, Excel won't change your formula). Excel is also on the ball when you insert and delete rows and columns.

If at any time Excel can't find your cells, the formula changes to show the error code #REF!. You can then take a closer look at the formula to find out what really went wrong. For example, if you delete column B from your spreadsheet (by selecting the column and using the Home→Cells→Delete command), the formula changes to this:

```
=#REF!+A2
```

Even though there's still a B1 cell in your worksheet (it's the cell that was formerly named C1), Excel modifies the formula to make it clear that you've lost your original data.

CHAPTER

9

Math and Statistical Functions

Excel is packed with dozens of mathematical functions. Some of these functions are for specialized audiences, like engineers or statisticians, while others are so useful they can turn up in almost any civilian's spreadsheet.

In this chapter, you'll take a tour of two of Excel's largest function categories: Math & Trig and Statistical. Rather than slog through each function one by one, this chapter covers the most useful functions in each category. It starts by looking at a bunch of functions that help round, add, and count numbers. Then it also explains how to find averages, medians, and percentiles, which are helpful when comparing groups of data. Toward the end of the chapter, you'll see some of the more specialized functions that showcase Excel's skill with trigonometry and advanced statistics—great for math lovers, accounting jockeys, or political-polling wonks interested in double-checking statistical significance claims.

Rounding Numbers

Most people don't devote enough thought to *rounding*, the process by which you adjust fractional numbers so they're less precise but more manageable. For example, rounding can transform the unwieldy number 1.984323125 to 2. Excel has two ways to use rounding:

- **Modify the number format of the cell**. With this method, Excel rounds the displayed value, but it doesn't change the underlying value. The advantage to this approach is that you can use the value in other calculations without losing any precision. When Excel rounds your numbers using this method, it simply rounds to the last displayed digit (rounding up if the next digit is 5 or greater).

For example, if you tell Excel to show the number 3.145 using two decimal places, then Excel displays the rounded value of 3.15. (Cell value formatting is described in Chapter 5.)

- **Use a rounding function**. This approach gives you more control. For example, you can round a number *before* you use it in another calculation, and you can round numbers to a multiple you choose, like 500 or 1,000. The drawback is that when you use a rounding function, you may lose precision. This doesn't happen when you change the number format, which simply tweaks the way the number is displayed.

With classic overkill, Microsoft includes no fewer than 10 functions designed specifically for rounding numbers, from the basic ROUND(), to the more flexible MROUND(), and to the quirky EVEN() and ODD().

ROUND(), ROUNDDOWN(), ROUNDUP(): Rounding Numbers

The most basic (and most commonly used) of Excel's rounding functions is ROUND(), which rounds a numeric value to whatever level of precision you choose. The ROUND() function needs two arguments: the actual number you want to round, and the number of digits you want to keep to the right of the decimal point. Here's what it looks like:

```
ROUND(number_to_round, number_of_digits)
```

For example, the following formula rounds the number 3.987 to two decimal places. The result is 3.99.

```
=ROUND(3.987, 2)
```

If you specify 0 for the number of digits, then Excel rounds to the nearest whole number. Interestingly, you can also round to the nearest 10, 100, 1000, and so on by using negative numbers for the second argument. For example, if you use –2 for the number of digits, then Excel rounds two digits to the *left* of the decimal place, which means your number gets rounded to the nearest 100.

Here's an example:

```
=ROUND(34655.7, -2)
```

This formula produces a result of 34,700.

Note: The ROUND() function always rounds the positive values 1 through 4 *down*; 5 through 9 get rounded *up*. If you round 1.5 to the nearest whole number, for instance, the result is 2. When dealing with negative numbers, Excel rounds the digits 5 through 9 down (toward the larger negative value). Similarly, –1 through –4 get rounded up. For example, –1.5 gets rounded to –2, while –1.4 gets rounded up to –1.

The ROUNDDOWN() and ROUNDUP() functions work similarly to ROUND(). Like ROUND(), they take two arguments: the number you want to round, and the number of decimal places you want the final, rounded number to use. The difference is that ROUNDDOWN() always rounds numbers down, while ROUNDUP() always rounds numbers up.

For example, the result of ROUNDUP(1.1, 0) is 2, even though 1.1 is only slightly above 1. Similarly, the result of ROUNDDOWN(1.9, 0) is 1, even though 1.9 is almost 2. The only time that ROUNDUP() and ROUNDDOWN() don't change a number is if it's already rounded to the appropriate precision. For example, the result of ROUNDUP(2, 0) and ROUNDDOWN(2, 0) is the same: 2.

When most people learn about ROUNDUP() and ROUNDDOWN(), they often wonder why anyone would want to use a rounding function that's *less* precise than ol' reliable ROUND(). The answer, not surprisingly, must do with making more money. If you're selling discounted Beanie Babies, say, then you might set the price at 60 percent of the original list price. However, this formula produces prices like $8.43423411 that you need to round to the nearest penny. Rather than rounding down (and giving up your fractional cents), you can use ROUNDUP() to make sure the price is always rounded up, ensuring that you keep every last penny and even collect a few extra. Never say Microsoft didn't try to help you make a buck.

MROUND(), CEILING(), and FLOOR(): More Rounding Functions

While you may not necessarily think of the ROUND() function as letting you round a number to any multiple of a *power of 10* you'd like, that's actually what's going on under the hood. But in some cases, you might want to round your numbers to something that *isn't* a power of 10. For example, you might want to round a number to the nearest multiple of 5 or 50 or 100.

In these cases, you need to use Excel's MROUND() function. MROUND() takes two arguments: the number you want to round, and the multiple you want to use:

```
MROUND(number_to_round, multiple)
```

Here it is in action, as the following formula rounds the number 653 to the nearest multiple of 5. The result's 655.

```
=MROUND(653, 5)
```

CEILING() and FLOOR() work like MROUND(), except they let you control whether the rounding goes up (in which case you'd use CEILING()) or down (FLOOR()). In other words, the CEILING() function's really a combination of MROUND() and ROUNDUP(). The FLOOR() function, meanwhile, is a combination of MROUND() and ROUNDDOWN().

Consider the following formula. It rounds the number 653 to the nearest multiple of 5, but, because it uses the FLOOR() function, it rounds *downward*. The result is 650:

```
=FLOOR(653, 5)
```

Note: The rounding functions don't always treat negative numbers the way you'd expect. ROUNDUP() and ROUNDDOWN() ignore the minus sign entirely. If you use ROUNDUP() to round to no decimal places with –4.5, then you'll end up with –5. The result is a larger negative number. MROUND() exhibits the same behavior (Excel treats positive and negative numbers identically), with one catch. When rounding negative numbers with MROUND(), you must supply a negative number for the multiple argument.

Compare that with the following formula, which produces a result of 700. This formula uses the CEILING function, which means it always rounds up, unless the number is already rounded to the appropriate multiple, of course.

```
=CEILING(652, 50)
```

The CEILING() function is handy if you need to calculate an item's price in more than one currency. Consider a U.S.-based company that's trying to figure out how to price its products in Canada (Figure 9-1). The company would probably start by setting the price in U.S. dollars, and then use a currency conversion ratio to produce a set of Canadian prices. However, this approach produces prices that, in effect, leave a little money on the table. (Most people willing to pay $153.16 for a product are probably also willing to pay $153.99—whether it's Canadian or American dollars.) Using a function like MROUND(), the company could make sure that prices are always rounded to the nearest 25 or 50 cent multiple. That's fine, but there's a possibility that margins might be shaved too thin if you round prices down—so the CEILING() function is a better choice. And if you want to get particularly fancy, why not use CEILING() to round up to the nearest dollar and then subtract a single cent, giving attention-grabbing prices that always end with .99?

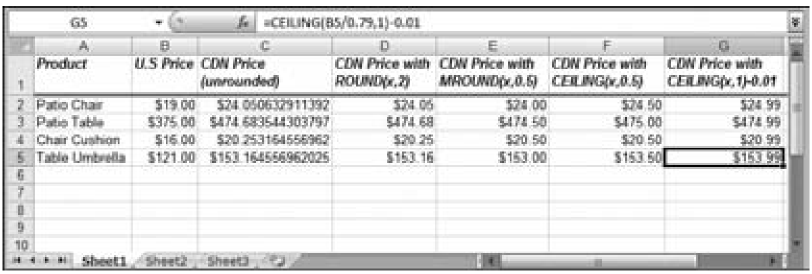
G5 =CEILING(B5/0.79,1)-0.01

	A	B	C	D	E	F	G
1	Product	U.S Price	CDN Price (unrounded)	CDN Price with ROUND(x,2)	CDN Price with MROUND(x,0.5)	CDN Price with CEILING(x,0.5)	CDN Price with CEILING(x,1)-0.01
2	Patio Chair	$19.00	$24.050632911392	$24.05	$24.00	$24.50	$24.99
3	Patio Table	$375.00	$474.683544303797	$474.68	$474.50	$475.00	$474.99
4	Chair Cushion	$16.00	$20.253164556962	$20.25	$20.50	$20.50	$20.99
5	Table Umbrella	$121.00	$153.164556962025	$153.16	$153.00	$153.50	$153.99

Sheet1 Sheet2 Sheet3

Figure 9-1: *The CEILING() function in action. Notice how different approaches to rounding net different profits.*

INT() and TRUNC(): Chopping Off Non-Whole Numbers

INT() and TRUNC() are two functions that simply shorten your numbers, removing extra digits without performing any rounding. INT() is the simpler of the two, as it always rounds to whole numbers. You need only specify the number you want rounded. The formula here, for example, works out to 2:

```
=INT(2.5)
```

You'd get the same answer for INT(2), INT(2.9), INT(2.7509630), and so on. In each case, Excel discards the decimal portion.

TRUNC() is similar to INT(), except it uses a second argument specifying the number of decimal places you want to *preserve*. This argument is optional, and if you leave it out, TRUNC() and INT() behave exactly the same with positive numbers.

```
TRUNC(number_to_round, [number_of_digits])
```

In real life, you'll find reasons to round numbers far more often than you have reasons to truncate them.

One of the few instances in which you'd use truncation is when you want to *ignore* the digits of a number that you don't want, rather than *change* them. For example, imagine you want to start memorizing the digits of the mathematical constant pi (3.14159265…). Using the following formula, you can truncate pi to the first three decimal places so you can get started:

```
=TRUNC(PI(), 3)
```

This formula gives you the number 3.141 (note that Excel lets you use the PI() function to insert the constant pi). If you used the ROUND() function, then you'd end up with 3.142 instead, because Excel would round the last digit.

You might also use TRUNC() to give you just the whole number portion of a value. This approach is reasonable when the decimal portion doesn't have any real meaning. For example, you might calculate that your furniture factory produces 5.6 chairs per day when you average the monthly production. Unfortunately, this still means you can rely on having only five chairs to sell at the end of the day. You can make a similar case for using the ROUNDUP() function. For example, maybe you know your friends tend to eat 1.3 pumpkin pies each time they come for a dinner party, so you'd do best to round up and make sure you slide two pies into the oven.

POWER USERS' CLINIC

Excel's Rounding May Not Be Good Enough

Even though Excel provides 10 rounding functions, they all work in exactly the same way. They use a technique called *arithmetic rounding*, which always rounds the number 5 up. Although this may seem like the only option, accountants and statisticians are well aware that a variety of different approaches to rounding numbers exist, and that they all have their own particular problems.

The key difficulty is deciding how to round the number 5. For example, consider 1.5, which lies exactly halfway between the numbers 1 and 2. In arithmetic rounding, 1.5 is always rounded up to 2. This convention can bias your results if you're rounding a large group of numbers and then adding them together (since the five digits 5, 6, 7, 8, and 9 always get rounded up, versus the four digits 1, 2, 3, and 4, which always get rounded down).

The best way to understand the problem is to think of the sales tax you pay on a typical restaurant bill. Sales tax is supposed to be calculated as a percentage of the total, and then rounded arithmetically. However, imagine what would happen if the sales tax were calculated separately on each item you ordered, *rounded up* separately, and then added. You'd pay more than if the tax were calculated from the total.

So how can you avoid rounding biases? The best option is to calculate first, and round later. Never add numbers that you've already rounded. You could also use a different type of rounding. One example is *banker's rounding*, which rounds 5 up sometimes and down other times, depending on whether it's paired with an even or odd number. For example, 1.5 is rounded up to 2, but 2.5 is rounded down to 2, 3.5 is rounded up to 4, and so on. You can use banker's rounding in Excel, but Microsoft makes it available only to the VBA (Visual Basic for Applications) programming language. Translation: You need to write a custom macro or function to access it.

You'll learn about how you can use banker's rounding with VBA on page 841. In the meantime, you can learn more about the various aspects of rounding, and see what types of rounding Microsoft uses in different products in the technical support note at *http://tinyurl.com/d6taz.*

EVEN() and ODD(): Rounding Up to Even or Odd Values

EVEN() and ODD() are the last of Excel's rounding functions. These functions accept a single number and round it up to the nearest even or odd number. As simple as this sounds, these functions can cause a bit of confusion because many people assume that the functions will return the *closest* odd or even number. But they don't. Since the functions round *up*, the closest correct number may be, numerically speaking, a confusingly long ways away.

To understand these quirky functions a little better, consider the following formula:

```
=ODD(2.6)
```

This formula produces the expected result: It rounds 2.6 up to the closest odd number, 3. Now consider:

```
=ODD(3.6)
```

This formula also rounds up to the nearest odd number, which in this case is 5.

In fact, ODD() always rounds up, unless you begin with a whole odd number. That means that the result of the following formula is also 5, even though 3.1 is clearly much closer to 3 than 5:

```
=ODD(3.1)
```

The EVEN() function behaves similarly. Thus, the result of the following formula is 4:

```
=EVEN(2.1)
```

The EVEN() and ODD() functions aren't useful too often. For most people, they simply represent an interesting footnote in Excel functions.

Groups of Numbers

You can use spreadsheets to distil a few important pieces of information out of several pages of data. For example, say you want to hunt through a column looking for minimums and maximums, in order to find the lowest-priced product or best sales quarter. Or maybe you want to calculate averages, means, and percentile rankings to help grade a class of students. In either case, Excel provides a number of useful functions. Most of these are part of the Statistical group, although the SUM() function is actually part of the Math & Trig group.

SUM(): Summing Up Numbers

Almost every Excel program in existence has been called on at least once to do the same thing: add a group of numbers. This task falls to the wildly popular SUM() function, which simply adds everything in it. The SUM() function takes over 200 arguments, each of which can be a single cell reference or a range of cells.

Here's a SUM() formula that adds two cells:

```
=SUM(A1,A2)
```

And here's a SUM() formula that adds the range of 11 cells from A2 to A12:

```
=SUM(A2:A12)
```

And here's a SUM() formula that adds a range of cells along with a separately referenced cell, and two literal values:

```
=SUM(A2:A12,B5,429.1,35000)
```

Note: The SUM() function automatically ignores any cells in its range with text content, or any blank ones. However, SUM() adds up calendar dates (which are actually specially formatted numbers, as you saw on page 70). Therefore, make sure you don't sum a range of cells that includes a date.

POWER USERS' CLINIC

Creating Formulas Quickly with AutoSum

People so often sum a row or column of values that Excel has a button dedicated to just that purpose: the Auto-Sum button. To find it, choose Formulas→Function Library→AutoSum. The button has a picture of the Greek letter Σ (capital sigma), which means sum to mathematicians.

When you click the AutoSum button, Excel makes an educated guess about what cells you want to total (as shown in Figure 9-2). For example, if you're at the end of a row, then Excel assumes you want to add all the numeric values in all the columns on the left. If you're at the bottom of a column of numbers, Excel assumes you want to add these values instead.

The AutoSum feature isn't just for summing. You can also use it to calculate averages, counts, maximums, and minimums. To do so, don't click the top part of the AutoSum button. Instead, click the drop-down arrow at the bottom of the AutoSum button to get a list of calculation options, including summing, averaging, counting, and finding the maximum or minimum. When you pick an option from the AutoSum menu, Excel inserts the appropriate function into your cell and uses the same auto-guess strategy to pick out the nearby cells for the calculation.

COUNT(), COUNTA(), and COUNTBLANK(): Counting Items in a List

Sometimes you don't need to add a series of cells, but instead want to know how many items there are (in a list, for instance). That's the purpose of Excel's straightforward counting functions, COUNT(), COUNTA(), and COUNTBLANK(). COUNT() and COUNTA() operate similarly to the SUM() function in that they accept over 200 arguments, each of which can be a cell reference or a range of cells.

The COUNT() function counts the number of cells that have numeric input (including dates). The COUNTA() function counts cells with any kind of content. And finally, the COUNTBLANK() function takes a single argument—a range of cells—and gives you the number of empty cells in that range (see Figure 9-3).

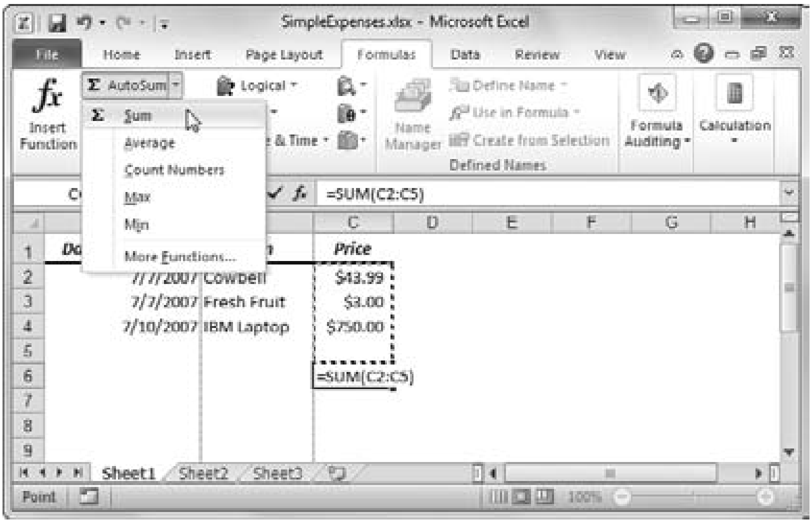

Figure 9-2: *When you click AutoSum, Excel creates a formula that uses SUM() with the cell range it thinks you need, and it highlights the reference with a marquee box (as shown here). At this point, you can hit Enter to accept the formula as is, type in a new reference, or resize the range with your mouse.*

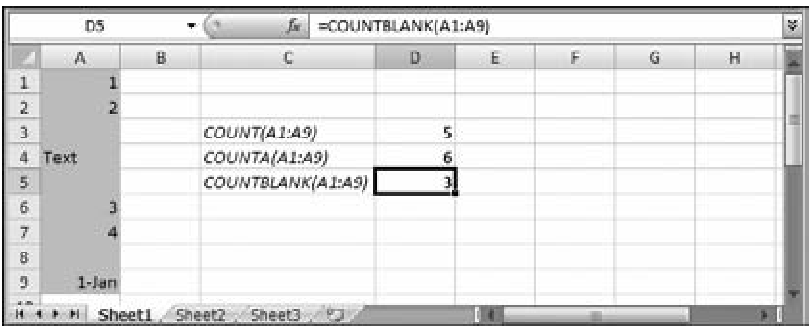

Figure 9-3: *This worksheet shows the counting functions in action on a range of cells, from A1 to A9. The function names are in column C; the formulas containing the functions are in column D. The COUNT() function includes the date cell A9, while the COUNTA() function includes all cells that have any kind of information in them.*

Here's how you could use the COUNT() function with a range of cells:

```
=COUNT(A2:A12)
```

For an example of when COUNT() comes in handy, consider the following formula, which determines the average of a group of cells, without requiring that you manually input the total number of values:

```
=SUM(A2:A12)/COUNT(A2:A12)
```

This formula finds the average of all the cells in the range A2:12 that have values. Because it uses COUNT(), the average remains correct even if the range includes empty cells (which Excel simply ignores). Of course, Excel already includes an AVERAGE() function that can perform this calculation for you, but the COUNT() function is still useful in long lists of data that may contain missing information. For example, imagine you create a worksheet with a list of customer charges. One of the columns in your list is Discount Amount. Using the COUNT() function on the discount column, you can find out how many customers got a discount in your list. When you want to get a little fancier, use COUNT() on the Customer Name column to find out how many rows are in the table. Then you can divide the number of discount sales by the number of total sales to find out the percentage of the time that you sell something for a reduced price.

You can also combine the counting functions to calculate some additional pieces of potentially useful information. If you subtract COUNT() from COUNTA(), for example (that is, if your formula reads something like *=COUNTA(A1:A10)–COUNT(A1:A10)*), you end up with the number of text cells in the range.

Tip: You can also use SUMIF() and COUNTIF() functions to sum or count cells that meet specific criteria. For more information on using conditional logic, as well as the SUMIF() and COUNTIF() functions, see page 367.

MAX() and MIN(): Finding Maximum and Minimum Values

The MAX() and MIN() functions pick the largest or smallest value out of a series of cells. This tool's great for picking important values (best-selling products, failing students, or historically high temperatures) out of a long list of information. As with COUNT() and SUM(), the MAX() and MIN() functions accept over 200 cell references or ranges.

For example, the following formula examines four cells (A2, A3, B2, and B3) and tells you the largest value:

```
=MAX(A2:B3)
```

The MAX() and MIN() functions ignore any nonnumeric content, which includes text, empty cells, and Boolean (true or false) values.

Excel includes dates in MAX() and MIN() calculations because it stores them internally as the number of days that have passed since a particular date. (For Windows versions of Excel, this date is January 1, 1900; see page 70 for details on how this works.) For this reason, it makes little sense to use MAX() or MIN() on a range that includes both dates and ordinary numeric data. On the other hand, you may want to use MAX() to find the latest (that is, the largest) date or MIN() to find the oldest (or smallest) date. Just make sure you format the cell containing the formula using the Date number format, so that Excel displays the date it has identified.

Excel also provides MAXA() and MINA() functions, which work just like MAX() and MIN(), except for the way in which they handle text and Booleans. MAXA() and MINA() always assume TRUE values equal 1 and FALSE values equal 0. Excel treats all text values as 0. For example, if you have a list of top professionals with a Salary column, and some cells say "Undisclosed," it makes sense to use MAX() and MIN() to ignore these cells altogether. On the other hand, if you have a list of items you've tried to auction off on eBay with a Sold For column, and some cells say "No Bids," you may use MINA() and MAXA() to treat this text as a 0 value.

LARGE(), SMALL(), and RANK(): Ranking Numbers

MAX() and MIN() let you grab the largest and smallest numbers, but what if you want to grab something in between? For example, you might want to create a top10 list of best-selling products, instead of picking out just the single best. In these cases, consider the LARGE() and SMALL() functions, which identify values that aren't *quite* the highest or lowest. In fact, these functions even let you specify how *far* from the top and bottom you want to look. Here's how they work.

Both the LARGE() and SMALL() functions require two arguments: the range you want to search, and the item's position in the list. The list position is where the item would fall if the list were ordered from largest to smallest (for LARGE()), or from smallest to largest (for SMALL()). Here's what LARGE() looks like:

```
LARGE(range, position)
```

For example, if you specify a position of 1 with the LARGE() function, then you get the largest item on the list, which is the same result as using MAX(). If you specify a position of 2, as in the following formula, then you get the second largest value:

```
=LARGE(A2:A12, 2)
```

Here's an example formula that adds the three largest entries in a range:

```
=LARGE(A2:A12,1) + LARGE(A2:A12,2) + LARGE(A2:A12,3)
```

Assuming the range A2:A12 contains a list of monthly expenses, this formula gives you the total of your three most extravagant splurges.

SMALL() performs the opposite task by identifying the number that's the smallest, second-smallest, and so on. For example, the following formula gives you the second-smallest number:

```
=SMALL(A2:A12, 2)
```

And finally, Excel lets you approach this problem in reverse. Using the RANK() function, you can find where a specific value falls in the list. The RANK() function requires two parts: the number you're looking for and the range you're searching. In addition, you can supply a third parameter that specifies how Excel should order the values before searching. (Excel automatically searches values from highest to lowest.) Here's what the RANK() formula looks like:

```
RANK(number, range, [order_type])
```

For example, imagine you have a range of cells from A2 to A12 that represent scores on a test. Somewhere in this range is a score of 77. You want to know how this compares to the other marks, so you create the following formula using the RANK() function:

```
=RANK(77, A2:A12)
```

If this formula works out to 5, you know that 77 is the fifth-highest score in the range you indicated. But what if more than one student scored 77? Excel handles duplicates in much the same way as it deals with tied times in races. If three students score 77, for example, then they all tie for fifth place. But the next lowest grade (say, 75) gets a rank of eight—three positions down the list.

If you want to rank values in the reverse order, then you need to include a third argument of 0. This orders them from highest to lowest (so that the lowest grade gets a rank of 1):

```
=RANK(77, A2:A12, 0)
```

Note: If you use RANK() to rank a number that doesn't exist in the data series, Excel gives you the #N/A error code in the cell, telling you it can't rank the number because the number isn't there.

You can do a neat trick and rank all your numbers in a list. For example, assume you have a list of grades in cells A2 to A12. You want to show the rank of each grade in the corresponding cells B2 to B12. Calculate the rank for A2 like so:

```
=RANK(A2, $A$2:$A$12)
```

This formula uses absolute references for the list. That means you can copy this formula to cells B3 through B12 to calculate the rank for the remaining grades, and Excel automatically adjusts the formula to what you need.

POWER USER'S CLINIC

RANK.EQ() and RANK.AVG()

Excel 2010 turns RANK() into a compatibility function (page 239) and adds two new versions, RANK.EQ() and RANK.AVG(), which differ in how they handle tied values.

RANK.EQ() works in the same way as RANK(): it assigns the next available rank to all the tied values. So if two students tie for second place, they both get a rank of 2 and the next student is ranked 4. No one gets the third-place ranking.

RANK.AVG() calculates an *average* rank for tied values. So if two students tie for second place, they both get a rank of 2.5 (because that's the average of rank 2 and the skipped rank 3). If three students tied for second place, the average rank is 3, because that's the result of the average rank calculation *(2+3+4)/3.*

The average ranking system has the advantage of keeping tied values from rating quite as highly as standalone performers. This lets you avoid the illusion that a large number of people are better than average because they've all tied for a high rank.

Just remember, if you use either RANK.EQ() or RANK.AVG(), your calculations won't work on older versions of Excel.

AVERAGE() and MEDIAN(): Finding Average or Median Values

Excel makes it easy for you to find the average value for a set of numbers. The AVERAGE() function doesn't accomplish anything you couldn't do on your own using COUNT() and SUM() together, but, well, you could also bake your own bread. Bottom line: The AVERAGE() function is a real timesaver.

The AVERAGE() function uses just one argument—the cell range you want to average:

```
=AVERAGE(A2:A12)
```

The AVERAGE() function ignores all empty cells or text values. For example, if the preceding formula had only three numbers in the range A2:A12, then Excel would add these values and divide them by 3. If this isn't the behavior you want, you can use the AVERAGEA() function, which counts all text cells as though they contain the number 0 (but it still ignores blank cells). (For an example of when text should and shouldn't be treated as 0 values, see the discussion about the MAXA() and MINA() functions on page 272.)

Tip: In some cases, you may want to perform an average that *ignores* certain values. For example, when determining the average score of all students on a test, you may want to disregard students who scored 0 because they may have been absent. In this case, you can use the conditional averaging technique described on page 367.

Excel can also help you identify the *median* value for a set of numbers. If you were to order your range of numbers from lowest to highest, then the median would be the number that falls in the middle. (If you have an even amount of numbers, Excel averages the two middle numbers to generate the median.)

You calculate the median (see Figure 9-4) in the same way that you calculate an average:

```
=MEDIAN(A2:A12)
```

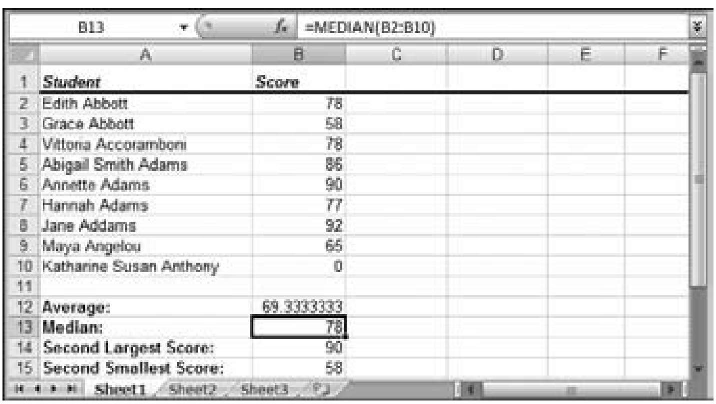

Figure 9-4:
When comparing test scores and class grades, people often use medians instead of averages. The reason? Averages can be unnaturally skewed if one or more values are extremely high or low. For example, if one student turns in a blank paper for a test and receives a 0, the average may be unnaturally low, and not truly representative of a class's performance. You get unnaturally low averages, for instance, when averaging a small set of numbers. Usually, however, the median's quite close to the average.

Tip: Remember, Excel has a quick shortcut for five of its most popular functions—SUM(), AVERAGE(), COUNT(), MAX(), and MIN(). Just choose Formulas→Function Library→AutoSum to perform a quick sum on any group of selected cells. Or, if you click the drop-down arrow on this button, you see a list where you can quickly unleash any of the other four functions.

PERCENTILE() and PERCENTRANK(): Advanced Ranking Functions

If you really want to dissect the test scores from the worksheet in Figure 9-4, you don't need to stop with simple averages, medians, and rankings. You can take a closer look at the overall grade distribution by using the PERCENTILE() and PERCENTRANK() functions.

People often use percentiles to split groups of students into categories. (Technically, a percentile is a value on a scale of 1 to 100, indicating the percentage of scorers in any given ranking who are equal to, or below, the scorer in question.) For example, you might decide that the top 10 percent of all the students in your class will gain admission into an advanced class the following semester. In other words, a student's final grade must be better than 90 percent of his or her classmates. Students at this prestigious level occupy the 90th percentile.

That's where the PERCENTILE() function comes in very handy. You supply the cell range, the percentile (as a fraction from 0 to 1), and the function reveals the numeric grade the student needs to match that percentile. For example, here's how you'd calculate the minimum grade needed to enter the 90th percentile (assuming the grades are in cells B2 to B10):

```
=PERCENTILE(B2:B10, 0.9)
```

Percentiles make the most sense when you have a large number of values because then the distribution becomes the most regular. In other words, there's a wide range of different test scores that are well spread out, with no "clumping" that may skew your results. Percentiles don't work as well with small amounts of data.

PERCENTRANK() performs the inverse task. It takes a range of data and a number from 1 to 100, and gives you the fraction indicating the number's percentile. The following formula, for instance, returns the percentile of the grade 81 (as it compares to a list of other grades):

```
=PERCENTRANK(B2:B13, 81)
```

Using the data shown in Figure 9-5, this formula gives you the fraction 0.727, indicating that 81 falls into the 72nd percentile (and almost hits the 73rd percentile). That is, it's a better score than 72 percent of all the other students, and a lower score than the other 28 percent. In this case, you can also apply the rounding functions explained earlier in the chapter because percentiles are (by convention) whole numbers. You can use ROUND() to round the result of the formula to the closest percentile (73rd), or use TRUNC() to return the highest percentile that the student actually reached (72nd), which is the conventional approach.

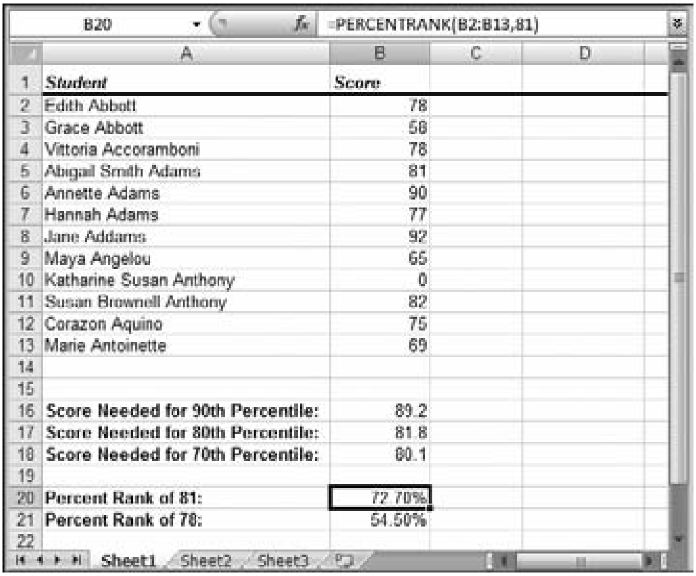

B20 =PERCENTRANK(B2:B13,81)

	A	B	C	D
1	*Student*	*Score*		
2	Edith Abbott	78		
3	Grace Abbott	58		
4	Vittoria Accoramboni	78		
5	Abigail Smith Adams	81		
6	Annette Adams	90		
7	Hannah Adams	77		
8	Jane Addams	92		
9	Maya Angelou	65		
10	Katharine Susan Anthony	0		
11	Susan Brownell Anthony	82		
12	Corazon Aquino	75		
13	Marie Antoinette	69		
14				
15				
16	**Score Needed for 90th Percentile:**	89.2		
17	**Score Needed for 80th Percentile:**	81.8		
18	**Score Needed for 70th Percentile:**	80.1		
19				
20	**Percent Rank of 81:**	72.70%		
21	**Percent Rank of 78:**	54.50%		
22				

Sheet1 Sheet2 Sheet3

Figure 9-5:
Percentiles make a lot of sense when you're trying to analyze how one value fits into a large range. For example, using these student scores, you can tell that students need an 89.2 to enter the 90th percentile. Ideally, you'd use a much larger series of data to calculate percentiles as a small sample like this generates widely varying percentiles each time you enter one new score. Keep in mind that percentiles also work wonders for figuring wages, survey data, and even physical characteristics like heights and weights.

POWER USER'S CLINIC

A Different Way to Calculate Percentiles

It's a shady secret in statistics that there's no universally accepted definition for percentiles. Excel's approach is to designate the smallest number in your sample as the 0th percentile and the highest number as the 100th percentile. In the student grading spreadsheet shown in Figure 9-5, that means the 0th percentile is represented by a test score of 0 and the 100th percentile is represented by a test score of 92.

Excel's approach makes sense for the student score example, because you have all the values. But it doesn't work as well for, say, a researcher gathering data in the field. In this situation, the data in the worksheet is a small sample of the total population, and the range of your sample probably isn't as great as the range of the full population.

For example, imagine you're studying the average length of a certain type of lake trout. In all likelihood, the biggest fish you capture won't be the biggest fish in the lake. Similarly, the smallest fish you capture won't be the smallest fish in the lake. To deal with this situation, statistical pros make an estimate of what the 100th and 0th percentile should be, based on the data they have. The 100th percentile value will be slightly higher than the largest value in the sample, and the 0th percentile value will be slightly smaller than the smallest sample value.

Excel 2010 introduces two functions that let you use this method of percentile calculation: PERCENTILE.EXC() and PERCENTRANK.EXC(). The *EXC* stands for excluded, because the 0th and 100th percentile values are *excluded* (they're not part of the sample data you use for your calculation). Excel 2010 also adds the PERCENTILE.INC() and PERCENTRANK.INC() functions, where the *INC* stands for *included*. They use the same logic as the traditional PERCENTILE() and PERCENTRANK() functions, which are now designated as compatibility functions.

FREQUENCY(): Putting Numbers into Grouped Ranges

The FREQUENCY() function is the last statistical function you'll learn about in this section. Like the functions you've seen so far, FREQUENCY() helps you analyze numbers based on the way they're distributed within a group of other numbers. For example, PERCENTILE() and PERCENTRANK() let you examine how a single test score ranks in comparison to other students' scores. Functions like AVERAGE(), RANK(), and MEDIAN() give you other tools to compare how different values stack up against one another.

FREQUENCY() is different from these functions in one key way: It doesn't compare values against each other. Instead, it lets you define multiple ranges, and then, after chewing through a list of numbers, tells you how many values on the list fall into each range. For example, you could use FREQUENCY() to examine a list of incomes in order to see how many belong in a collection of income ranges that you've identified.

FREQUENCY() is a little more complicated than the other functions you've seen in this chapter because it gives you what's known as an *array* (a list of separate values). When a function gives you an array, it returns multiple results that Excel must display in different cells.

In order to use the FREQUENCY() function properly, you first need to create an *array formula*, which displays its results in a group of cells. (Put another way: A plain-vanilla formula occupies *one* cell; an array formula spans *multiple* cells.)

A worksheet filled with student test scores, like the one shown in Figure 9-6, provides an ideal place to learn how the FREQUENCY() function works, showing you how many students aced a test and how many botched it. Here's how to use the FREQUENCY() function to generate the results shown in cells F6 through F10:

1. **Figure out which cells have the numbers you want to analyze.**

   If you were creating this worksheet, after filling the first two columns with student names and scores, then you'd enter the words *Grade, Maximum*, and *Frequency* in columns D, E, and F. You'd also enter the relevant grade ranges (in cells D6 to D10) and maximum scores (E6 to E10). The ranges and scores entered here are the traditional values, but of course you could enter any range you wanted.

   The FREQUENCY() function is interested in two groups of numbers. B2 through B12 are the student scores you want to analyze. E6 through E10 represent the ranges against which Excel's comparing each student score. (E6 covers scores between 0 and 49, E7 covers 50 to 59, and so on.) Step 3 explains how these cell ranges become part of the FREQUENCY() function as arguments. The array formula itself goes into cells F6 to F10, as explained in the next step.

Tip: If you want to calculate the values that fall in a range with no upper maximum (one that simply captures all values from the previous maximum up), just leave the cell that defines the maximum value blank. This category must be the last one you define. For instance, in the student grade example, you could leave cell E10 blank, although there's no reason to take this step, because you know a student can't score higher than 100.

2. **Select the cells where you want to enter the formula (which is also where the results appear).**

   One result appears in each cell. In this example, you'd select cells F6 to F10.

3. **Type in the formula.**

   The FREQUENCY() function requires two arguments: the cells you want to analyze (B2:B13) and the value ranges the scores you're comparing against (E6:E10). You can manually type in the function arguments, or you can enter them by clicking the worksheet itself.

   Here's the actual formula used in this example:

   ```
   =FREQUENCY(B2:B13,E6:E10)
   ```

4. **When you've finished entering the formula, press Ctrl+Shift+Enter.**

   You can't simply press Enter, because that just enters the formula into the first selected cell. Instead, you need to press Ctrl+Shift+Enter to copy the formula into *all* the result cells. Once you've done this, you'll see the results shown in Figure 9-6.

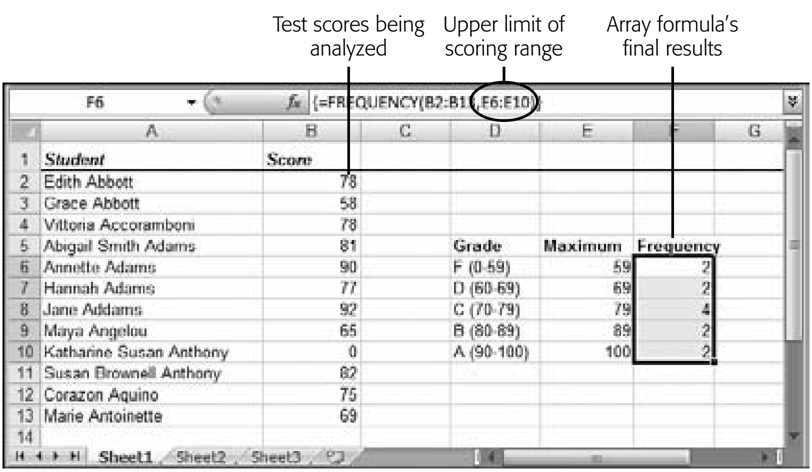

Figure 9-6: *This worksheet shows the completed analysis performed by the FREQUENCY() function. It indicates that four students received a C, while only two received an A. As with any other formula, if you modify the scores or the category maximums, you see the results change automatically. If you want to change your FREQUENCY() formula, make sure you select all the cells in the array, so that you update all of them.*

MODE(): Finding Numbers That Frequently Occur in a List

MODE() is an interesting function that gives you the most frequently occurring value in a data series. For example, if you use MODE() with the worksheet shown in Figure 9-6 like this: *=MODE(B2:B10)*, you get the number 78. That's because the score 78 occurs twice, unlike all the other scores, which appear just once.

MODE() does have a few limitations, though. It ignores text values and empty cells, so you can't use it to get the most common text entry. It also returns the error code #N/A if no value repeats at least twice. And if more than one value repeats the same number of times, then you get the highest number.

This last limitation led the creators of Excel to turn MODE() into a compatibility function and introduce two new variants in Excel 2010. MODE.SNGL() works the same way as the existing MODE(), which means it always gives just a single result. MODE.MULT() is craftier: it returns as many results as you want. The trick is that you need to enter MODE.MULT() as an array function. Here's how to do it:

1. **First, decide how many results you want. Select that many cells, either in a continuous row or column.**

   For example, if you select three cells, you'll get three results: the value that repeats the most, the value that repeats the second most, and the value that repeats third-most. If, for example, only two values repeat, the final cell will get the error code #N/A, which indicates that there's no result for it to show.

2. **With the cells selected, begin typing the formula that uses the MODE.MULT() function.**

   For example, *=MODE.MULT(B2:B10)* searches for repeated values in the column of cells from B2 to B10.

3. **Press Ctrl+Shift+Enter to insert an array formula in all of the selected cells.**

   The list of repeated values will appear, in order of most repeated to least repeated. As with the MODE() function, if the number of repetitions is a tie, the higher value is shown first.

General Math Functions

Excel's Math & Trig category contains functions that go a step beyond the basic mathematical operators (like +,–, /, *, and ^) to help you figure out things like absolute values and greatest (and least) common denominators.

Tip: Poker fan alert: the COMBIN() and PERMUT() functions may help you keep your game sharp even when you're away from the table.

This section covers 14 of the most commonly used functions in this category.

PRODUCT(), FACT(), POWER(), and SQRT(): Products, Factorials, Powers, and Square Roots

These four functions actually offer the same services provided by many of the simple operators covered in the last chapter, but, in some cases, they come with a few twists that may save you time.

The PRODUCT() function takes a list of numbers, multiplies them together, and gives the result. For example, the following formula multiplies 2*3*3, and arrives at 18:

```
=PRODUCT(2,3,3)
```

One interesting characteristic of the PRODUCT() function is that in addition to accepting *individual* cell references, it also supports cell *ranges*. You can use PRODUCT() to multiply all the numbers in a range in the same way you use SUM() to add numbers. Here's an example:

```
=PRODUCT(A2:A12)
```

You can use the FACT() function to calculate a *factorial*, which is a sequence of numbers (starting at 1), each of which is multiplied together to arrive at the result. For example, 5! (pronounced "five factorial") translates into 5*4*3*2*1, which works out to 120. Rather than typing in all these numbers and their accompanying multiplication symbols, you can use the following convenient formula:

```
=FACT(5)
```

People use factorials in many areas of math, including probability theory and calculus. On the other hand, you aren't likely to find factorials cropping up in financial statements or expense reports.

POWER() is a straightforward function used for working with exponents. For example, if you want to calculate 4^3, you can use one of two approaches. You can use the exponentiation operator:

```
=4^3
```

Or you can opt for the POWER() function:

```
=POWER(4,3)
```

The POWER() function is useful if you frequently use cube roots—or any other exponential expressions whose exponent is a fraction—because it lets you write clearer formulas. For example, imagine you want to find the cube root of 4, which, on paper, would be written as $4^{1/3}$. Here's the operator-based expression for that:

```
=4^(1/3)
```

With the POWER() function, you don't need to use the parentheses (and, more important, you have no chance of getting the wrong answer because you forgot to use them):

```
=POWER(4,1/3)
```

If you simply want to calculate the square root of a number, you can use the POWER() function, or you can summon the SQRT() function, which requires only a single argument. Here's a sample formula that gives you 3:

```
=SQRT(9)
```

QUOTIENT() and MOD(): Higher Division

QUOTIENT() and MOD() (see Figure 9-7) are handy for fancy division problems.

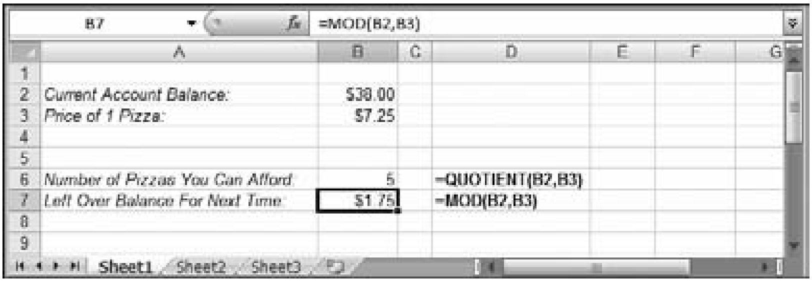

Figure 9-7: *The QUOTIENT() and MOD() functions let you break a number down into two pieces. In this case, the worksheet uses QUOTIENT() to find out how many pizzas you can afford with a set amount of money, and MOD() to find out how much money you'll have left over.*

QUOTIENT() performs *integer* division, which means it ignores the remainder, and just gives you the whole number result. For example, consider the following formula:

```
=QUOTIENT(7,2)
```

This formula performs the calculation 7/2, and gives you a result of 3. Excel discards the remainder of 1. If you performed this division using the division operator, then you'd end up with the fractional value 3.5. Sure, you could truncate the remainder portion using a function like TRUNC(), but the QUOTIENT() function saves you the trouble.

The MOD() function gives you the remainder (also known as the *modulus*) that's left over when you divide a number. For instance, the following formula gives you 1, which is what remains after dividing 7 by 2 using integer division:

```
=MOD(7,2)
```

ABS() and SIGN(): Absolute Value and Determining a Number's Sign

ABS() is a straightforward function that gives you the *absolute value* of a number. The absolute value is the number stripped of any negative sign. Thus, the absolute value of -3 is 3. Here's how that function would look:

```
=ABS(-3)
```

The ABS() function doesn't change a positive number. For example, the absolute value of 8 is just 8. Absolute value calculations work equally well with whole numbers or fractional values. Often, people use ABS() to make sure that a number isn't negative. For example, Excel's financial functions (see Chapter 10) sometimes use negative numbers when you don't expect it—like when you calculate the monthly payment on a loan—to represent the fact that money's leaving your hands. However, to avoid confusing the person reading the spreadsheet, you can use ABS() to make sure that all the information you display is positive, which tends to be what most people understand best.

SIGN() is a related, but less useful function that lets you quickly determine whether a number is positive or negative. SIGN() gives you 1 if the number's positive, 0 if the number's 0 (or the cell's blank), and –1 if the number's negative (as is the following case):

```
=SIGN(-3)
```

SIGN() isn't used too often, because it's usually just as easy to use a condition like A1<0 to check whether a number's negative.

RAND() and RANDBETWEEN(): Generating Random Numbers

The RAND() function gives you a random fractional number that's less than 1, but greater than or equal to 0 (like 0.778526443457398).

```
=RAND()
```

You can use this function to calculate values in any range. For example, if you want a random value between 0 and 50, multiply the RAND() result by 50, as in this formula:

```
=RAND()*50
```

Note: Technically, random numbers from a computer aren't really random. Instead, they're either generated from an algorithm or drawn from a list of pseudo random numbers that the computer has stored in its hardware. When generating a random number, most programs look at the current millisecond value of your computer's clock to decide where to start in this list. Despite these limitations, however, Excel's random numbers still suffice for most spreadsheets—just don't count on using them for advanced statistical models or secret-sauce encryption tools.

Converting a RAND()-generated number into an integer is pretty easy: just use it in conjunction with the INT() function. For example, here's how to get a random whole number from 1 to 6:

```
=INT(RAND()*6)+1
```

For added efficiency, you can save a step with the RANDBETWEEN() function. This function takes two arguments, indicating the minimum and maximum between which you want your randomly selected value to fall. For example, consider the following:

```
=RANDBETWEEN(1,6)
```

This formula generates a random number from 1 to 6 (crucial if your boss is onto your noontime Craps game).

The RAND() and RANDBETWEEN() functions are *volatile functions*, meaning their results change every time Excel recalculates one of the formulas in the worksheet. This behavior can be quite distracting, as actions like inputting new data, editing a cell, or even undoing an action can trigger recalculation—at which point any random values in your worksheet change. If you want to use the RAND() or RANDBETWEEN() function to generate a set of random numbers, and then keep these numbers fixed, then you need to do a little more work. Here's the deal:

1. **Create one or more random numbers using the RAND() or RANDBETWEEN() function.**

   For example, you might want to fill a column using RANDBETWEEN(1,6) to simulate the outcome of rolling a die multiple times.

2. **Select the cells with the random values, and copy them by selecting Home→Clipboard→Copy (Ctrl+C).**

   The next step pastes the cell content back into the same location—but with a twist.

3. **Choose Home→Clipboard→Paste→Paste Values.**

   This step overwrites the formulas and permanently inserts the calculated random numbers—permanent, that is, until you decide to change them.

One neat trick involves using the RAND() function to perform a random sort, which is helpful, say, if you need help deciding which of your 12 sisters to call. To do this, you begin by adding a new column next to the column listing their names. Then, put a random number into each cell, as shown in Figure 9-8.

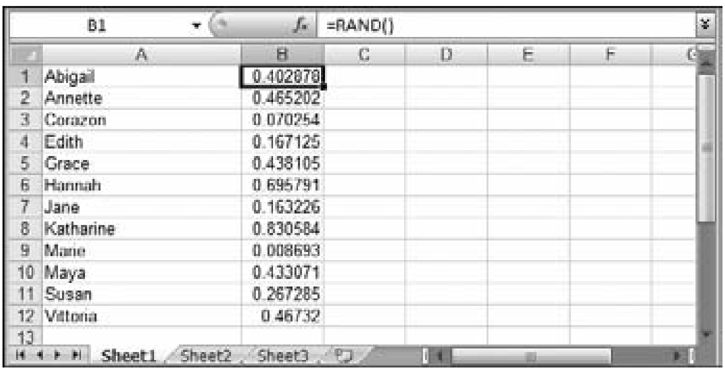

Figure 9-8: *Insert a random value next to each item you want to sort randomly. To speed this process up, you can use the AutoFill feature (page 75) to copy your formula.*

Next, select all the cells you want to sort and the accompanying random values. (In this example, that's cells A1 through B12.) Choose Data→Sort & Filter→Sort, and then create a sort order that uses the random values; Figure 9-9 explains how to perform this sort-order maneuver.

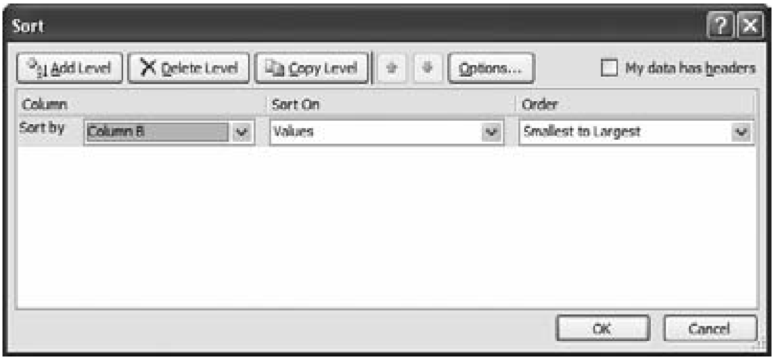

Figure 9-9: *In the "Sort by" box, choose the column with the random value (in this case, Column B). Leave the rest of the settings untouched, and then click OK. (You can also check out page 405 for more detail about Excel's sort feature.) The final result is shown in Figure 9-10.*

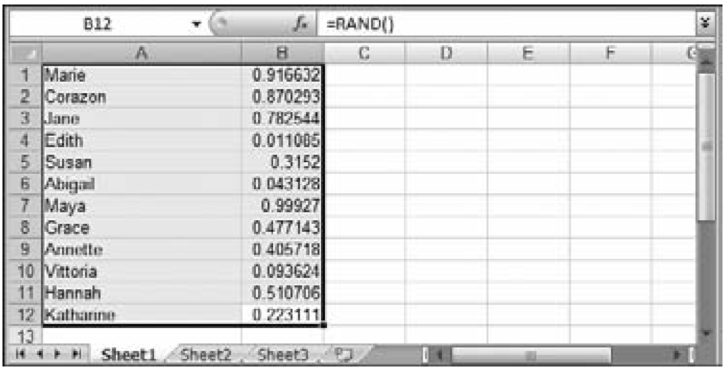

Figure 9-10: *The final result is a randomly ordered list of names. If you want a different random sort, then you need to re-sort the list (choose Data→Sort & Filter→Sort Again). This method works because Excel recalculates the worksheet and changes your random values at the end of every sort operation.*

GCD() and LCM(): Greatest and Least Common Denominator

The GCD() function returns the greatest common divisor that you can use to divide two different numbers. For example, consider the numbers 12 and 8. Both of these numbers share the divisor 4 (12/4 = 3 and 8/4 = 2). Therefore 4 is the greatest common divisor of 8 and 12. The number 12 can also be divided by 6, but it isn't possible to divide 8/6 without generating a fraction.

Here's how you'd arrive at this conclusion using the GCD() function:

```
=GCD(12,8)
```

Consider a few points about GCD(). If you use fractional values, GCD() simply ignores the fractional part of the number. That means it calculates GCD(12.5,8.1) in the same way that it calculates GCD(12,8). Also, Excel always finds a common divisor, as all whole numbers are divisible by 1, even if they don't have any other divisors.

The LCM(), or least common multiple function, works slightly differently. This function answers the question: Among all the numbers being processed by this function, what's the smallest number that you can divide them all *into*? For example, consider again 12 and 8. You can divide both numbers evenly into 24, which makes it their lowest common multiple. Here's the formula that gleans this insight:

```
=LCM(12,8)
```

As with GCD(), LCM() discards fractions in the numbers it's evaluating. The function always finds a common multiple, as you can always multiply the arguments together. For example, 12*8 produces a number that's guaranteed to be a multiple of both 12 and 8, although, in this case, it's not the lowest multiple.

COMBIN() and PERMUT(): Figuring Combinations and Permutations

People use the COMBIN() and PERMUT() functions in probability theory, to count the number of possible outcomes when randomly choosing items from a set. COMBIN() and PERMUT() are short for *combinations* and *permutations*, which are mathematical terms for the number of distinct combinations you can make from a set of items. The main difference lies in the way they consider the order of items in a set (more on that in a moment).

One reason you might use COMBIN() and PERMUT() is to assess the odds in a game of chance. You could use COMBIN() to help you figure out, say, the odds of being dealt an ace of spades from a deck of cards. And people often use both these functions to calculate the occurrence of other random events, like winning the lottery or being hit by a falling anvil.

Both COMBIN() and PERMUT() take two arguments. The first argument is the number of different items in your set. The second argument is the number of items you're randomly selecting each time you "draw" from the set. For example:

```
COMBIN(number_in_set, number_chosen)
```

The key difference between COMBIN() and PERMUT() is that PERMUT() assumes order is important, while COMBIN() assumes it isn't. For example, say you want to count how many possible ways you can draw five cards from a 52-card deck.

You'd use the PERMUT() function as shown here:

```
=PERMUT(52,5)
```

PERMUT() generates a whopping result of *311,875,200,* because it counts every group of five cards as unique, even if the difference is only a matter of card order. Thus, the set 1, 2, 3, 4, 5 is different, for PERMUT(), from 5, 4, 3, 2, 1.

By contrast, when COMBIN() evaluates those same parameters, it returns a value of *2,598,960*.

As you can see, permutations are far more plentiful than combinations. In fact, you can convert the number of combinations into the number of permutations by multiplying it by 5! (5*4*3*2*1), which is a measure of how many different ways you can arrange the order of five different items.

To go much further with PERMUT() and COMBIN(), you'll need a refresher course in probability theory. However, you may be interested to know that you can find the *probability* of an event by dividing the number of satisfactory outcomes (the number of outcomes that meet your criteria) into the number of total outcomes (the total number of possibilities). For example, the chance of your getting a specific hand of cards is expressed like so:

```
=1/COMBIN(52,5)
```

Here, you're only interested in a single outcome (a hand that has every card you want). Thus, there's one satisfactory outcome. You can use the COMBIN() function to find the total number of possibilities by counting the different possible draws.

To see your percentage chance of getting the cards you want, you can multiply the result by 100 or use the Percentage number style.

Trigonometry and Advanced Math

People sometimes describe trigonometry as the science of circles and angles. It's trigonometry that helps you calculate the hypotenuse of a triangle or the diameter of a circle. However, when you use trigonometry in Excel, you probably won't be worrying about shapes; instead, you'll be using some type of formula from a scientific field that requires common trigonometric calculations like the cosine or tangent.

Students of the sciences know that trigonometry turns up anywhere you need to think about space, including geography, astronomy, kinematics, and optics. Less direct applications of trigonometry turn up in just about every other scientific field—from chemistry to social science.

You can use Excel's trigonometric functions quite easily. Just choose the formula you need and the appropriate argument, much as you'd use a pocket calculator. (Knowing when to use these trigonometry functions is another issue, and it may require a refresher course in mathematics.) Table 9-1 lists the functions you can use. Figure 9-11 shows how you can use trigonometry to solve a common high-school math problem: you have the distance to a tree (which you can easily measure), and the angle between your current position and the top of the tree (which you can estimate). The problem is you have no way to judge the height of the tree.

Table 9-1. *Trigonometry functions*

Function	Syntax	Description
ACOS()	ACOS(number)	Calculates the inverse cosine (also known as the arccosine) of a number.
ACOSH()	ACOSH(number)	Calculates the inverse hyperbolic cosine of a number.
ASIN()	ASIN(number)	Calculates the inverse sine (also known as the arcsine) of a number.
ATAN()	ATAN(number)	Calculates the inverse tangent (also known as the arctangent) of a number.
ATANH()	ATANH(number)	Calculates the inverse hyperbolic tangent of a number.
COS()	COS(angle)	Calculates the cosine of an angle.
COSH()	COSH(angle)	Calculates the hyperbolic cosine of an angle.
DEGREES()	DEGREES(angle)	Converts an angle in radians into degrees. 360 degrees equals 2π radians.
EXP()	EXP(power)	Calculates e (the natural logarithm constant), raised to the power you specify.
LN()	LN(number)	Calculates the natural logarithm of a number (which is the logarithm with the base of e).
LOG()	LOG(number, base)	Calculates the logarithm of a number with the base you specify.
LOG10()	LOG10(number)	Calculates the base 10 logarithm of a number.
PI()	PI()	Returns the mathematical constant pi (π), which is approximately 3.14.
RADIANS()	RADIANS(degrees)	Converts an angle in degrees into an angle in radians. 360 degrees equals 2π radians.
SIN()	SIN(angle)	Calculates the sine of an angle.
SINH()	SINH(angle)	Calculates the hyperbolic sine of an angle.
TAN()	TAN(angle)	Calculates the tangent of an angle.
TANH()	TANH(angle)	Calculates the hyperbolic tangent of an angle.

In addition to its trigonometric functions, Excel also includes a number of specialized mathematical functions, like those used with matrices. Table 9-2 summarizes these functions.

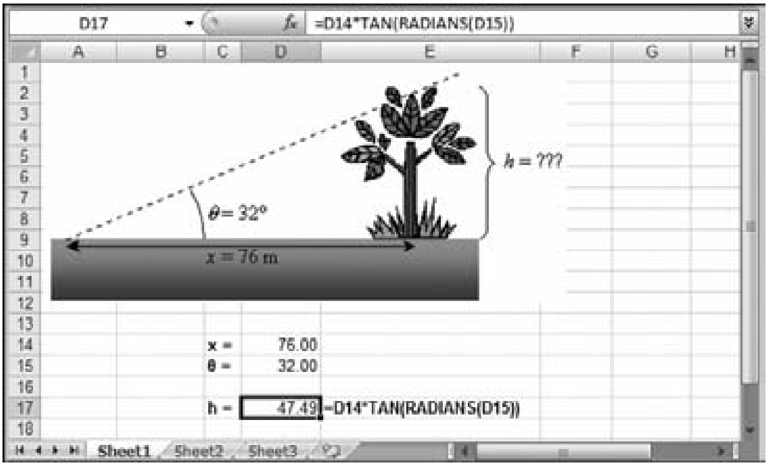

Figure 9-11: *You can solve this problem by using the trigonometric identity tan = h/x. You can rearrange this formula to find h (the height of the tree) so that h = x*tan. Be aware that Excel calculates the tangent using radians (radians are a measure of an angle, like degrees). Fortunately, Excel provides a RADIANS() function for converting degrees into radians, as well as a DEGREES() function for converting radians into degrees. In this example, the RADIANS() function was first used to convert degrees into radians; that result was then used as the argument in the TAN() function. The image was inserted using the Insert→Illustrations→Picture command, which you'll learn about in Chapter 19.*

Table 9-2. *Miscellaneous math functions*

Function	Syntax	Description
FACTDOUBLE()	FACTDOUBLE(number)	Calculates the double factorial of a number. The double factorial's like a factorial, but it skips every second number. Thus, 5!! is 5*3*1 (whereas 5! is 5*4*3*2*1).
MDETERM()	MDETERM(range)	Calculates the determinant of a matrix.
MINVERSE()	MINVERSE(range)	Returns the inverse of a matrix. This function's an array function that returns multiple results.
MMULT()	MMULT(range1, range2)	Returns the product of multiplying two matrixes. This function's an array function that returns multiple results.
MULTINOMIAL()	MULTINOMIAL(range)	Returns the multinomial of a set of numbers. The multinomial's the ratio of the factorial of a sum of values to the product of factorials.

ROMAN()	ROMAN(number, [form])	A bizarre function that converts an ordinary number (from 1 to 3999) into Roman numerals. For example, ROMAN(4) returns the text *IV*. You may find large Roman numerals very confusing, so use the optional form argument to tell Excel to simplify the Roman numeral by applying a kind of rounding off effect. Form can range from 0 (which generates the unrounded Roman numeral) to 4 (which generates a much shorter but slightly less precise version). Search Excel's Help for *ROMAN* for a complete explanation.
SERIESSUM()	SERIESSUM(x, n, m, coefficients)	Calculates the sum of a power series, where x is the base of the power series, *n* is the initial power to raise x to, *m* is the increment to the power for each term in the series, and *co* are the coefficients by which each successive power of x is multiplied. For more gory mathematical details, check out a detailed example of the series formula in Excel's function reference (choose Formulas→Function Library→Insert Function, find the function, and then click the "Help on this function" link).
SQRTPI()	SQRTPI(number)	Multiplies a number by pi (π), and then calculates the square root of the result.

Advanced Statistics

Excel offers a smorgasbord of number-crunching tools; in fact, it has more than 50 advanced statistical functions that this chapter hasn't even mentioned. Many of these functions will appeal only to those who already understand concepts like standard deviation, theoretical distributions, and dimensionless Pearson product moment correlation coefficients (a subject of keen interest to all six people in the world who currently use such numbers).

But even if you don't understand *any* of those concepts (and if you're still wondering whether your inner mathematician may yet emerge), Excel has two advanced statistical functions—FORECAST() and TREND()—that give you a taste of all that awaits the advanced statistician. While both FORECAST() and TREND() rely on some complex, behind-the-scenes computations, you can use them quite easily, even if you're not a specialist.

The FORECAST() and TREND() functions predict missing values based on existing data. These "missing" values may represent experimental data you didn't collect, or future sales numbers for dates that haven't happened yet. Excel examines the data you have so far (the more the better), and uses a mathematical technique known as the *method of least squares* to generate a trend line. Excel then uses this trend line to generate the other values you've requested.

Both FORECAST() and TREND() work similarly and produce the same answers. However, FORECAST() is designed to calculate a single value. The TREND() function, on the other hand, is an array function that gives you multiple values. It works a little like the FREQUENCY() function described on page 277.

The FORECAST() and TREND() analyses work by trying to spot a relationship between two sets of numbers, designated as the X-axis and Y-axis. To predict a new value in the FORECAST() function, for instance, you supply the function with a new X value. It then uses the known relationship between the X and Y values to guess the corresponding Y value. The syntax is as follows:

```
FORECAST(x, known_y_values, known_x_values)
```

For example, imagine you're trying to study the relationship between the number of product inquiries received during one month and the actual sales recorded in the next month. You have plenty of historical data, and you'd like to use that to predict how many sales you should expect next month (based on the number of inquiries you've received during the current month). The existing inquiry numbers are the first set of known values, so they make up the X-axis values (the cells A2:A9, as shown in Figure 9-12). The second set of known values is the actual recorded sales in the following month (the Y-axis values, shown in cells B2:B9). You want to predict a new Y-axis value for expected sales based on the current inquiries (cell A12). Here's the final formula:

```
=FORECAST(A12,B2:B9,A2:A9)
```

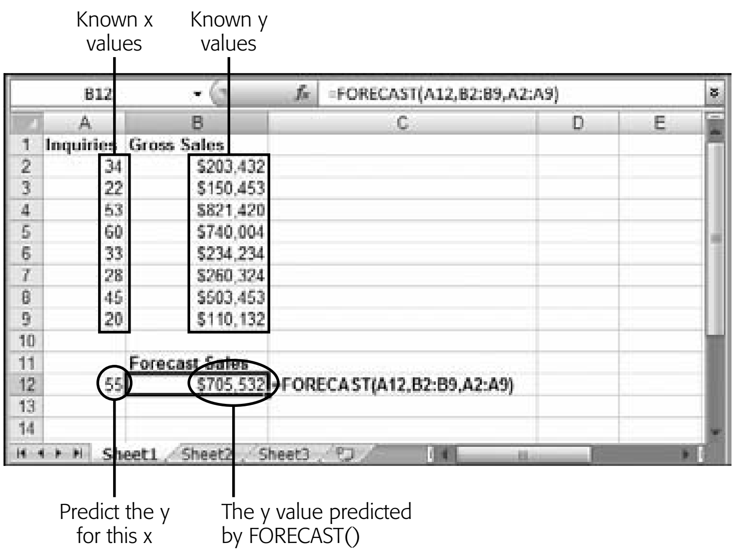

Figure 9-12: *This worksheet shows how you can forecast a future value based on a historical relationship by using the FORECAST() function. The example here predicts sales based on an analysis of the previous relationship between inquiries and gross sales.*

Note: The FORECAST() and TREND() functions can always find some sort of relationship using your data. However, if your data doesn't have a real correlation, then the values they predict will be meaningless. For example, the number of product inquiries may not be an effective predictor of future sales, especially when you consider other factors. Determining where correlations exist and what relationships are statistically significant is a much more subtle and complex issue than just using the FORECAST() and TREND() functions. An understanding of basic statistics will set you down the right path.

You can also use the TREND() function to generate expected sales based on multiple inquiry numbers (as opposed to the single sales value that FORECAST() generates). Using the previous example, you could theoretically use TREND() to try and predict sales for the next six months. In truth, that approach doesn't really make sense because, while you have previous sales data to draw from, you don't know how many inquiries you'll receive in future months.

You're better off using TREND() to project future sales based on a slightly different set of data. You still use previous sales, but rather than using *inquiries* as your second known factor, you use the *progression of time*. In effect, you're asking TREND() if your sales over periods 1, 2, and 3 were 100, 150, and 200, what would sales look like over periods 4, 5, 6? Without further ado, here's what you need to generate new, predicted sales figures using the TREND() function:

1. **Make sure your worksheet has all the data your TREND() formula needs.**

   As shown in Figure 9-13, you need three groups of numbers: known y-values (historical sales figures), known x-values (previous time periods), and the new x-values (known, new time periods). You want the formula to figure out the new y-values (that is, the sales figures).

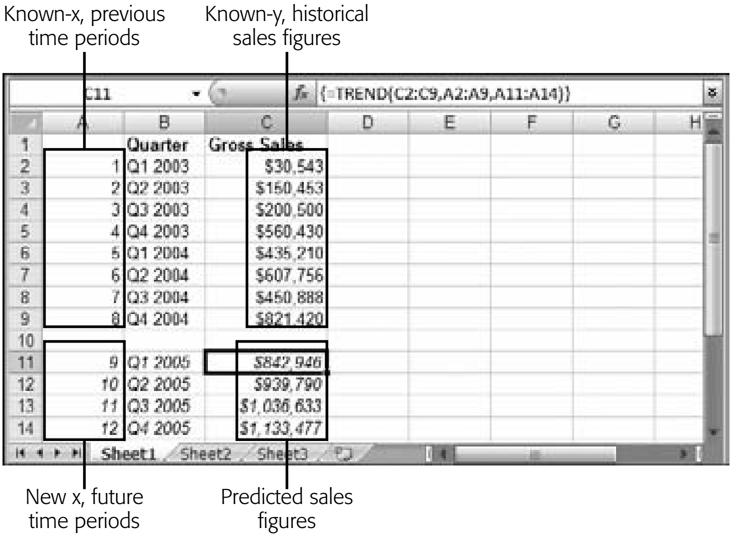

Figure 9-13: *All signs point to a strong year. Based on the previous growth of sales, Excel cheerily predicts four more quarters of escalating growth. (Excel has predicted the italicized cells.) Keep in mind, of course, that Excel generated these numbers using the pattern of the last two years. Various other factors could make this trend line completely irrelevant.*

2. **Select the range of cells where you want to display the predicted values.**

   In this example, four new sales figures appear in cells C11:C14.

3. **Enter the array formula using the TREND() function.**

   Here's the formula used in this example:

   ```
   =TREND(C2:C9, A2:A9, A11:A14)
   ```

4. **When you've finished typing in the formula, hit Ctrl+Shift+Enter.**

   You can't simply press Enter, as that enters the formula only into the first selected cell. Instead, you need to press Ctrl+Shift+Enter to copy the array formula into all the result cells. You can tell that an array formula is at work by the curly braces {} in the formula bar, which Excel adds automatically.

Excel provides many more advanced statistical functions. For the full list, refer to the Formulas→Formula Library→More Functions→Statistical group. Table 9-3 outlines some of the more commonly used functions in this group.

POWER USERS' CLINIC

An Advanced Engineering Toolkit

Excel includes an Engineering category of functions with some very specialized mathematical functions. To take a look, select Formulas→Function Library→More Functions→Engineering.

These functions tackle two primary problem areas: imaginary, or complex numbers (numbers that incorporate *i*, the square root of –1); and number conversions between binary, decimal, and octal representations. If you need to use these functions for advanced mathematical spreadsheets, they're quite understandable. On the other hand, if you don't know the difference between a complex conjugate and a vector, they probably aren't very useful.

Excel doesn't corner the market on advanced mathematical functions. You can get many more specialized functions if you install add-ins from other vendors. (Try Googling "Excel add-ins" to start searching.) And in Chapter 29, you'll learn how you can write the custom logic for your own functions using Visual Basic for Applications (VBA).

Table 9-3. *Advanced statistical functions*

Function	Syntax	Description
AVEDEV()	AVEDEV(range)	Calculates the average deviation, which is the average of how much a series of data points deviate from the mean.
CORREL()	CORREL(range1, range2)	Calculates the correlation coefficient between two sets of data, which is a measure of how the sets are related. The correlation coefficient's always between –1 and 1 (the closer to 1 the stronger the correlation, while 0 means no correlation at all).
COVAR()	COVAR(range1, range2)	Calculates the covariance, which is the average of the products of deviation for each data point pair. Covariance determines if a relationship exists between data sets.
LINEST()	LINEST(known_y_values, known_x_values)	Calculates the best-fit line using the least squares method. This array function returns both the slope and the y-intercept.
NORMDIST()	NORMDIST(x, mean, standard_dev)	Calculates the normal cumulative distribution for a specified mean and standard deviation.
NORMINV()	NORMINV(probability, mean, standard_deviation)	Calculates the inverse of the normal cumulative distribution for a specified mean and standard deviation.
SLOPE()	SLOPE(known_y, known_x)	Calculates a regression line through the supplied data points, and returns its slope.
STDEV()	STDEV(range)	Calculates the standard deviation of a data set. Excel ignores blank and text cells.
STDEVA()	STDEVA(range)	Calculates the standard deviation of a data set. Excel uses text cells as 0 values, but ignores blank cells.
VAR()	VAR(range)	Estimates the variance of a sample population. Excel ignores blank and text cells.
VARA()	VARA(range)	Estimates the variance of a sample population. Excel uses text cells as 0 values, but ignores blank cells.

CHAPTER
10

Financial Functions

Calculating cosines and analyzing statistical trends may seem like a great way to tap Excel's brainpower. But what if you want help tracking the movement of small pieces of green paper from one bank account to another? Fortunately, Excel's no slouch when it comes to dealing with money. In fact, it comes with a slew of financial functions that can help you figure out the bottom line like a pro.

Microsoft designed the majority of Excel's financial functions to help you determine how numbers change over time. You can use them to track soaring assets or mounting debts. In this chapter, you'll start by learning the basic concepts, like *present value* and *payment period,* which Excel bases all of its financial functions on. Afterwards, you'll take a close look at some of the most useful financial functions, complete with examples that answer common questions about mortgages, loans, and investments.

The World of Finance

Before you start using Excel's financial functions, it helps to understand the financial concepts that lie at the heart of many of these operations. Here are some terms that those wacky accountants love to use:

- **Present Value (PV)**. The value of an investment or loan at the very beginning of its life. (Hopefully, after this point, the investment will rise or you'll begin to pay off the loan.) This number's also called the *principal.*
- **Future Value (FV)**. The value of an investment or loan at some point in the future.

- **Rate**. The rate at which an investment or loan will increase or decrease. A typical investment could have an annual interest rate of five percent, which means that after one year, the future value will be five percent larger than the present value.
- **Payment (PMT)**. An amount of money that's being contributed to an investment or loan. It's a regular contribution that's usually made at the same time the interest is calculated.
- **Number of Payment Periods (NPER)**. The total number of payment periods between the present value and the future value of an investment or loan. If you've got a 3-year car loan (with payments due monthly), then the NPER equals 36—in other words, there are 12 payment periods each year, for 3 years.

Usually, you'll use one of Excel's financial formulas when you have *some* of the information from the previous list, but not all of it. Maybe you want to discover how much an investment will increase in value over a set amount of time. Or perhaps you want to determine what rate of savings you need to maintain in order to collect a certain amount of money for retirement.

UP TO SPEED

Knowing When to Use Positive and Negative Numbers

Since Excel lets you use the same functions whether you're calculating investments or loans, the big thing to remember–and one of the hardest–is whether to use positive or negative numbers in your formulas. The rules of thumb are as follows:

- If it's money that's leaving your hands, whether it's a deposit *to* an account or a payment *for* a loan, then the number should be negative.
- If it's money that's coming *to* you, whether you're receiving a loan or an investment that's matured, then the number should be positive.

Be on the lookout for numbers bearing an incorrect sign. If the results of your formula indicate an unusually large monthly payment, a shrinking investment, or a growing loan balance, numbers with incorrect signs may be at fault. The easiest way to test for this problem is to slightly modify your formula, and see how it changes. If you *decrease* the amount of money you want to borrow and the monthly payment *increases*, then one of your numbers probably has the wrong sign.

Financial Functions

Excel includes several dozen financial formulas, but non-accountants use only a handful of these regularly. These all-purpose functions, described in the following sections, are remarkably flexible. You can use them to make projections about how investments or loans will change over time, or to answer hypothetical questions about investments or loans. Perhaps you want to determine the interest rate or length of time you need to reach an investment goal or pay off a loan.

FV(): Future Value

The FV() function lets you calculate the future value of an investment, assuming a fixed interest rate. Perhaps the FV() function's most convenient feature is that it lets you factor in regular payments, which makes it perfect for calculating how money's accumulating in a retirement or savings account.

To understand how FV() works, it helps to start out by considering what life would be like without FV(). Imagine that you've invested $10,000 that's earning a fixed interest rate of 5 percent over one year. You want to know how much your investment's going to be worth at the end of the year. You can tackle this problem quite easily with the following formula:

```
=10000*105%
```

This calculation provides you with the future value—that is, the initial 100 percent of the principal plus an additional 5 percent interest.

It's just as easy to determine what happens if you keep your money invested for two years, reinvesting the 5 percent interest payment for an additional year. Here's the formula for that:

```
=10000*105%*105%
```

In the end, you wind up with a tidy $11,025.

Clearly, for these simple calculations, you don't need any help from Excel's financial functions. However, real accounts aren't always this simple. Here are two common problems that can make the aforementioned calculation a lot more difficult to solve when using a do-it-yourself formula:

- You've invested your money in a savings account that pays interest monthly. Even though the annual rate of interest's the same as in the earlier examples, your money accumulates faster thanks to *compound interest* (that is, interest earned on the interest you've previously earned).
- You make regular deposits to the investment account. You can't calculate this extra amount of money separately because it also accumulates interest, starting from the date on which you deposited it.

All of a sudden, this calculation isn't so easy to perform. Of course, you could solve it on your own, but the formula you'd need to write is startlingly complex. At this point, the FV() function becomes a lot more attractive. Here's how the function breaks down (remember, brackets indicate optional arguments):

```
FV(rate, nper, payment, [pv], [type])
```

- The *rate* is the interest rate your investment's earning.

- The *nper* is the number of interest payments. If your account receives interest once a year and you invest your money over a two-year period, then the nper will be 2. If you're making regular contributions, this value also specifies the number of contributions you're making. FV() assumes that every contribution's made on the same day as the interest's generated, which is an unfortunate limitation. (You can get around this limitation by using FV() several times to calculate the value of individual contributions, and then adding these together.)
- The *payment* is the amount of the contribution you want to make regularly. Set this to 0 if you don't want to add anything.
- The *pv* is the present value, or the initial balance of your account. If you omit this value, then Excel assumes you start with nothing. As a result, you'll need to include something other than 0 for the payment.
- The *type* indicates the timing of the payment. If you specify 0 (or omit this value), then the payment's made at the end of the period. If you specify 1, then the payment's made at the beginning, giving your interest just a little bit more time to compound.

The only trick—some would say Bizarro World element—with FV() is that you need to make sure both the payment and the initial balance *(pv)* are *negative* numbers (or zero values). Huh? Isn't this money that's *accumulating* for you? Here's what's going on: In Excel's thinking, the initial balance and the regular contributions are money you're *handing over,* so these numbers, consequently, need to be negative. The final value is positive because that's the total you're getting back.

Continuing with the earlier example, here's how you'd rewrite the formula that calculates the return on a $10,000 investment after one year earning five percent annual interest:

```
=FV(5%, 1, 0, -10000)
```

Note: The formula knows that this is a one-year investment because the 5 percent is an *annual* interest rate, and you've indicated (through the second argument) that there's only one interest payment being made.

This calculation returns the expected value of $10,500. But what happens if you switch to an account that pays monthly interest? You now have 12 interest payment periods per year, and each one pays a 12th of the total 5 percent interest (see Figure 10-1). Here's the revised formula:

```
=FV(5%/12, 12, 0, -10000)
```

Note that two numbers are different from the original formula: the interest (divided by 12 because it's calculated per payment period) and the number of payment periods. The new total earned is a slightly improved $10,511.62.

Note: The interest rate and the number of payment periods always must use the same time scale. If payments are monthly, then you must use a monthly interest rate. (Remember, to calculate the monthly interest rate, just take the yearly interest rate and divide by 12.)

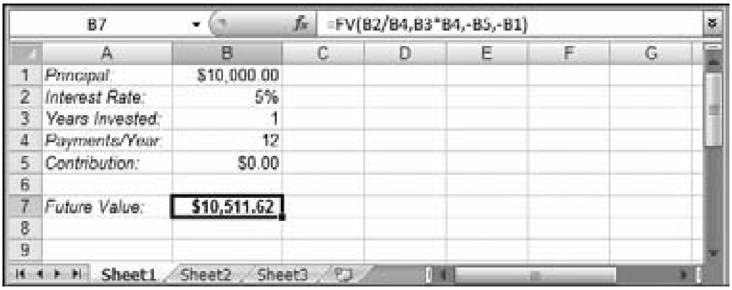

Figure 10-1: *This worksheet depicts a future value calculation using the FV() function and the information in cells B1 to B5. In this example, the initial balance is $10,000, the interest is five percent (paid monthly), and you don't make any additional contributions. Note that the nper (number of interest payments) is calculated by multiplying Years Invested by Payments/Year.*

Finally, here's the calculation over two years. The only number that changes is the number of payment periods.

```
=FV(5%/12, 24, 0, -10000)
```

And while you're at it, why not check what happens if you make a monthly contribution of $100 to top up the fund? In this case, you'll make the payment at the beginning of the month. The required formula is shown here:

```
=FV(5%/12, 24, -100, -10000, 1)
```

Now, the total tops $13,000 ($13,578.50 to be precise). Not bad for a little number crunching!

Incidentally, FV() works just as well on loan payments. Say you take out a $10,000 loan and decide to repay $200 monthly. Interest is set at 7 percent and calculated monthly. The formula that tells you your *outstanding* balance—that is, the amount that you still owe—after three years is as follows:

```
=FV(7%/12, 3*12, -200, 10000)
```

Note that the balance begins positive, because it's money that you've *received,* while the payment's negative, because it's money you're paying out. After three years, you'll probably be disappointed to learn that the loan hasn't been paid off. The FV() function returns –4,343.24, which is the balance remaining to be paid.

Note: The FV() function simply can't solve some of the more complex problems. It can't take into account, for example, payments that change, interest rates that change, or payments that aren't made on the same date as the interest date. You can sometimes deal with an interest rate that varies from one year to another by using the FV() function more than once, to calculate the years separately. Excel includes some more advanced functions that can help you out–if you have an advanced accounting degree. If not, you may be interested in checking out a heavyweight reference book like *Financial Analysis with Microsoft Excel 2007*, (South-Western College Pub, 2009).

FREQUENTLY ASKED QUESTION

Monthly Versus Annual Interest Rates

How do I compare monthly and annual interest rates?

A dollar may always be a dollar, but a five percent interest rate may not always pay what you expect. The important factor is how often your interest's *compounded,* which all depends on how often the interest payments are paid to your account (or calculated against your loan). As you've seen with the FV() function examples, if 5 percent interest is compounded monthly, then you have a chance to make a little extra cash. That's because your interest payments make it into your account before the year's over, giving you the chance to earn some additional interest on the amount of interest you've received so far.

When comparing different types of investments, it's often important to get a feeling for how much difference compound interest makes. Say you need to know which investment option is better–one that pays 5 percent after one year, or a savings account that pays a monthly interest of 4.75 percent. To answer this question, you can run the numbers with FV(), or you can use a shortcut, available courtesy of the EFFECT() function. This financial function calculates the effective annual interest for an interest rate that's compounded multiple times. All you need to do is specify two arguments: the rate and the number of times it'll be compounded. In the case of a 4.75 percent rate compounded 12 times over the course of the year, here's your formula:

```
=EFFECT(4.75%, 12)
```

This formula returns a result of 0.485, or 4.85 percent. In other words, 4.75 percent compounded monthly is equivalent to 4.85 percent compounded annually, provided you don't make any more payments during the lifetime of the investment. Assuming you don't need early access to your money, the 5 percent fixed investment's still the better deal. How much better? Well, once again, you could use FV() to calculate the future value. But seeing as you've used EFFECT() so far, why not just try the following shortcut?

```
=(5%-EFFECT(4.75%,12)) * 10000
```

Keeping in mind Excel's precedence rules (page 229), which specify that expressions inside parentheses get evaluated first, here's what's happening: first, the EFFECT() function repeats the calculation just performed, returning a value of 0.485. Excel then subtracts this value from 5 percent (0.5), leaving you with the difference between the two rates: about 0.15 percent. Once you multiply that by the total amount of money you've invested ($10,000), you'll get the extra bit of money you make.

In this example, the difference between the fixed investment and the savings account is a modest $14.52, which may be a price worth paying for flexibility.

PV(): Present Value

The PV() function calculates the initial value of an investment or a loan (which is also called the present value). The PV() function takes almost the same arguments as the FV() function, except the optional *pv* argument is replaced with an optional *fv* (or future value) argument:

```
PV(rate, nper, payment, [fv], [type])
```

At first glance, the PV() function may not seem as useful as FV(). Using future information to calculate a value from the past somehow seems counterintuitive. Can't people just dig up the principal value from their own records?

Yes, but the real purpose of PV() is to answer hypothetical questions. If you know what interest rate you can get for your money, how long you'll be invested, and what future value you hope to attain, you can pose the following question: What initial amount of money do I need to come up with? The PV() function can provide the answer.

Consider this formula:

```
=PV(10%/12, 25*12, 0, 1000000)
```

The question Excel answers here is: In order to end up with $1,000,000, how much money do I need to invest initially, assuming a 10 percent annual interest rate (compounded monthly) and a maturation period of 25 years? The PV() function returns a modest result of $82,939.75.

If you don't have that much cash on hand, you can supplement your principal with a regular investment. The following formula assumes a monthly payment of $200, paid at the beginning of each month. Note that you should type in a negative number, because it's money you're giving up:

```
=PV(10%/12, 25*12, -200, 1000000)
```

This change decreases the principal you need to a cool $60,930.30.

See Figure 10-2 for a much larger collection of PV formulas designed to answer the popular question: Who wants to be a millionaire? These formulas assume no monthly payment, and display the initial balance you need to reach a cool million. Sharp eyes will notice that the formula uses partially fixed references (like $A9 instead of A9). You can copy the formula from one cell to another without mangling the cell reference. $A9 tells Excel it can change the row but not the column. As a result, no matter where you copy it, this reference always points to the cell in column A of the current row (which holds the interest rate).

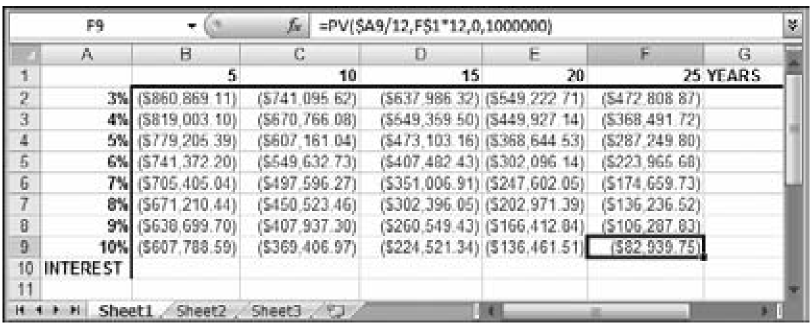

F9 =PV($A9/12,F$1*12,0,1000000)

	A	B	C	D	E	F	G
1		5	10	15	20	25 YEARS	
2	3%	($860,869.11)	($741,095.62)	($637,986.32)	($549,222.71)	($472,808.87)	
3	4%	($819,003.10)	($670,766.08)	($649,359.50)	($449,927.14)	($368,491.72)	
4	5%	($779,205.39)	($607,161.04)	($473,103.16)	($368,644.53)	($287,249.80)	
5	6%	($741,372.20)	($549,632.73)	($407,482.43)	($302,096.14)	($223,965.68)	
6	7%	($705,405.04)	($497,596.27)	($351,006.91)	($247,602.05)	($174,659.73)	
7	8%	($671,210.44)	($450,523.46)	($302,396.05)	($202,971.39)	($136,236.52)	
8	9%	($638,699.70)	($407,937.30)	($260,549.43)	($166,412.84)	($106,287.83)	
9	10%	($607,788.59)	($369,406.97)	($224,521.34)	($136,461.51)	($82,939.75)	
10	INTEREST						
11							

Sheet1 Sheet2 Sheet3

Figure 10-2: *With the right amount of money and a sufficient amount of time, anyone can become a millionaire. This table shows the effect of different interest rates and investment time periods in an effort to make $1,000,000. The selected cell shows that if you assume your investment has 25 years to grow at 10 percent per year, you need to start out with $82,939.75 to hit the millionaire mark.*

If you don't have the patience to wait for accruing money, you may be more interested in using PV() to determine how much money you can afford to borrow *now*. Assuming you can pay $250 a month for a three-year loan at a 7 percent annual interest rate, here's how you'd calculate the size of the loan:

```
=PV(7%/12, 3*12, -250, 0)
```

The answer? $8,096.62. In this example, the future value *(fv)* is 0, because you want to pay the loan off completely.

PMT(), PPMT(), and IPMT(): Calculating the Number of Payments You Need to Make

The PMT() function calculates the amount of the regular payments you need to make, either to pay off a loan or to achieve a desired investment target. Its list of arguments closely resembles the FV() and PV() functions you've just learned about. You specify the present value and future value of the investment and the rate of interest over its lifetime, and the function returns the payment you'd need to make in each time period. Here's how the function breaks down:

```
PMT(rate, nper, pv, [fv], [type])
```

If you don't specify a future value, then Excel assumes it's 0 (which is correct if you're performing the calculation to see how long it'll take to pay off a loan). Once again, the *type* argument indicates whether you make payments at the beginning of the payment period (1) or at the end (0).

To consider a few sample uses of the PMT() function, you simply need to rearrange the formulas you've been using in the FV() and PV() sections. If you have a 7 percent interest rate (compounded monthly) and a starting balance of $10,000, how much do you need to pay monthly to top it up to $1,000,000 in 30 years? The PMT() function provides your answer:

```
=PMT(7%/12, 12*30, -10000, 1000000)
```

The result—$753.16—is a negative number, because this is money that you're *giving up* each month.

A loan calculation is just as easy, although, in this case, the present value becomes positive, since it represents money you received when you took out the loan. To determine the payments needed to pay back a $10,000 loan (that comes with a 10 percent annual interest rate) over five years, you need this formula:

```
=PMT(10%/12, 12*5, 10000, 0)
```

Assuming you make payments at the end of each month, the monthly payment is $212.47. If you add a *type* argument of 1 to pay at the beginning of the month, then this amount decreases to $210.71.

The PPMT() and IPMT() functions let you take a closer look at how you're repaying your loan, as they both analyze a single loan payment. PPMT() calculates the amount of a payment that's being used to pay down the loan's principal, while IPMT() calculates the amount of a payment that's being used to pay back accrued interest. You'll find both functions extremely useful if you need to figure, for tax purposes, what portion of your monthly loan payment is paying interest, versus what portion's paying off a loan's principal (see Figure 10-3).

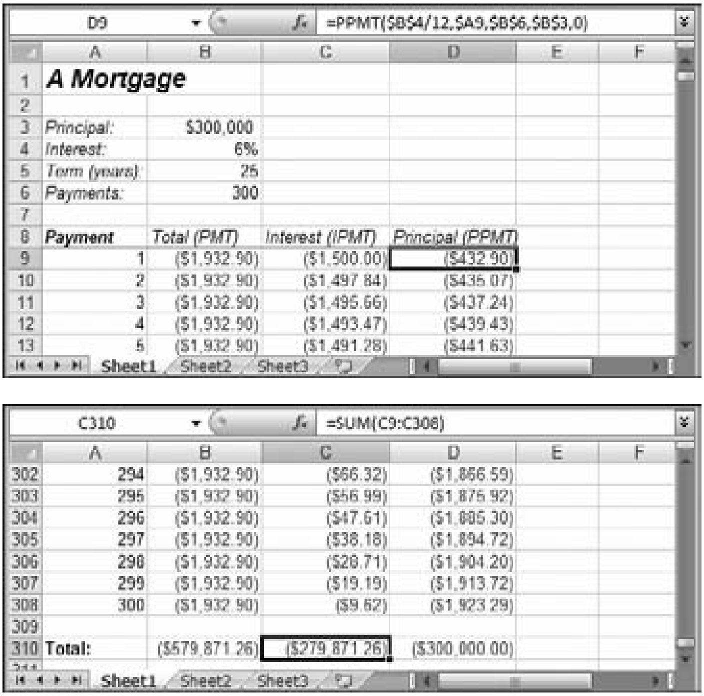

Figure 10-3:
Top: A home mortgage is an ideal reason to use the PMT(), PPMT(), and IPMT() functions. This worksheet depicts each of the 300 mortgage payments you'd make over the course of a 25-year $300,000 mortgage. The first payment contributes a measly $432.90 to paying off the principal.

Bottom: This table shows that the total of all payments made equals $579,871.26—the amount you actually spent to pay off the $300,000 loan. This type of table is an amortization table.

Note: Over the course of a loan, the principal will gradually decrease, as will the amount of each payment going to pay the interest. But for each payment, it's always true that the *PMT = PPMT + IPMT*. In other words, your total monthly payment is always going to remain the same.

The PPMT() and IPMT() functions take the same arguments as the PMT() function with the only difference being that you need to specify one additional argument: *per* (short for period). This argument indicates which payment you're analyzing. For example, a *per* of 1 examines the first payment. A *per* of 6, meanwhile, analyzes the sixth payment, which, assuming you're paying the loan on a monthly basis, happens halfway through the year. Here are both functions:

```
PPMT(rate, per, nper, pv, [fv], [type])
IPMT(rate, per, nper, pv, [fv], [type])
```

As a quick example, consider the first payment of the $10,000 loan analyzed earlier using the PMT() function. You already know that each payment's $212.47. But what portion of this amount is actually paying down the principal? For the first payment, you can calculate the answer as follows:

```
=PPMT(10%/12, 1, 12*5, 10000, 0, 0)
```

The answer's a relatively minute $129.14. On the last payment, however, the situation improves:

```
=PPMT(10%/12, 12*5, 12*5, 10000, 0, 0)
```

Now a full $210.71 is paying off the last remaining bit of the balance.

Note: Clearly, the PPMT() and IPMT() functions don't work if you specify an argument of *per* that's greater than *nper*. In other words, you can't analyze a payment you don't make! Should you ever make this mistake, the function returns a #NUM! error value.

POWER USERS' CLINIC

Cumulative Interest and Principal Payments

The PPMT() and IPMT() functions let you analyze how much of a single payment pays a loan's principal, as well as the amount that pays interest. But what if you want to calculate the principal or interest payments over a *range* of payments? You can build an amortization table, like the one shown in Figure 10-3, or you can opt for a simpler choice: the *CUMIPMT()* and *CUMPRINC()* functions (which stand for cumulative interest payments and cumulative principal, respectively).

The CUMIPMT() function calculates the amount of money paid toward interest over a range of payments. Just specify the first and last payment you want to consider:

```
CUMPIPMT(rate, nper, pv, first_per,
last_per, type)
```

The following formula calculates the amount of interest paid over the first year (payments 1 to 12) of a $300,000 mortgage with a 25-year term.

```
=CUMIPMT(6%/12, 25*12, 300000, 1, 12, 0)
```

The CUMPRINC() function performs a similar, but complimentary task. It has the same syntax, but instead calculates the amount of money paid toward the *principal* over a range of payments.

NPER(): Figuring Out How Much Time You'll Need to Pay Off a Loan or Meet an Investment Target

The NPER() function calculates the amount of time it will take you to pay off a loan or meet an investment target, provided you already know the initial balance, the interest rate, and the amount you're prepared to contribute for each payment. Here's what the function looks like:

```
NPER(rate, pmt, pv, [fv], [type])
```

If you're ready to contribute $150 a month into a savings account that pays 3.5 percent interest, you can use the following formula to determine how long it'll take to afford a new $4,500 plasma television, assuming you start off with an initial balance of $500:

```
=NPER(3.5%/12, -150, -500, 4500)
```

The answer is 25.48 payment periods. Remember, a payment period in this example is one month, so you need to save for over two years.

A similar calculation can tell you how long it'll take to pay off a line of credit. Assuming the line of credit's $10,000 at 6 percent, and you pay $500 monthly, here's the formula you'd use:

```
=NPER(6%/12, -500, 10000, 0)
```

In this case, the news isn't so good: It'll take 21 months before you're rid of your debt.

RATE(): Figuring the Interest Rate You Need to Achieve Future Value

The RATE() function determines the interest rate you need to achieve a certain future value, given an initial balance, and a set value for regular contributions. The function looks like this:

```
RATE(nper, pmt, pv, [fv], [type], [guess])
```

The math underlying the RATE() function is trickier than the calculations used in the other financial functions. In fact, there's no direct way to determine the interest rate if there's more than one payment made. Instead, Excel uses an *iterative approach* (otherwise known as *trial and error*). In most cases, Excel can quickly spot the answer, but if it comes up empty after 20 iterations, the formula fails and returns the dreaded #NUM! error code.

You can always supply an optional *guess* argument, which tells Excel what interest rate to try first. If you don't specify a rate, then Excel assumes 10 percent, and works from there. The guess must use the appropriate time scale, so if you're using monthly payments, remember to divide the yearly interest rate by 12. The number returned by the RATE() function also uses the same time scale.

Tip: It's all too easy to unintentionally ask an impossible question using the RATE() function. So make sure you're using realistic numbers when employing this function. Say you want to find out what interest rate you need to pay off a $10,000 loan in two years by making monthly payments of $225. If you attempt this calculation, then you'll get a negative value, indicating that there's no way to meet your goal unless the loan pays *you* interest.

You could have avoided this mix up if you calculated the total amount of your monthly contributions ($225*24), which add up to only $5,625.00. Clearly, this amount of money isn't enough to even pay off the $10,000 principal, let alone the interest!

Imagine you have a starting balance of $10,000. You make regular payments of $150, and you're hoping to double your money to $20,000 in three years. To determine the interest rate you need to make this a reality, use the following formula:

```
=RATE(3*12, -150, -10000, 20000)
```

This calculation returns a monthly interest rate of 0.88 percent. You can generate the *annual* rate from this number by multiplying it by 12 (which gives you 10.5 percent, compounded monthly).

RATE() lends itself similarly well to loan calculations. Say you want to pay off a $5,000 loan in two years by making $225 payments each month. Here's the formula that determines the maximum annual rate you can afford.

```
=RATE(2*12, -225, 5000, 0)*12
```

In this case, you need to find a loan that charges less than 7.5 percent annual interest.

NPV() and IRR(): Net Present Value and Internal Rate of Return

NPV() is a more specialized function that can help you decide whether to make an investment or embark on a business venture by calculating *net present value*. To understand net present value, you first need to understand the concept of *present value,* which is the value that a projected investment must*day*. If you have an investment that earns 5 percent monthly interest and is worth $200 at maturity (after one month), its present value is $190.48.

You tackled present values with the PV() function described earlier in this chapter. The net present value's the same concept, except it applies to a series of *cash flows* (the profit, or loss, generated by an investment), rather than an investment with a fixed interest rate.

Practically speaking, people almost always use the NPV() function to compare the present value of an investment or business (sometimes called the *venture investment*) to an investment with a fixed rate of return. The basic idea is simple: In order to be worthwhile, the venture investment must exceed a specific rate of return. If you know that you can get a 5 percent fixed rate of return from your bank, you won't consider using the money to open a coffee shop that's projected to make only 3 percent annually. (As always, Excel's financial functions take only cold, hard money into account—if opening a coffee shop is a lifelong dream, you may think differently.)

The NPV() function makes its calculation by examining cash flow over a series of years. You specify how much the venture investment cost you, or how much you received during each period, and NPV() calculates how that compares with an investment with a fixed rate. To conduct this comparison, you must choose an interest rate for the venture investment you're comparing against the fixed investment. This investment rate is called the *discount rate*. If the final NPV() value is negative, you'd have been better off with the fixed investment. If the NPV() value is positive, your venture is exceeding the fixed investment option.

Imagine you want to start a new coffee shop business. You're prepared to invest $25,000 to start off, and you expect to realize profits of $2,000 in the first year. In the following years, your profit projections are $6,500, $10,000, and $12,500. Simple addition tells you that this business will earn $31,000 in four years (and a net profit of $6,000). However, you also need to account for the interest you could receive by investing the cash profits at 5 percent as soon as they become available. Based on this scenario, how would your business (the venture investment) compare to a fixed security that earns 5 percent each year? You can answer that question using NPV().

The NPV() function (see Figure 10-4) requires the discount rate, and then a series of payments (or a cell range that contains a series of payments). Here's how to calculate the net present value of the four-year investment:

```
=NPV(5%, 2000, 6500, 10000, 12500)
```

Excel counts each cash flow as one period. Thus, this formula goes through four periods. Because this formula uses an annual interest rate, each period represents a full year. The formula returns a result of $26,722.61. In other words, in order to generate the same amount of money as your business will, you'd need to invest $26,722.61 initially at an annual interest rate of 5 percent. This amount's more than the outlay required by our hypothetical business (which needs only $25,000 to get off the ground), so the business is a better place to put your money.

Or to put it another way, the present value of the business *today* is $26,722.61, even though you haven't ordered a sign or chosen a name. Because it costs only $25,000 to build the business, you'll be ahead of the game as soon as you set up shop (assuming, of course, you meet your profit projections).

Another way to look at net present values is to calculate the total money you've gained by subtracting the cost of your initial investment, as shown here:

```
=NPV(5%, 2000, 6500, 10000, 12500) - 25000
```

In this case, as long as the result's greater than $0, it's a go. In this example, the difference is $1,722.61. The business is worth this much today, over and above the outlay required to get it started.

Note: In this example, the business made money each and every year. However, you could still use NPV() even if some years show a negative cash flow (although this won't help your rate of return).

TROUBLESHOOTING MOMENT

Subtle Problems with NPV()

No one ever said that financial accounting is easy. In fact, the NPV() function is a good example of accounting's limitations, as it's littered with caveats that can easily trip you up.

For instance, you'll see that in order to calculate the real net present value, you need to subtract the initial investment cost after you use the function. You might wonder why you can't include a cost of $25,000 as one of the cash flows for NPV(), like this:

```
=NPV(5%, -25000, 2000, 6500, 10000, 12500)
```

This formula assumes that the $25,000 cost happens *after* the first period, which is incorrect. Also, the NPV() function now assumes that five years have elapsed, instead of four, further skewing the answer. These problems happen because NPV() is sensitive to the *time value* of money. Time value is an expression economists use to emphasize the fact that $95 dollars today isn't worth the same as $95 next year—because the earlier you get the money, the more time you must earn interest. Adding arguments or changing their order can cause a different result.

Another problem with NPV() is that it assumes you're able to reinvest your business's earnings at a rate comparable to the set interest rate you're using to calculate NPV (in this case, 5 percent). However, if you're unable to invest the profit at 5 percent (perhaps because you need a shorter term investment that earns less), then you're not generating as much profit as the NPV() function indicates.

The NPV() function is particularly useful for evaluating prospective real estate investments. Consider what happens if you buy a condominium, rent it out for several years, and then sell it for the original price. If you take the rent money you earn (or predict you're going to earn), the NPV() function can tell you whether you'd be better off investing your money at a fixed interest rate. Figure 10-4 shows an example.

In Figure 10-4, NPV() suggests that the rent you stand to earn (a total of $62,448.48, calculated by summing the rent collected in each year) is worth about $53,569 today. But you need to factor in the price of the property, which may increase or decrease. In this case, the property's estimated selling price ($160,000, which represents a measly gain of $10,000) makes it equivalent to having $125,364.19 to invest today at 5 percent interest. Thus, you're losing money on the value of the property (versus the rate of return you could get in the bank), but the rent you're earning more than makes up for it. By adding these two factors (property value and rent), you get a total net present value of $179,333.71. And because you purchased the property with only $150,000, this is very good news. If you still have doubts, look at the internal rate of return in cell B11. It shows that you're earning the equivalent of 8.08 percent—significantly more than 5 percent—from your real estate dealings.

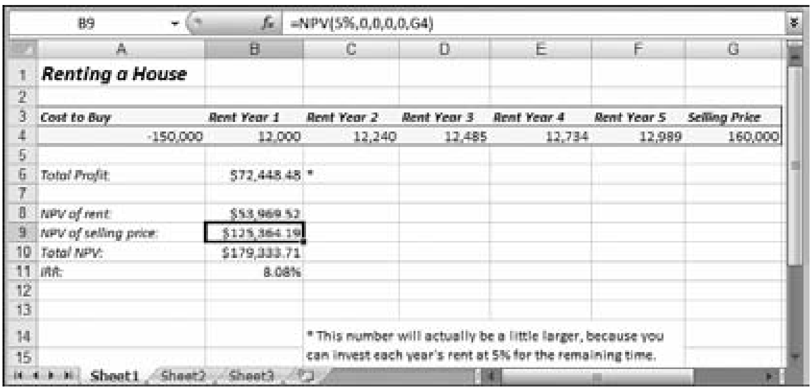

Figure 10-4: *This worksheet helps figure out whether it makes more sense to invest $150,000 at a fixed rate, or to use that money to buy and rent out a property.*

The IRR() function is closely related to NPV(). The main difference is that while NPV() calculates the value of the business or investment today, IRR() calculates how fast the business or rate *appreciates* in value (its rate of return). Technically, IRR() calculates the *internal rate of return,* which is the effective return rate based on your cash flows. In order to make this calculation, you must include the initial investment number with the cash flow. In fact, you need to supply all the values as a range, like so:

```
=IRR(A4:G4)
```

This calculation returns 7.35 percent for the coffee shop business introduced on page 306, and 8.08 percent in the rental property scenario shown in Figure 10-4.

Warning: The IRR() function is dangerous! While the NPV() function assumes you can reinvest your money at the rate you've specified (in the first argument), the IRR() function assumes that you can reinvest your money in the business for the same rate of return, which often isn't possible. Both NPV() and IRR() calculations are just estimates, and you need to carefully evaluate them.

Depreciation

Another common calculation in the world of finance is *depreciation*. Simply stated, depreciation is how much the value of an asset decreases over time. All assets that depreciate begin with a certain value (which you determine) and then depreciate over the course of a lifetime (which you specify). At the end of an asset's life, from an accounting perspective, the asset is deemed to be useless and without value.

POWER USERS' CLINIC

A Better NPV()

Excel also provides versions of NPV() and IRR() that don't assume that payments are made (or profits realized) on fixed dates, once per period. These functions are XIRR() and XNPV(). Using them is a little bit more complicated because you need to supply not only the cash flows, but also the corresponding dates when each transaction happens. You do this by submitting two cell ranges: the first range has the cash flows, and the second range has the dates.

```
XNPV(rate, cashflow_range, date_range)
XIRR(cashflow_range, date_range)
```

The two ranges must be exactly the same size (that is, they must include the same number of cells), and must line up logically. In other words, Excel assumes that the first cash flow in the cash flow range is deposited on the first date in the date range.

Excel has four basic depreciation functions (explained below), which can help you determine how much value your asset has lost at any given point in time. You'll use these values if you want to sell the asset, or if you're attempting to calculate the current net worth of a business. Depreciation can also figure into tax calculations and business losses; a company could write off a loss based on the value an asset has lost.

Note: The depreciation of an asset isn't as straightforward as interest calculations. You use a certain amount of guesswork when you're deciding what an asset is worth and how rapidly its value declines. An asset can include almost anything, from a piece of equipment or property to a patented technology. It may depreciate due to wear and tear, obsolescence, or market conditions (like decreased demand and increased supply).

In order to assign a value to a depreciated asset, Excel makes a logical guess about the way in which the asset's depreciating. You can use any of a number of accepted ways to make this guess (or estimate, as financial types prefer to call it). The easiest way of depreciating an asset is the *straight-line depreciation* method, where the value of the asset decreases regularly from the starting value to the final value. However, this approach isn't necessarily realistic for all assets.

Excel supports four basic depreciation functions, each of which figures depreciation in a slightly different way:

- **SLN()** uses simple *straight-line depreciation,* where the cost of the asset minus its *scrap value* (the value of the asset if sold purely as raw material) is simply divided by the life span of the asset. In other words, if the life span's 10 years, the book value of the asset is depreciated by 10 percent of the original value each year. SLN() is the only form of depreciation that proceeds regularly. All other types of depreciation are *accelerated depreciation* functions, because they assume that the asset's greatest loss in value happens early on, rather than evenly over several years. As a result, these accelerated depreciation functions are often more realistic.
- **SYD()** uses the *sum-of-years-digits* depreciation method. This method starts with a larger depreciation rate, which gradually lowers as the asset becomes less valuable. SYD() is a good all-around choice for most depreciation calculations.
- **DDB()** uses the *double declining balance* depreciation method, which is like straight-line depreciation on steroids. It reduces an asset by double the percentage of the SLN() method, which makes for a fast reduction in value (and a hefty tax write-off). In other words, if the life span is 10 years, the book value of the asset depreciates by 20 percent of the original value each year.

Note: This approach has a problem, however: Eventually the asset will dip below its final value (and even become a negative number). Accountants handle this problem by imposing an additional rule. As soon as the amount of depreciation in any given year is *less* than the straight-line depreciation amount, you need to immediately abandon the DDB() method and start calculating depreciation with SLN(). But it's still up to you to be aware of this and check for this condition. If this sounds messy—well, it is.

- **DB()** uses the *declining balance* method, where an asset declines in value by a fixed percentage of its book value every year. Like the SYD() and the DDB() approaches, this method is an accelerated depreciation function. However, it has a different limitation: You can't specify the final value the asset should have after its lifetime's over (the salvage value).

These functions work with more of less the same arguments, but they produce different results. Here's a look at the syntax of all four functions:

```
SLN(cost, salvage, life)
SYD(cost, salvage, life, period)
DB(cost, salvage, life, period, [month])
DDB(cost, salvage, life, period, [factor])
```

- The *cost* is the initial value of the asset.
- The *salvage* argument is the value of the asset at the end of its depreciation. If it's lost all value, this is 0.
- The *life* is the number of periods over which the asset is depreciating. It's also known as the *useful life* of the asset.

- The *period* is a number that indicates the period in which you want to calculate depreciation; it must use the same unit as the life. If you're depreciating an asset over three years, a period of 3 will get you the amount of money that the asset depreciated in the third year. The SLN() function doesn't require a period argument because depreciation is the same in every period.
- The *month* is an optional argument supported by the DB() function. It specifies the number of months that the asset's in use for the first year. If you leave this out, then Excel assumes 12 months of service.
- The *factor* is an optional argument supported by the DDB() function. You multiply the factor by the expected straight-line depreciation to get the actual depreciation. If you omit this argument, then Excel assumes a value of 2 (which is why it's known as *double* declining depreciation).

As an example, imagine a company purchases a top-of-the-line computer for $10,000. The computer has a salvage value of $200 after its useful life of five years. To compute its second-year depreciation using the popular SYD() method, you'd use the following calculation:

```
=SYD(10000, 200, 5, 2)
```

This calculation indicates that the computer loses $2,613.33 of its value in the second year. However, to calculate the total depreciation so far, you need to add the depreciation in both the first and second year, and subtract that from the initial value, as shown here:

```
=10000-SYD(10000, 200, 5, 1)-SYD(10000, 200, 5, 2)
```

This formula takes the $10,000 total, and subtracts the first- and second-year depreciation ($3,266.67 and $2,613.33, respectively), which gives you a current value of $4,120. You'll notice that with the SYD() function, most of the depreciation takes place in the first year, and then each following year the asset loses a smaller amount of its value.

The easiest way to see depreciation at work is to build a worksheet that compares these four different functions. Figure 10-5 shows an example.

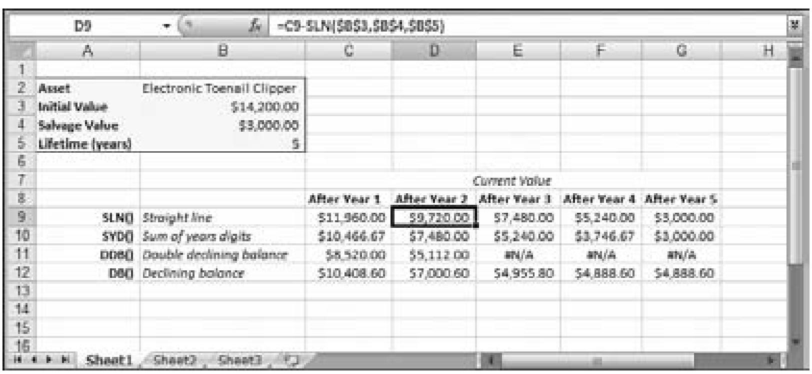

Figure 10-5: *This worksheet shows depreciation on a sample investment. Once you set the initial value, lifetime value, and salvage value, you can calculate the depreciation. But as you can see, some methods produce faster depreciation than others.*

In Figure 10-5, the straight-line depreciation is the slowest of all, which makes it less realistic for dealing with the plummeting values of aging high-tech parts. The sum-of-years-digits depreciation provides faster depreciation, yet still ends up with the desired salvage value of $3,000. The last two methods are a little trickier. You can use the double-declining balance approach only for the first two years. In the third year, the depreciation it calculates ($2,044.8) is less than the sum-of-years depreciation value ($2,240), so it's time to switch methods. To underscore that fact, the following cells display the #N/A code. The declining-balance approach works well initially, but then tops out in the last year. Probably, you'll ignore the last calculated value and just use the salvage value of $3,000.

Note: When using double-declining balance depreciation, it's up to you to spot when the method no longer applies. Excel doesn't tell you; instead, it cheerily continues calculating depreciation values, even if the depreciation exceeds the total value of the asset. In Figure 10-5, the #N/A values were typed in by hand.

Other Financial Functions

Excel includes dozens more financial functions. People like stockbrokers, bond traders, and chief financial officers, who regularly depend on the functions' ability to deal with complex financial transaction procedures, use most of these. Table 10-1 provides a quick overview.

Tip: For more info on non-Excel-specific investing, depreciation, and financial concepts, check out *http://beginnersinvest.about.com/od/investinglessons*.

Table 10-1. *Specialized financial functions*

Function	Description
ACCRINT()	Returns the accrued interest for a security that pays periodic interest.
ACCRINTM()	Returns the accrued interest for a security that pays interest at maturity.
AMORDEGRC()	Returns the depreciation for each accounting period.
COUPDAYBS()	Returns the number of days from the beginning of the coupon period to the settlement date.
COUPDAYS()	Returns the number of days in the coupon period that contains the settlement date.
COUPDAYSNC()	Returns the number of days from the settlement date to the next coupon date.
COUPNCD()	Returns the next coupon date after the settlement date.
COUPNUM()	Returns the number of coupons payable between the settlement date and maturity date.
COUPPCD()	Returns the previous coupon date before the settlement date.
DISC()	Returns the discount rate for a security.

DURATION()	Returns the annual duration of a security with periodic interest payments.
FVSCHEDULE()	Returns the future value of an initial principal after applying a series of compound interest rates.
INTRATE()	Returns the interest rate for a fully invested security.
MDURATION()	Returns the Macaulay modified duration for a security with an assumed par value of $100.
MIRR()	Returns the internal rate of return where positive and negative cash flows are financed at different rates.
NOMINAL()	Returns the annual nominal interest rate.
ODDFPRICE()	Returns the price per $100 face value of a security with an odd first period.
ODDFYIELD()	Returns the yield of a security with an odd first period.
ODDLPRICE()	Returns the price per $100 face value of a security with an odd last period.
ODDLYIELD()	Returns the yield of a security with an odd last period.
PRICE()	Returns the price per $100 face value of a security that pays periodic interest.
PRICEDISC()	Returns the price per $100 face value of a discounted security.
PRICEMAT()	Returns the price per $100 face value of a security that pays interest at maturity.
RECEIVED()	Returns the amount received at maturity for a fully invested security.
TBILLEQ()	Returns the bond-equivalent yield for a Treasury bill.
TBILLPRICE()	Returns the price per $100 face value for a Treasury bill.
TBILLYIELD()	Returns the yield for a Treasury bill.
VDB()	Returns the depreciation of an asset for a specified or partial period using a declining balance method.
YIELD()	Returns the yield on a security that pays periodic interest.
YIELDDISC()	Returns the annual yield for a discounted security, like a Treasury bill.
YIELDMAT()	Returns the annual yield of a security that pays interest at maturity.

CHAPTER

11

Manipulating Dates, Times, and Text

Most of the functions you've learned so far focus on crunching raw numbers. However, Excel also provides functions that work with *other* types of data, including dates, times, and ordinary text. For example, you may want to pull first and last names from different columns and join them together in a single column. Or, you may want to remove a word that appears in a bunch of column titles. Similarly, you may want to replace a character in a word, capitalize a name, or count the number of letters in a cell. Excel provides specialized functions for all these tasks, and you'll learn about them in this chapter.

Dates and times also have their own specialized functions. These functions perform some indispensable tasks, like retrieving the current time and determining what day of the week a given date falls on. In addition, Excel lets you perform calculations with dates and times just as you would with ordinary numbers. This chapter introduces these techniques, and explains how Excel stores dates and times behind the scenes.

Manipulating Text

You can't use arithmetic operators like + and – with text. If you try to perform a calculation by referring to one or more cells containing text, Excel displays a *#VALUE* error message. However, there's one operator you can use: the *concatenation* operator (&), which joins together text. For example, imagine you have an individual's first name in cell A1, and a last name in cell A2. You could join the values from these two cells to create a full name with this formula:

```
=A1 & A2
```

This approach has one drawback. In all likelihood, the first and last name cells don't include any leading or trailing spaces. This trait means when you join them, they'll get fused into one word, like JohnSmith. One solution is to explicitly add a space (between quotation marks) into your formula, like so:

```
=A1 & " " & A2
```

The important concept in this example is that you can enter *string literals*—fixed pieces of text (including spaces)—as easily as you can enter literal numbers. The only difference between entering literal text and literal numbers is that you have to place all text between quotation marks. You can stitch together as many pieces of text together at a time as you want; there's no limit. The next group of functions showcases the many different ways that Excel lets you manipulate text.

Tip: Concatenation also works with cells that contain numbers. In these cases, the text is simply the cell content formatted with the General number format, no matter what formatting you've used in the cell. For example, even if you've formatted the number 43.2 so it appears as the currency value $43.20, when you join this number to a piece of text, it automatically reverts to the ordinary 43.2. This often isn't what you want, particularly if the cell you want to use contains date information, which Excel displays as a serial number in the General number format. (For more on how number formats affect the appearance of dates, see page 70.) To avoid these problems, you can use the TEXT() function described on page 321.

CONCATENATE(): Joining Strings of Text Together

The CONCATENATE() function lets you join text in exactly the same way the & operator does. CONCATENATE() joins all the parameters that you supply into one long piece of text, in the order you specify them.

Here's an equivalent way to rewrite the aforementioned name-joining formula by using CONCATENATE() with two pieces of text:

```
=CONCATENATE(A1, " ", A2)
```

LEFT(), MID(), and RIGHT(): Copying Portions of a Text String

Just as you can join pieces of text, you can also split up a string of text. The LEFT(), MID(), and RIGHT() functions let you extract a portion from a larger text string. For example, the LEFT() function takes two arguments: the text you want to examine, and the number of characters that Excel should extract, starting from the string's left side:

```
LEFT(text, num_characters)
```

To take the first four letters from the text in cell A1, you'd use the formula:

```
=LEFT(A1, 4)
```

Assuming the cell contains the text *tofurkey*, the result of this formula would be the shortened text *tofu*.

The RIGHT() function performs the same operation, but it extracts letters starting from the right side of a string. For example, consider the following formula:

```
=RIGHT(A1, 5)
```

If you use this function with the same string, you'll simply end up with the text *urkey*.

The MID() function is more powerful than the LEFT() and RIGHT() functions, as it has the ability to extract a consecutive series of characters from anywhere inside a string. When using the MID() function, you need to supply *three* arguments: the text you're evaluating, the starting position, and the number of characters. Excel numbers each letter in a string (it also counts spaces), starting with 1 for the first letter, 2 for the second letter, and so on. That means if you specify a starting position of 3 and a length of 2, Excel extracts the third and fourth characters. The basic formula looks like this:

```
MID(text, start_position, number_of_characters)
```

Here's an example that copies characters from the middle of a string. If the cell A1 contains the text *Swanky Franks*, the following formula returns the value Frank.

```
=MID(A1, 8, 5)
```

Note: LEFT(), MID(), and RIGHT() all pluck out the strings you specify, but leave the original contents of the cell you're examining unchanged.

LEN(), FIND(), and SEARCH(): Counting Characters in a String

The LEFT(), RIGHT(), and MID() functions let you copy specified segments of a string. But what happens if you don't know the exact length of the string you're searching? For example, if you're interested in retrieving last names from a column that contains full names, none of these functions help you since last names vary in length.

Fortunately, Excel provides some other tools that can help you out. Three of these, LEN(), FIND(), and SEARCH(), give you numeric information about text. This section first explains how to use each of these functions, and then it shows how you can combine them with the LEFT(), RIGHT(), and MID() functions to perform some really powerful operations.

To begin with, LEN()—short for LENgth—counts the number of characters in a string of text. For example, the result of the following formula is 5:

```
=LEN("Hello")
```

The FIND() function is more sophisticated. It gives you a number representing the position of a given character or series of characters. For example, the function can tell you where a space is located in a phrase. If there's more than one match, the FIND() function returns the position only of the first match. The FIND() function can also take an optional third parameter indicating the position where the search should begin—if left out, Excel starts at the beginning of the text. Here's what the basic formula looks like:

```
FIND(find_text, within_text, [start_position])
```

Now consider the following example. It gives you 5, indicating that the first space in the phrase "Mind the Gap" is in position 5.

```
=FIND(" ", "Mind the Gap")
```

SEARCH() works in almost exactly the same way, and it takes the same two or three arguments. The only difference is that FIND() performs a *case-sensitive* search, which means it looks at whether upper- and lowercase letters match, whereas SEARCH() doesn't care about inconsistencies between the case of the text you're looking for and the text you're evaluating.

The LEN(), FIND(), and SEARCH() functions all become even more useful when used in conjunction with the LEFT(), RIGHT(), and MID() functions. For example, say you've got a column filled with full names, as shown in Figure 11-1, and you want to copy the first and last names and paste them into their own new columns (called "First Name" and "Last Name"). To copy the first name from a cell containing a full name, here's what to do.

Although you don't know how long the first name is, you can search for the position of the first space by using the FIND() function. Next, subtract one from the number you get from the FIND() function. You'll now end up with the length you need to copy the first name (using the LEFT() function). Here's a formula that puts all these steps together:

```
=LEFT(A2,FIND(" ", A2)-1)
```

This formula's used for cell B2 in Figure 11-1.

You can use a similar trick to get the last name using the RIGHT() function:

```
=RIGHT(A2,LEN(A2)-FIND(" ", A2))
```

This formula is also put to work in Figure 11-1—it's in cell C2. At first glance it looks a bit overwhelming, so it helps to remember the two arguments that the RIGHT() function is processing: first, which cell to evaluate (A2) and, second, the number of spaces over that it should count from the right (LEN(A2)–FIND(" ", A2)). The first argument is straightforward. Here's how the second argument breaks down: First, the FIND() function finds the location of the first space (this is the space separating the first and last name). Next, the LEN() function determines how many characters (including spaces) are in the full name string. Excel then subtracts the FIND() result from the LEN() result, which helps identify exactly where the last name begins. Then the RIGHT() function can use this number to do its job: copying out the last name.

Tip: The previous example uses a custom formula to split text. Excel also has another feature designed for just this problem–the Text to Columns feature. Simply select the cells you want to change, and then choose Data→Data Tools→Text to Columns. Excel shows the Convert Text to Columns Wizard, which lets you choose how you want to carve your text up into separate columns (either by fixed position, or by using a recognized delimiter character, like a space or a comma). The Convert Text to Columns Wizard is based on the Text Import Wizard, which you'll explore in detail on page 787.

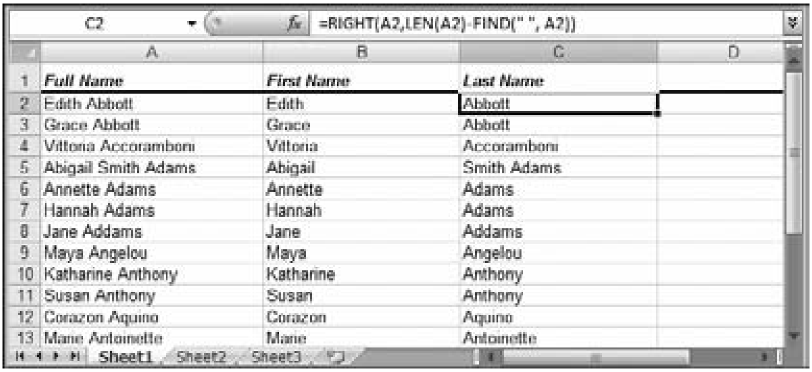

C2 =RIGHT(A2,LEN(A2)-FIND(" ", A2))

	A	B	C	D
1	*Full Name*	*First Name*	*Last Name*	
2	Edith Abbott	Edith	Abbott	
3	Grace Abbott	Grace	Abbott	
4	Vittoria Accoramboni	Vittoria	Accoramboni	
5	Abigail Smith Adams	Abigail	Smith Adams	
6	Annette Adams	Annette	Adams	
7	Hannah Adams	Hannah	Adams	
8	Jane Addams	Jane	Addams	
9	Maya Angelou	Maya	Angelou	
10	Katharine Anthony	Katharine	Anthony	
11	Susan Anthony	Susan	Anthony	
12	Corazon Aquino	Corazon	Aquino	
13	Marie Antoinette	Marie	Antoinette	

Sheet1 Sheet2 Sheet3

Figure 11-1: *This worksheet shows the LEFT(), RIGHT(), FIND(), and LEN() functions at work. Excel uses these functions to extract the first and last names from the full names in column A.*

UPPER(), LOWER(), and PROPER(): Changing Capitalization

Another way to manipulate text is by changing capitalization. Excel provides three functions for this purpose (see Figure 11-2):

- **UPPER()** converts text to all capitals.
- **LOWER()** converts text to all lowercase.
- **PROPER()** converts text to initial-case. That means every letter is lowercase, except for the first letter in each word. (Excel identifies where words start by looking for the spaces in between.)

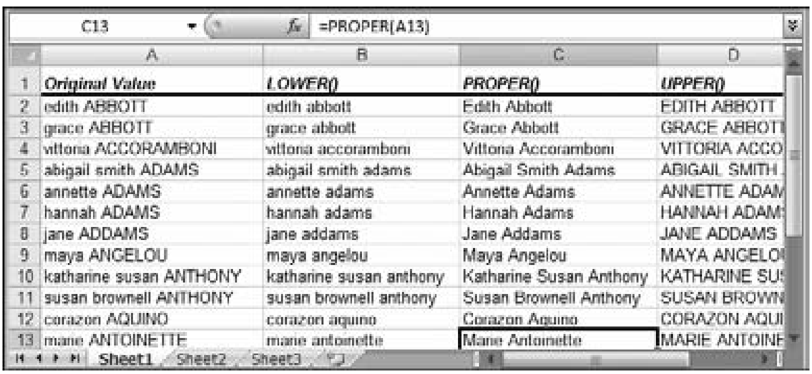

C13 =PROPER(A13)

	A	B	C	D
1	*Original Value*	*LOWER()*	*PROPER()*	*UPPER()*
2	edith ABBOTT	edith abbott	Edith Abbott	EDITH ABBOTT
3	grace ABBOTT	grace abbott	Grace Abbott	GRACE ABBOTT
4	vittoria ACCORAMBONI	vittoria accoramboni	Vittoria Accoramboni	VITTORIA ACCO
5	abigail smith ADAMS	abigail smith adams	Abigail Smith Adams	ABIGAIL SMITH
6	annette ADAMS	annette adams	Annette Adams	ANNETTE ADAM
7	hannah ADAMS	hannah adams	Hannah Adams	HANNAH ADAM
8	jane ADDAMS	jane addams	Jane Addams	JANE ADDAMS
9	maya ANGELOU	maya angelou	Maya Angelou	MAYA ANGELO
10	katharine susan ANTHONY	katharine susan anthony	Katharine Susan Anthony	KATHARINE SU
11	susan brownell ANTHONY	susan brownell anthony	Susan Brownell Anthony	SUSAN BROWN
12	corazon AQUINO	corazon aquino	Corazon Aquino	CORAZON AQUI
13	marie ANTOINETTE	marie antoinette	Marie Antoinette	MARIE ANTOINE

Sheet1 Sheet2 Sheet3

Figure 11-2: *This worksheet shows how the UPPER(), LOWER(), and PROPER() functions change the capitalization of the names in column A.*

All three of these functions need just one part, which is a string (a short bit of text, anywhere from a few letters to a lengthy phrase). You can use a reference to a cell that contains text, or a piece of literal text—as long as you remember the quotation marks. For example, the following formula displays the text contained in cell A1, but changes all characters to lowercase:

```
=LOWER(A1)
```

GEM IN THE ROUGH

Using PROPER() to Change Names in All Caps to Initial Caps

Usually, functions like UPPER(), LOWER(), and PROPER() transform the appearance of text that's already in your spreadsheet. But you could have a tricky time figuring out exactly how to do this transformation.

Say you've got a list of first and last names in column A (beginning in cell A1) whose letters are all uppercase. You want to change these names so that only the first letter of each name's capitalized. Here's what to do. First, insert a new column B. Next, enter the formula *=PROPER(A1)* in cell B1, and then copy this formula to the rest of the cells in column B. Because this formula uses a *relative* cell reference (as explained on page 252), Excel automatically adjusts the formula for each cell that you paste it into.

Once you've taken this step, Excel displays the properly capitalized names in column B. But you can't just copy and paste these cells into column A (because they hold formulas that reference the cells you'd be pasting them into). Instead, you need to select the corrected names from column B, copy them (Home→Clipboard→Copy, move to cell A1, and then select Home→Clipboard→Paste→Values (not the plain-vanilla Home→Clipboard→Paste command).

This technique provides a quick way to clean up a number of problems. For example, you can also use this approach with many other text functions like TRIM(), CLEAN(), SUBSTITUTE(), TEXT(), FIXED(), DOLLAR(), all of which are explained in this chapter.

TRIM() and CLEAN(): Removing Unwanted Spaces and Non-Printing Characters

The TRIM() and CLEAN() functions perform minor cleanup on any strings of text you run through them. TRIM() removes any leading and trailing spaces; it also changes any series of more than one space to a single space. Thus, if you use TRIM() on the text string " Hello There " the altered text becomes "Hello There." TRIM() can be quite handy for fixing erratic spacing.

CLEAN() simply removes non-printable characters from a text string. Non-printable characters, which usually appear as empty-box icons in your text, tend to appear only if you import some text from another file format that Excel has difficulty understanding.

SUBSTITUTE(): Replacing One Sequence of Characters with Another

The SUBSTITUTE() function replaces a sequence of characters in a string with another set of characters. The function has three parts: the text you want to modify, the characters you're looking to replace, and the replacement text you want to insert. In addition, you can supply an optional *occurrence number* parameter, which Excel uses if it finds more than one match. For example, if your search text's matched three times, and you supply 2 for your occurrence number, Excel changes only the second occurrence. If you don't supply the occurrence number, Excel changes all occurrences. Here's what the function looks like:

```
SUBSTITUTE(text, old_text, new_text, [occurrence_number])
```

Consider the case where cell A1 contains the text: *It was the best of times; it was the worst of times*. You could use the following formula:

```
=SUBSTITUTE(A1, "times", "nanoseconds")
```

The result is the string: *It was the best of nanoseconds; it was the worst of nanoseconds*.

On the other hand, the following formula explicitly replaces just the second occurrence. The resulting string is: *It was the best of times; it was the worst of crimes*.

```
=SUBSTITUTE(A1, "times", "crimes", 2)
```

Note: The SUBSTITUTE() function always performs a case-sensitive search. That means if you try using SUBSTITUTE() to replace the word *it*, in the previous example, Excel won't match *It*.

TEXT(), VALUE(), FIXED(), and DOLLAR(): Converting Text to Numbers and Vice Versa

Sometimes, you may need to convert text into a number, or vice versa. For example, imagine you have a cell that contains the sentence "A good sandwich costs $5.95." Using the MID() function, you could copy just the part of this text that has the price—"5.95." However, even though this text contains numeric information, to Excel it's still a piece of text, so you can't perform mathematical operations with it (like adding tax). On the other hand, you may have the reverse problem: You might want to show a string of text that includes a number from another cell. In these cases, the data conversion functions TEXT() and VALUE() are useful.

TEXT() converts an ordinary number into formatted text using the format that you specify. It always requires two arguments:

```
TEXT(number, format)
```

The first argument is the number you're converting; the second argument's the format you want to use. You can use any of the date, time, or numeric formatting codes described in Chapter 5 (in particular, Table 5-2 on page 155). For example, the following formula converts the number 434.2 to the formatted text $434.20.

```
=TEXT(434.2,"$#,##0.00")
```

On its own, this method may not seem very practical. After all, you've already learned how you can control cell formatting using the Format Cells dialog box. What's the point of doing the same job with a formula? The answer is that you may want to do so when you're looking to perform some fancy text processing with your number.

For example, imagine you have the price of a product in cell A1 (which happens to be 300). You could use the following formula to change a number to formatted text, and then put it into a complete sentence, like *$300.00 is way too expensive*.

```
=TEXT(A1,"$#,##0.00") & " is way too expensive."
```

Now compare what happens if you try to do the same thing without using the TEXT() function:

```
=A1 & " is way too expensive."
```

You'd end up with an unformatted result: *300 is way too expensive*. This outcome is true regardless of what formatting you've used in cell A1.

In other words, you can use text-based functions with numbers, but unless you explicitly perform the conversion using the TEXT() function, Excel always uses General formatting—which may not be what you want.

Tip: You can also use the FIXED() or DOLLAR() functions (explained below) to convert to specific formats without needing to find the right format string.

The VALUE() function performs the reverse transformation—it converts a piece of text into a number you can manipulate in a formula. Here's an example:

```
=VALUE(A1)
```

This trick becomes useful if you need to extract a price out of a string of text, and then perform a calculation with it. Here's an example that gets the number 300 out of a sentence using the MID() function, and then converts it to a number using the VALUE() function:

```
=VALUE(MID("I suggest a price of $300.00 for the chair.",23,6))
```

Note: In many cases, you can get by without the VALUE() function, because you can coax Excel into converting text into numbers. However, with some formulas, you need it–and it never hurts to make your formulas clearer by using it.

The VALUE() function is fairly simple to use, but it isn't terribly bright. If you use it with content that contains both numeric characters and letters, like the string of text *42 bananas*, it fails. You'll see the error message #VALUE! instead of the desired content.

The FIXED() and DOLLAR() functions also convert numbers to text. The difference is that these functions are customized to use a set format (described shortly), so you don't need to specify the desired format as you would with the TEXT() function.

The FIXED() function actually performs several steps. First it rounds a number to a specified number of decimal places. Next, it formats the number with a decimal point and, optionally, with commas to separate the thousands. Finally, it converts the number to text. Here's what the function looks like:

```
FIXED(number, [number_of_decimals], [no_commas])
```

You need only the first argument. If you don't specify the other arguments, the FIXED() function automatically uses two decimal places and includes commas if your number's large enough to warrant them. If you don't want commas, set the third argument to TRUE.

For example, the following formula gives you the text *5,450.59*:

```
=FIXED(5450.586, 2)
```

The DOLLAR() function automatically applies the currency format before it converts a value, so your text appears with a currency symbol before it. If you want, you can specify the number of decimal places, or you can just accept the default of two decimal places:

```
DOLLAR(number, [number_decimals])
```

Note: Because the FIXED() and DOLLAR() functions give you text, not numbers, you'll notice that cells bearing these functions are left-aligned (Excel's default alignment for all text). If this bothers you, you can explicitly change the cell alignment. You'll also notice that if you change the number format for any of these cells, it won't have any effect. That's because the cell contains text, not a number.

Other Text Functions

So far you've had a tour of the most useful text manipulation functions. Excel also provides a few lesser-used functions, outlined in Table 11-1.

Table 11-1. *Miscellaneous text functions*

Function	Syntax	Description
CHAR()	CHAR(number)	Returns the character for a specific character code. For example, CHAR(100) corresponds to the lowercase letter d.
CODE()	CODE(text)	Returns the numeric code for the first character in a text string. For example, CODE("d") returns 100.
EXACT()	EXACT(text1, text2)	Compares two pieces of text, and returns TRUE if they match, or FALSE if they don't. Usually, it's just easier to use the IF() function with the equal sign for comparison, as described in Chapter 13.
REPT()	REPT(text, number_of_times)	Creates a text string by repeating the text you specify the number of times you specify. This function's a quick way to repeat one or more characters in a cell.
T()	T(value)	T() ensures that certain content is in text form. If you use T() with a piece of text, it returns that piece of text. If you use T() with a number, it returns an empty string. Thus, =T("Hello") returns the text *Hello*, while =T(56) just returns a blank value.

N()	N(value)	N() ensures that certain content is a number. If you use N() with a numeric value, it returns that number. If you use N() with a text string, it returns 0, even if the text could be converted to a number. Thus, =N(2) returns the number 2, while =N("2") returns 0.

Manipulating Dates and Times

To understand how you can use Excel's functions to manipulate dates and times, you need to understand a little more about how Excel stores these values. The reality is that even though many dates look, at least partly, like text entries (for instance, *Mar. 5, 2010*), Excel actually stores them as *serial numbers*. Under this system, Excel designates the date January 1, 1900 as day 1; January 2, 1900 is day 2; and so on.

Thus, if you use Excel to subtract one date from another, you actually end up calculating the difference in days, which, it turns out, is exactly what you want. But this system of date storage leads to some interesting side effects. For example, you can't enter a date in Excel that's earlier than January 1, 1900. Try it—if you do, Excel treats your date just like text.

Excel also supports an alternate date system where it designates January 1, 1904 as day 1. This format is included to ensure compatibility with the Macintosh version of Excel, which uses that date system. To change your worksheet to use dates that start at 1904, select File→Options, and then choose the Advanced section. Scroll down the page of settings until you see the heading "When calculating this workbook". Underneath, turn on the checkbox next to "Use 1904 date system". Dates don't look any different on your worksheet, but their internal representations have changed, and you can't enter a date earlier than January 1, 1904, without Excel converting it to plain text.

Tip: The fact that the PC version of Excel uses a different date format than the Mac version of Excel won't cause a problem when you transfer files, because Excel's smart enough to adjust to the difference between the two formats. However, in a few rare cases, it could be trouble. First, if you cut-and-paste between Excel files that use different date systems, you could get some glitches. And if you use General format on your dates, and then transfer the files, you could wind up with errors. These cases aren't common—but if they affect you, be on the lookout.

Ordinarily, you won't see these underlying serial values, because Excel always displays dates using the Date number format. However, you can take a look at the underlying number by changing the format of any cell that contains a date. For example, type *1/1/2010* into a worksheet, right-click the cell, and then choose Format Cells. In the Category list, choose General instead of Date. Now you'll see the number 40179 instead of the text date. (Incidentally, you can also perform the reverse trick to display a normal number as a date, although doing this doesn't usually make much sense.) For more information about formatting date values, see Chapter 5.

Excel also stores times as numbers behind the scenes. In fact, Excel stores every time value internally as a fractional number from 0 to 1. The number 0 represents 12:00 AM (the start of the day) and 0.999 represents 11:59:59 PM (the end of the day). Because Excel stores times as a single number, it's easy to subtract one time value from another. However, time values can have varying degrees of precision. If your time's accurate down to the millisecond, it includes more decimal places (up to eight). For example, Excel stores 10:30 AM as 0.4375; whereas 10:30:32.34 is 0.437874306.

So now that you know dates are really whole numbers and times are fractions, what happens if you enter a number like *40179.50,* and then apply the Date number format? In this case, Excel uses the whole number part (40179) for the date, and the fractional part (0.50) for the time. Excel therefore interprets 40179.50 to mean 40,179 days from January 1, 1900. Thus, the resulting value is the combined date and time of 1/1/2010 12:00:00 PM (see Table 11-2). If you change the number to *40179.40*, you end up with 1/1/2010 9:36:00 AM. At this point, you're probably realizing that there's really no difference between dates and times—they're just different components of a single number.

Table 11-2. *The internal representation of dates and times*

What you see	What Excel sees
January 1, 2010	40179
12:05 PM	0.503472
January 1, 2010 12:05 PM	40179.503472

Math with Dates and Times

On Planet Excel, since dates and times are really just special types of numbers, you can use them in calculations like addition, subtraction, and so on. In the case of dates, the most common operation is to subtract one date from another to calculate the number of days in between. For example, consider this formula:

```
=A2-A1
```

If A2 contains the value 10/30/2010, and A1 contains the value 3/20/2010, the result is 224, which is the number of days between these two dates.

Using Dates and Times with Ordinary Functions

Of course, you can also use other Excel functions in your calculations. For example, you could take a number that represents a combined date and time value, extract just the fractional time portion, and then round it to a number of whole hours using the following formula:

```
=ROUND((A1-INT(A1))*24, 0)
```

This formula works on a date in cell A1. Here's how it breaks down. The formula begins by calculating *INT(A1)*. This function truncates the date number, removing the fractional time portion. Then, the formula evaluates *A1-INT(A1)*. Now you've got just the fractional time portion of the date. Next, multiply this fractional part by 24 (the number of hours in a day), and then round it to a whole number. The end result is that you extract the number of hours in a date/time value. Figure 11-3 shows some examples of date calculations similar to this one.

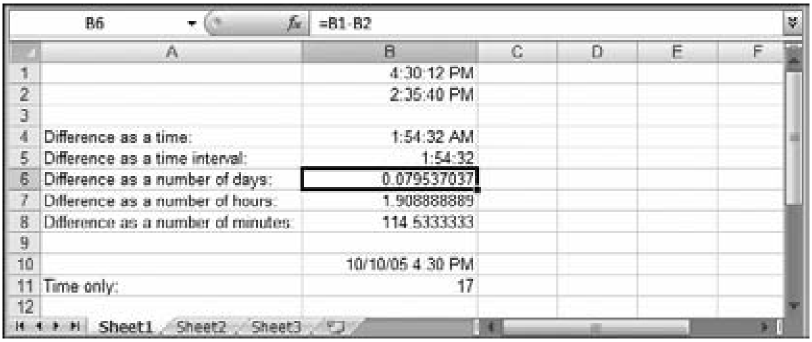

Figure 11-3: *This worksheet shows different ways to find the difference between two times.*

In Figure 11-3, the calculation shown in cell B4 (*=B1–B2*) uses the default format, which incorrectly interprets the time difference as a chronological time value. The second calculation (cell B5) uses the same formula but an interval time format. As a result, Excel displays the result as an interval of time. The next example (cell B6) uses the General format to display the result in terms of the number of fractional days. Finally, the last two examples take the additional step of multiplying the number of fractional days by 24 (to calculate fractional hours) and by 60 (to calculate fractional minutes).

As you'll see in the rest of this chapter, Excel has plenty more specialized date functions that often make calculations easier than this.

Note: Remember, if you change the number format, you may end up hiding part of the date's information. For example, if you choose to display a combined date and time value using a time-only format, you won't see the date information. However, Excel still uses this date information in any calculations you perform, which could lead to unexpected data. For best results, perform your date manipulation using the date functions described in the next section, rather than writing your own formulas from scratch.

Date and Time Functions

You aren't limited to arithmetic operators for manipulating dates and times. Excel also provides some invaluable ready-to-use functions to help you do things like identify the current date, extract just part of the date, and perform calculations with a date.

TODAY() and NOW(): Inserting the Current Date and Time

The TODAY() function automatically retrieves the current date. It doesn't require any arguments because it simply examines your computer's internal clock. The TODAY() function is extremely useful for creating spreadsheets that continuously update themselves (sometimes called *dynamic spreadsheets*). For example, you could create a formula that determines the number of days a payment is overdue. Here's an example:

```
=TODAY()-A1
```

This formula assumes that cell A1 contains the date a payment was due. As a result, it calculates the number of days between the two dates, which shows up as an ordinary number (like 14 for a payment that's two weeks late). Remember, you'll need to display the result as an ordinary number (representing the number of days), not as a date.

The NOW() function is similar to the TODAY() function, except it retrieves the current date along with the current time. If you use NOW() to display a value in a cell that doesn't have any special formatting applied, Excel uses a custom format that shows the date and time (listed in the 24-hour format; for example, 10/5/2010 19:06).

You can use other formats to hide some of this information. For example, a cell you've formatted using the custom number format [h]:mm:ss means you'll see only the time portion 19:06, not the date information. (For more information about custom number formats, see page 153.) The following formula shows one handy way to calculate the current time by completely removing the date:

```
=NOW()-TODAY()
```

Remember to format the displayed result as a time value.

Note: Excel recalculates both TODAY() and NOW() when you reopen a spreadsheet or when you explicitly refresh the worksheet by pressing F9. But, sometimes, you may want to insert the current date and make sure Excel never updates it again. In these cases, you can use the TODAY() or NOW() function, but you need to convert the result into a static date. Hit F2 to activate edit mode for the cell, and then press F9 to replace the cell with the calculated result. At this point, you'll see the serial number appear in the cell. Finally, press Enter to commit this value.

DATE() and TIME(): Dates and Times in Calculations

As you've learned over the last few chapters, formulas can include cell references and *literals*, or fixed values. So far, you've seen examples of literal numbers (like 1+1) and literal text. Unfortunately, it's not as easy to insert literal date values into a formula. Although Excel makes it easy to type a single date into a cell (just use a format like 1/7/2010), you can't use the same syntax *inside* a formula.

Note: If you simply want to calculate the difference between two dates that are stored in two different cells, you can use an ordinary formula, like =A2-A1 (assuming the dates are in cells A1 and A2).

For example, if you enter the formula *=2/9/2010–1/14/2010*, Excel won't recognize this act as an attempt to calculate the number of days between January 14, 2010 and February 9, 2010. Instead, Excel sees the whole chain of numbers and backward slashes, as well as the minus sign, as one long arithmetic operation involving division. And while you *can* enter a literal date value by typing in the corresponding serial number (as explained on page 325), this technique's confusing and error-prone—unless you're keeping track of the number of days that have elapsed since January 1, 1900.

An easier way to enter literal date values is to use the DATE() function. The DATE() function is a quick and easy way to insert a date into a formula. DATE() accepts three numbers, each of which represents a different component of the date. The function gives you the date's serial number. Here's what it looks like:

```
DATE(year, month, day)
```

In other words, if you enter this formula into a cell:

```
=DATE(2008, 1, 1)
```

Excel displays the date *1/1/2008*. However, you need to watch the formatting you use for the cell. If you use the DATE() function in a cell that uses the General number format, Excel automatically adjusts the cell to use the Date number format. However, if you've already applied a specific non-General number format to the cell (like Currency), the DATE() function won't change the number format. That means you'll see the familiar date serial number (in this example, *39448*), which isn't what you want.

Note: For maximum clarity, you should always use a four-digit year with the DATE() function. If you enter a two-digit year, Excel assumes you're referring to a year that begins with 19, meaning that 04 becomes 1904, not 2004.

Similarly, you can use the DATE() function to create a formula that combines date literals and cell references. Here's an example that determines the number of days between 12/30/2005 and another date (the one contained in cell A1):

```
=A1-DATE(2005, 12, 30)
```

Note: You can use the DATE() function to take a year number, month number, and day number from three different cells and calculate the resulting date. However, it's unlikely that you'll find a worksheet that splits date information into more than one cell.

The TIME() function performs similar magic with time values. It requires three components: an hour number, a minute number, and a second number. If you want, you can use fractional seconds to indicate milliseconds. With this function, you have to enter hours using a 24-hour clock. Here's what the function looks like:

```
TIME(hour, minute, second)
```

For example, the following function calculates how much time exists between now and a deadline at 9:00 p.m. (21:00 in a 24-hour clock):

```
=TIME(21,0,0)-(NOW()-TODAY())
```

As you'll see soon, constructing this formula requires several steps. (For a refresher on the precedence rules Excel uses when determining order of operations, see page 229.)

- First, the formula creates the time literal for 9:00 p.m. using the TIME() function.
- Next, the formula determines the current time using the calculation NOW()–DATE(). This calculation works because it takes the current date and time (the result of the NOW() function), and removes the current date portion (the result of the DATE() function).
- Finally, it calculates the difference between 9:00 p.m. and the current time. This calculation's result, like all time values, is expressed as a fractional number of days. You could multiply this number by 24 to get the final result as number of hours. But a better choice is to simply apply a Time number format to the cell.

DAY(), MONTH(), and YEAR(): More Date Calculations

The DAY(), MONTH(), and YEAR() functions are great when you want to calculate dates that fall before or after a certain date you already know. All three functions take a date argument and give you a number representing the day (1 to 31), the month (1 to 12), or the year (1900 to 9999), respectively. For example, if you place the date 1/1/2010 in cell A1, the following formula displays a result of 2010:

```
=YEAR(A1)
```

If you want to use the DAY(), MONTH(), and YEAR() functions with date literals, you need to use these functions in conjunction with the DATE() function. For example, the following formula displays the number 5:

```
=MONTH(DATE(2010,05,20))
```

Note: The DAY(), MONTH(), and YEAR() functions all require that you put their arguments in the form of a date's underlying serial value, so a formula like *=YEAR(1/1/2010)* won't give you the value you're presumably looking for (that is, 2010). Instead, Excel calculates the result of the division 1/1/2010 (which is 0.000498), and then passes that to the YEAR() function, which has no real meaning. To correct this problem, use the DATE() function to create the serial number for the date you need.

The DAY(), MONTH(), and YEAR() functions really shine when you need to take an existing date and move it forward or backward a set number of days, months, or years. For example, say that, given any date, you want to know what the date will be in two weeks. You could use the DAY(), MONTH(), and YEAR() functions in conjunction with the DATE() function to find out. Here's the formula you'd use (assuming your base date was in cell A1):

```
=DATE(YEAR(A1), MONTH(A1), DAY(A1)+14)
```

Here's how this formula breaks down, assuming that cell A1 contains the date 5/14/2010. The YEAR() and MONTH() functions both generate pretty straightforward results for the DATE() function's first two arguments: 2010 and 5, respectively. The DAY() function gives you a value of 14, which, when added to 14, results in 28, which serves as the third argument in the DATE() function. The final result of the formula, therefore, is 5/28/2010.

Impressively, Excel handles these calculations correctly even if the DAY argument's *greater* than the number of days in the month. For example, if you pass 33 as the third argument to the DATE() function, and the current month has only 31 days, Excel simply rolls over to the next month. (Excel knows how many days are in every month.) The DATE() function performs a similar trick to increment the year if the second argument you use is greater than 12.

Note: As you may expect, the TIME() function, which uses three arguments (hour, minute, second), works similarly. Namely, if you try to pass the function an hour value that's 24 or greater, or a number value that's 60 or greater, TIME() automatically rolls these values over. Thus, *=TIME(25, 10, 00)* results in 1:10 a.m..

These DAY(), MONTH(), and YEAR() functions are also useful for building conditions. For example, you can use MONTH() to check if a given date falls on a specific month or day, regardless of the year. The following formula shows a simple example that displays the heading "First of the month" when the day component of the current date is 1. In all other cases, Excel leaves the cell blank.

```
=IF(DAY(TODAY())=1, "First of the month", "")
```

For much more on how to use conditional logic, see Chapter 13.

HOUR(), MINUTE(), SECOND(): More Time Calculations

The HOUR(), MINUTE(), and SECOND() functions work the same way as the DAY(), MONTH(), and YEAR() functions, except they each generate a number representing the number of hours, minutes, or seconds in a given time. The hour component always uses a 24-hour clock. For example, if you type 9:30 PM into cell A1 (which Excel stores as the serial number 0.89583), the following formula displays the number 21:

```
=HOUR(A1)
```

And this function shows the minute component, which is the number 30:

```
=MINUTE(A1)
```

WEEKDAY(): Determining the Day of the Week

The WEEKDAY() function takes a date, and returns a number that represents which day of the week that date falls on. For example, if the date falls on Sunday, the number's 1, on Monday it's 2, right up through Saturday, which is 7. If you need this sort of information, the WEEKDAY() function is indispensable.

UP TO SPEED

Excel's Intentional Date Bug

Most software programmers spend their waking hours struggling to eradicate bugs from their software. So it may come as a bit of a surprise that Microsoft programmers have deliberately *inserted* at least one bug into all versions of Excel.

People know this bug affectionately as the *date leap year bug*. As leap year aficionados know, the year 1900 is *not* a leap year. According to Excel, however, the year 1900 *is* a leap year. So if you enter the following formula, Excel won't complain, even though February 29, 1900 isn't an actual date:

```
=DATE(1900,2,29)
```

The reason for this error is compatibility. When Microsoft released Excel, Lotus 1-2-3 was by far the most popular spreadsheet software. Lotus 1-2-3 contained this bug, and by emulating it, Excel could use the same date numbering system, enhancing compatibility.

This bug also means that the days of the week prior to March 1, 1900, are incorrect, and that date calculations that stretch over this date will also be off by one day.

It's unlikely that Microsoft will ever fix this problem. Most conceivable fixes would change the internal numbering for every date after February 28, 1900, potentially causing problems with every existing spreadsheet that uses dates. Fortunately, this problem's also fairly benign, because few spreadsheets use dates before February 29, 1900.

Here's what the function looks like:

```
WEEKDAY(date, [return-type])
```

The first argument in the WEEKDAY() function is the date serial number. The WEEKDAY() function also accepts an optional second parameter, which you can use to change the numbering system. The second parameter can take one of three preset values: 1, 2, or 3. If you specify number 1, Excel uses the default numbering system described earlier. Number 2 instructs Excel to start counting at 1 with Monday and end at 7 with Sunday. If you specify number 3, Excel starts counting at 0 with Monday and ends at 6 with Sunday.

Here's a sample formula that returns 4, indicating that January 6, 2010 falls on a Wednesday:

```
=WEEKDAY(DATE(2010,1,6))
```

It's quite easy to change things if you prefer to show the day's name rather than its number. You simply need a little help from the TEXT() function (page 321), which can convert a number to any number format. In this case, you want to use a custom date format that shows the day name. Here's how:

```
=TEXT(WEEKDAY(DATE(2010,1,6)), "dddd")
```

Here, Excel shows the full day name (*Wednesday*). You can also tweak the format string to show the day's three-letter abbreviation (like *Wed*).

```
=TEXT(WEEKDAY(DATE(2010,1,6)), "ddd")
```

DATEDIF(): Calculating the Difference Between Dates

You've already learned how to use the DATE() function (page 327) to subtract one date from another and display the results in terms of days. But what if you want to calculate the difference in years or months? You could try and break the date up into components using the DAY(), MONTH(), and YEAR() functions, but these types of calculations can get surprisingly complicated. You're better off using Excel's little-known DATEDIF() function. Despite this being a useful gem for many date calculations, Excel's own Help tool neglects to cover this function.

FREQUENTLY ASKED QUESTION

Solving the DATEDIF() Mystery

Why isn't DATEDIF() mentioned in Excel's Help?

DATEDIF() is an Excel oddity. It's used by Excel gurus and unknown by just about everyone else. But DATEDIF() isn't anything new–in fact, it was introduced in Excel 5.0. Microsoft even documented the DATEDIF() function in Excel 2000, but the documentation was removed in later versions.

The DATEDIF() function has its origins in Lotus 1-2-3. Apparently, Microsoft included it in Excel for compatibility purposes. Microsoft won't say why it has removed the documentation, but it's most likely due to a legal issue. On the other hand, because it has a similar name to the built-in VBA (Visual Basic for Applications) function DateDiff(), it can create needless confusion. Either way, DATEDIF() is unlikely to ever disappear from Excel because it's used in countless spreadsheets to fill a gap that other date functions don't.

If you're still curious, you can read some unofficial DATEDIF() documentation at *www.cpearson.com/excel/datedif.htm.*

DATEDIF() accepts three arguments. The first two are the dates that you want to find the difference between. The smaller (older) date goes first, and the other date comes second. The third argument is a string that indicates the way you want to measure the interval, such as in days, months, years, and so on. Table 11-3 explains your options. Remember, as with all string arguments, you need to put quotation marks around whatever value you choose. The formula looks like this:

```
DATEDIF(start_date, end_date, interval_type)
```

Note: Remember, you can't use literal dates in any formula (ever). Instead, use tools like the DATE() function to first transform your date into the type of number you can use inside a formula.

Table 11-3. *Interval Strings for DATEDIF()*

Code	Meaning	Description
"m"	Months	The number of complete months between the two dates.
"d"	Days	The number of days between the two dates.
"y"	Years	The number of complete years between the two dates.

"ym"	Months Excluding Years	The number of months between the two dates, as if they were in the same year.
"yd"	Days Excluding Years	The number of days between the two dates, as if they were in the same year.
"md"	Days Excluding Months and Years	The number of days between the two dates, as if they were in the same month and the same year.

For example, here's how to calculate the number of months that separate a date in the future (stored in cell A1) from the current date:

```
=DATEDIF(TODAY(), A1, "m")
```

Remember that when using the DATEDIF() function with the "m" argument, you're counting *complete* months. That means Excel counts the interval from, say, January 6, 2010 to February 6, 2010, as one month, but the interval from January 6, 2010 to February 5, 2010, is still a day shy, and Excel therefore counts it as zero months. Usually, this is the behavior you want. However, you do have another, somewhat more complicated option, if you want intervals like January 6–February 5 to register as one month: use the YEAR() and MONTH() functions.

For example, here's the DATEDIF() approach (which has a result of 0):

```
=DATEDIF(DATE(2010,1,6), DATE(2010,2,5), "m")
```

And here's the YEAR() and MONTH() approach, which has the result of 1 (it works by subtracting the one month number from the other):

```
=MONTH(DATE(2010,2,5))-MONTH(DATE(2010,1,6))
```

And here's a revised approach that works with dates that aren't in the same year:

```
=(YEAR(DATE(2011,2,5))-YEAR(DATE(2010,1,6)))*12 +
MONTH(DATE(2011,2,5))-MONTH(DATE(2010,1,6))
```

Although this formula looks more complicated at first glance, it's really not that difficult. It's so long because Excel calculates the year and month components separately. Once you find the difference in year numbers, Excel multiplies that number by 12, and then adds it to the month component. You then end up with the total number of months.

Unfortunately, this formula assumes that every year is 365 days, which neglects leap years. This formula is probably right most of the time, but it fails in the days just before or after a person's birthday.

DATEVALUE() and TIMEVALUE(): Converting Dates and Times into Serial Numbers

DATEVALUE() and TIMEVALUE() convert dates and times that Excel has stored as text into serial date numbers. Of course, in order for this conversion to work, the text must be interpretable as a date or time. So if you type *1-1-2010* or *January 1, 2010* into a cell, these functions would be able to convert these values, but if you entered something like *1st January, 2010* or *1,1,2010*, you'd get an error message.

POWER USERS' CLINIC

Getting a Birthday Right

Now that you've learned about a wide variety of Excel date functions, it probably seems like it wouldn't be very difficult to calculate a person's age based on the current date and his or her birth date. But, in fact, it can be surprisingly tricky.

One approach is to use Excel's support for date subtraction. The following formula is a good first try (assuming the birthday is stored in cell A1):

```
=INT((NOW()-A1)/365)
```

Unfortunately, this formula assumes that every year is 365 days, which neglects leap years. This formula is probably right most of the time, but it fails in the days just before or after a person's birthday.

The YEAR(), MONTH(), and DAY() functions don't provide a solution–they all suffer from the same problem of not taking leap years into account. The only real solution is using DATEDIF(), which *x* take leap years into account. Here's the DATEDIF() formula you'd need:

```
=DATEDIF(A1,NOW(),"y")
```

You can even get a little fancier with the following formula, which uses some of the text manipulation techniques you saw earlier in this chapter. It displays a person's age in years, months, and days:

```
=DATEDIF(A1,NOW(),"y") & " years, " &
DATEDIF(A1,NOW(),"ym") & " months, " &
DATEDIF(A1,NOW(),"md") & " days"
```

Note: Keep in mind that the DATEVALUE() and TIMEVALUE() functions don't change the formatting of the cell. Therefore, if you want to see something other than the date serial number, you'll need to choose a Date number format.

The rules Excel uses to convert a piece of text to a date with DATEVALUE() are the same ones that Excel uses for recognizing dates and time values when you enter them into a cell. To refresh your memory, see page 63. The only difference is that DATEVALUE() and TIMEVALUE() ignore the initial apostrophe in a text value.

For example, you can type the following text in cell A1. Remember that the apostrophe tells Excel to treat this entry as text even though it's clearly recognizable as a date:

```
'1/1/2004
```

Now another cell can perform a date calculation by first converting this text to a real date:

```
=NOW()-DATEVALUE(A1)
```

This formula calculates the number of days elapsed since January 1, 2004. Because the result is a number of days, you'll want to use the General number format (not a Date number format).

DATEVALUE() and TIMEVALUE() also work with literal text strings. That means you can use these functions as an alternative to DATE() and TIME() if you want to create a date or time literal. The difference? Whereas DATE() and TIME() create dates and times based on several numeric components you supply, DATEVALUE() and TIMEVALUE() create dates and times from string literals. Here's an example of TIMEVALUE() in action:

```
=TIMEVALUE("19:30:00")
```

This formula converts a static piece of text into the time value 19:30:00 (which Excel represents internally as the fractional number 0.8125). The only drawback to using DATEVALUE() and TIMEVALUE() is that you'll run into trouble if Excel can't interpret the text you supply. For example, the text string *January 1,2010* isn't formatted properly, because there's no space between the comma and the number 2010. This minor glitch is enough to completely stymie the DATEVALUE() function. It's much less likely that you'll make an error if you rely on the DATE() function instead.

DAYS360(): Finding Out the Number of Days Between Two Dates

DAYS360() gives you the number of days between two dates. Unlike simple date subtraction, however, DAYS360() assumes that every month has 30 days (giving the year a total of 360 days). So is this the kind of function you're supposed to use when following some radically alternative calendar system? Sort of. The only reason you'd use DAYS360() is if you were using an accounting system that performs its calculations based on 360-day years, which is common in some payroll systems. (Also, some systems calculate investments by assuming 12 months with 30 days each, for a total of 360 days.) Read on if you're still interested.

Here's what the function looks like:

```
=DAYS360(start_date, end_date, [European_method])
```

The smaller (older) date goes first (unless you want a negative number). The second argument is the date following the first date.

The third argument in DAYS360() is an optional value that you can set to TRUE to use the European version of the DAYS360() calculation. The only difference is how the calculation deals with start or end dates that happen on the 31st day of the month:

- In the default U.S. NASD (National Association of Securities Dealers) method, a starting day of 31 changes to 30. An ending day of 31 is handled differently depending on the start day. If the start day's 30 or 31, then the end date changes to day 30 in the same month. If the start day's earlier than 30, then the end date changes from 31 to the first day of the next month.
- In the European method, both a starting and ending day of 31 change to 30.

EDATE(): Calculating Future Dates

The EDATE() function (short for Elapsed Date) calculates a future date by adding a certain number of months to a date you supply. You specify two parameters: the starting date, and the number of months you want to move forward (use a negative number to move one or more months backward).

Here's an example that calculates a date one month from today:

```
=EDATE(NOW(),1)
```

The EDATE() function would be more useful if it provided a similar ability to advance a date by a set number of days or years. If you need to do that, then you need to resort to using the DAY(), MONTH(), and YEAR() functions, in conjunction with the DATE() function, as described earlier.

YEARFRAC(): Calculating the Percentage of a Year Between Two Dates

The YEARFRAC() function (short for Year Fraction) lets you take a range of days between two dates in the same year and determine what percentage this represents out of the whole year. For example, if you pay for a monthly fitness club membership and cancel it after a few weeks, this function may help determine what portion of your money the club should refund (provided they'll spare you from the usual Draconian contract clauses).

The YEARFRAC() function requires two parameters: the start date and the end date (which can come from similar or different years). In addition, you can specify a third parameter to indicate how Excel should calculate the fraction. Here's what the function looks like:

```
YEARFRAC(start_date, end_date, [basis])
```

The basis has to be a number from 0 to 4. The meaning of the number is as follows:

- **0.** Excel performs the calculation in the same way as the DAYS360() function. The calculation assumes that every month has 30 days, and every year has 360 days. This system is primarily useful for accounting systems that use 360-day years in their calculations.
- **1.** The calculation gives you the fractional amount of a year that exists between two days. This basis counts the actual number of days in the range, and then divides this figure by the number of days in the year. It correctly takes leap years into account.
- **2.** The calculation counts the number of days between your two dates, but divides this total by 360 instead of the actual number of days in the appropriate year.
- **3.** The calculation counts the number of days between your two dates, but divides this total by 365 instead of the actual number of days in the appropriate year. This calculation provides the same answer as basis 1, assuming the year isn't a leap year.
- **4.** The calculation uses the European version of the DAYS360() function.

If you don't specify a basis number, Excel uses a basis of 0.

The following formula shows an example that calculates the fraction of the year represented by the range from January 1, 2010 to February 14, 2010:

```
=YEARFRAC(DATE(2010,1,1),DATE(2010,2,14),1)
```

This formula gives you approximately 0.12, or 12 percent, of the year.

EOMONTH(): Finding the Last Day of Any Month

The EOMONTH() function (short for End of Month) calculates the last day of any month in any year. However, it doesn't work quite the way you'd expect. EOMONTH() is designed so that you can look into the future and answer questions like "Two months from now, what's the last day of the month?" This quirk seems confusing at first, but business people like it because it helps them set payment periods and due dates for invoices.

When using EOMONTH(), you need to supply two parameters: the starting date, and the number of months you want to look into the future. If you want to look into the past, use a negative number for the number of months (like –2 to go back in time 2 months). Here's what the basic EOMONTH() function looks like:

```
EOMONTH(start_date, number_of_months)
```

The following sample formula takes the date 1/1/2010, moves it to the next month (February), and then provides the last day of that month (February 28, 2010):

```
=EOMONTH(DATE(2010,1,1), 1)
```

If you want to find the last day in the current month, just specify 0 for the second parameter. Here's a formula that uses this approach to find the last day in the current month:

```
=EOMONTH(TODAY(), 0)
```

Like all of Excel's date functions, the EOMONTH() function is intelligent enough to handle leap years correctly.

NETWORKDAYS(): Counting the Number of Business Days

The NETWORKDAYS() function counts the number of business days in a given range. It requires two parameters: a start and an end date. In addition, you can supply a third parameter that specifies holidays that Excel shouldn't count. The basic function looks like this:

```
=NETWORKDAYS(start_date, end_date, [holidays])
```

The following formula gives you the number 21, which is the number of working days in January 2010:

```
=NETWORKDAYS(DATE(2010,1,1), DATE(2010,1,31))
```

You can specify holidays in several ways. When you have a single holiday to specify, you can set it as a date literal. Here's how you'd rewrite the earlier example so that January 15, 2010 is designated as a holiday and, therefore, not included in the calculation:

```
=NETWORKDAYS(DATE(2010,1,1), DATE(2010,1,31), DATE(2010,1,15))
```

If you want to specify more than one date, you have to use a messier syntax, which writes each date as a text string in the month/day/year format. Here's an example that excludes January 2, 2010, and January 5, 2010. Note that you need to surround the holidays with curly braces:

```
=NETWORKDAYS(DATE(2010,1,1), DATE(2010,1,31), {"1/2/2010","1/5/2010"})
```

Your final option is to use the holidays specified in a range of cells on the worksheet. For example, if you've entered a holiday schedule in the cells C1:C11, you could use it with the NETWORKDAYS() function, as shown here:

```
=NETWORKDAYS(DATE(2010,1,1), DATE(2010,1,31), C1:C11)
```

WORKDAY(): Figuring Out When Days Will Fall in the Future

The WORKDAY() function lets you perform date calculations that take workdays, weekends, and holidays into account. Essentially, the WORKDAY() function takes a date you specify, and moves it into the future by a certain number of business days (skipping over weekends automatically). One of the most common reasons to use WORKDAY() is to estimate a due date by taking the current date, and adding the number of business days it should take to complete a given task. The WORKDAY() function accepts three parameters: the start date, the number of days you want to move into the future, and any days that Excel should ignore because they represent holidays. The function looks like this:

```
WORKDAY(start_date, days, [holidays])
```

You can use the optional holiday parameter to specify days that Excel would ordinarily treat as workdays, but which you don't want to consider as such. You can also supply a reference to a range of cells that contain all the holiday dates.

For example, imagine you want to find out when you will complete a project if you start working today, and it requires 30 days of work. You can use the following formula to estimate the completion date, assuming you don't work on weekends:

```
=WORKDAY(TODAY(), 30)
```

As with many of the date functions, you'll need to format the cell so that the value Excel generates (the date serial number) appears as a date.

WEEKNUM(): Figuring Out in Which Week a Date Falls

The WEEKNUM() function accepts a date and gives you a number from 1 to 52 to indicate where the date falls in the year. For example, if the cell that WEEKNUM() refers to contains the date 1/1/2010, then the function gives you the value of 1 (indicating the first week in the year). If WEEKNUM() gives you 52, the date falls on the last week of the year.

CHAPTER

12

Lookup, Reference, and Information Functions

Excel's lookup, reference, and information functions are quite a bit different from the mathematical functions discussed in the last few chapters. Most math functions take your data as a starting point and use it to calculate some new result. Lookup, reference, and information functions, on the other hand, don't generate any *new* data. Instead, they let you search for and extract important bits of content from your worksheet and then reuse it in ways guaranteed to delight spreadsheet lovers worldwide.

All three types of functions play different but complementary roles. A *lookup function* finds and then copies data from a particular cell in a worksheet. A *reference function* retrieves more general information about groups of cells (like how many columns' worth of data is in your table). It's common to use lookup and reference functions together. You can use a reference function to find out which cell has the data you want to use, and a lookup function to actually retrieve the cell's contents.

Finally, *information functions* let you determine what *kind* of content resides in an individual cell (for instance, a number or text). Knowing that can help you construct extremely powerful *conditional formulas*, which behave differently depending on the type of data they encounter.

All these functions may not immediately seem useful, but they become indispensable in a variety of situations. Consider a few common examples:

- You want to let spreadsheet viewers look up items, and then see these items used in a series of calculations. Say you have a worksheet with a long list of customer records (each of which includes info like their mailing address and how much money they owe your company). You not only want to make the spreadsheet searchable, but you also want to present your searchers with a series of calculations about the items they've found, like how much money is owed to your company after 30 days plus interest, 60 days plus interest and penalty, and so on. Lookup and reference functions let you turn your spreadsheets into extremely useful, interactive databases. It's kind of like having an accountant-in-waiting for anyone who wants to ask questions about the information in your spreadsheet.

- You want to extract some data from a table and use it in another worksheet. Maybe you have a list of mailing addresses, and you need to extract a specific address and place it at the top of a form letter. Lookup functions are perfect for that.
- Your worksheet has more than one table, and they're all designed to work together. Say you want to pluck out a product's name and price from a massive table of product information, and then use it to create an invoice. This example represents one of the most powerful ways you can use lookup functions.

Tip: Excel user groups and online bulletin boards are full of problems that you can solve only with lookup functions. Try, for instance, the Microsoft Office Discussion groups at *www.microsoft.com/office/community/en-us/flyoutoverview.mspx* or third-party discussion groups like *www.msofficeforums.com/excel*.

In this chapter, you'll learn how to use the lookup and information functions, as well as how to overcome some of their limitations. Finally, you'll consider the oddest lookup and reference function of them all: the HYPERLINK() function, which lets you create Web-style hyperlinks in a worksheet cell.

The Basic Lookup

VLOOKUP() and HLOOKUP() are the two most popular lookup functions. They perform the same task—finding and copying data—but they look in different directions:

- **VLOOKUP() is the vertical lookup function**. You use it to find a specific row in a large table of data. VLOOKUP() works by scanning the values in a single column from top to bottom. Once it finds the entry you're looking for, it can then retrieve other information from the same row. VLOOKUP(), examined in closer detail in the next section, is the more commonly used of the two lookup functions.
- **HLOOKUP() is the horizontal lookup function**. You use it to find a specific column in a large table of data. HLOOKUP() works by scanning the values in a single row from left to right. Once it finds the entry you're looking for, it can then retrieve other information from the same column.

Both VLOOKUP() and HLOOKUP() require three parameters, and they have one optional parameter:

```
VLOOKUP(search_for, table_range, column_number, [range_lookup])
HLOOKUP(search_for, table_range, row_number, [range_lookup])
```

The first parameter, *search_ for*, is the actual cell content that you're trying to find. The second parameter, *table_range*, is the rectangular grid of cells that contains all the data you're searching through. The *column_number* and *row_number* parameters are known as *offsets*. Once the VLOOKUP() or HLOOKUP() function finds the requested data, the offset tells Excel how many cells to move over or down in order to find related data you want to retrieve.

Imagine a spreadsheet filled with customer records, each listed in its own row. You may initially use VLOOKUP() to find a customer's last name, but what you're really interested in is how much money she owes you. To find that crucial piece of info, you'd use the offset argument to retrieve the value from a cell that's a few columns away from the last name column.

Finally, the *range_lookup* parameter is a Boolean (true or false) value. If you specify TRUE, or leave out this parameter, then Excel finds approximate matches. If you specify FALSE, Excel gives you either an exact match or, if it finds no match, the error value #N/A. (You can also substitute 1 for TRUE or 0 for FALSE, if you find that's easier.)

VLOOKUP(): Vertical Lookups

The simplest way to understand how lookup functions work is to study one of these functions in action. Looking at VLOOKUP() is a good place to start. Figure 12-1 shows a worksheet that uses VLOOKUP() to let spreadsheet viewers learn more about the products contained in a long list of exotic grocery store items. The idea here is to offer a kind of interactive information dashboard on top of the spreadsheet so that anyone can enter a product's ID number, and then view specific information about the product (like the value of the inventory in stock).

True, you could simply construct your worksheet so this information appears alongside the product data that's already listed. But when you have multiple scenarios that you want to let people play around with, your spreadsheet can get more crowded than a cross-town bus at 5:30 p.m. The functions in this chapter offer a better way to offer customizable views for large amounts of information.

Here's a blow-by-blow breakdown of what's happening in the Figure 12-1 spreadsheet. The key cells are B2 and cells C4 through C11. B2 is where whoever's reading the spreadsheet types in the product ID. C4 through C8 all contain different versions of the LOOKUP() function, which retrieve Product Name, Unit Price, the number of units In Stock, the number of units On Order, and the Reorder Level, respectively. The actual calculations, which are pretty straightforward formulas, take place in cells C9, C10, and C11.

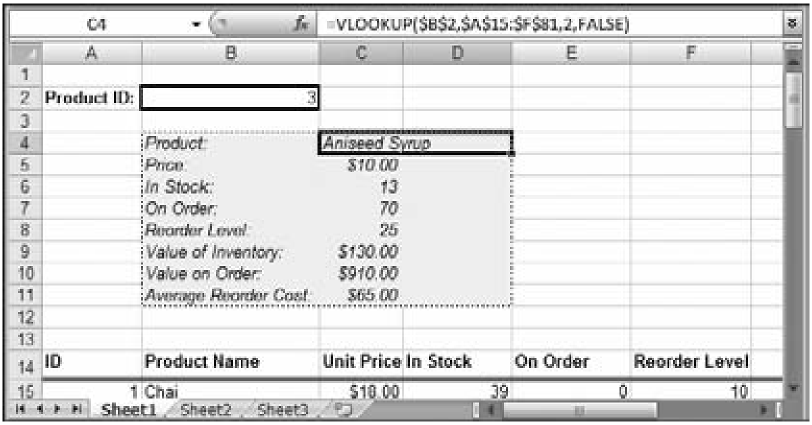

Figure 12-1: *VLOOKUP() in action: The person using this spreadsheet simply types in the appropriate product number in cell B2, and then the lookup formulas in cells C4 through C8 automatically retrieve different kinds of price and inventory data (which starts in row 15). The cells C9 through C11 don't need any lookup functions–they simply use the values from the lookup cells to perform additional calculations.*

Consider how Excel retrieves the product name, shown in cell C4. The formula is:

```
=VLOOKUP(B2, A15:F81, 2, FALSE)
```

Here's what you're looking at:

- The first parameter, B2, is the product ID number (as typed in by the person reading the spreadsheet). This is the value VLOOKUP() is seeking.
- The second parameter is the most important: It gives the cell range where Excel should search. This example has 66 product rows (from row 15 to 81) and six columns of data (from column A to F). The first column in this cell range must be the column that contains the data you're trying to locate (in this case, the product ID). The other columns contain the remaining information Excel is going to retrieve once it finds a match.
- The third parameter, 2, indicates the position of the data you want to retrieve. Excel numbers each column in the range that you supply from left to right, starting at 1. Thus, 2 represents the Product Name column.
- Finally, FALSE indicates that you're insisting on an exact match to ensure you have the correct product.

Now, if someone typed the product ID *21* into cell B2, for instance, then each of the VLOOKUP() functions contained in cells C4 through C8 finds the requested row, and then gives you the data they're designed to retrieve (Product, Price, In Stock, and On Order, respectively).

Tip: If you supply a range lookup value of TRUE, or if you omit this parameter altogether, Excel accepts a partial match when it performs the search. That is, it attempts to find an ID value that matches exactly, but if it can't find a match, it uses the closest ID value that's *less* than the lookup value B2. As a result, it's always a good idea to explicitly turn off range matching if you're trying to look up specific items.

Once you've retrieved the information you need, it's fairly easy to use it in other ordinary calculations. For example, the formula in cell C9 (Figure 12-1) calculates the total value of all the products on order by multiplying the retrieved unit price with the retrieved number of units on order:

```
=C6*C7
```

Note: Perhaps the most convenient feature of a worksheet that uses lookup functions is that the worksheet has to store only one copy of the data. So, in the previous example, if you modify the product list, the information you get from any linked lookup formulas updates automatically.

POWER USERS' CLINIC

Advanced Text Lookup

The product lookup example shown in Figure 12-1 searches by product ID, which is a numeric value. You can also use lookup formulas to search for *text*. You could modify the product lookup example to conduct searches using a product's name. The price lookup formula would then look like this:

```
=VLOOKUP(B2, B15:F81, 2, FALSE)
```

This version of the price lookup formula has two changes (compared to the one on page 342). First, the range of searched data starts at B15 instead of A15. That's because the column that Excel searches must always be the leftmost column. The second change is the offset. Because the first column is B, the VLOOKUP() function numbers that column as 1, and the adjacent Price column as 2 (instead of 3). Consequently, you need to modify your offset accordingly.

Text lookups are useful, but you may find them a bit cumbersome. As a rule of thumb, when you perform a text search, you should always specify FALSE for the range lookup parameter to prevent some quirky behavior that happens with unsorted lists. In this example, the person using the spreadsheet would need to type in the entire product name exactly as it appears in the spreadsheet, in order for the VLOOKUP() function to match it successfully.

You have one workaround for this limitation: the ever-popular asterisk (*) wildcard. If you add this to the end of the search value, then the VLOOKUP() function finds any cell that starts with this text. If you type in *Queso**, then the VLOOKUP() function matches the Queso Cabrales product. Actually, you can go one step further, and change your lookup formulas so they always accept partial matches. The trick is to add the asterisk onto the end of the cell value in your formula, as shown here:

```
=VLOOKUP(B2&"*", B15:F81, 2, FALSE)
```

This formula uses the ampersand (&)—a special operator for joining text. For more information about manipulating text, see Chapter 11.

HLOOKUP(): Horizontal Lookups

Now that you've mastered vertical lookups, it's a snap to perform horizontal lookups. The key distinction is that HLOOKUP() searches *across* the first row in the cell range you supply. You use the offset to retrieve data from other rows in the matching column.

Figure 12-2 shows an example of HLOOKUP() used with a grid of product information:

```
=HLOOKUP(C11,A1:E8,2)*C12
```

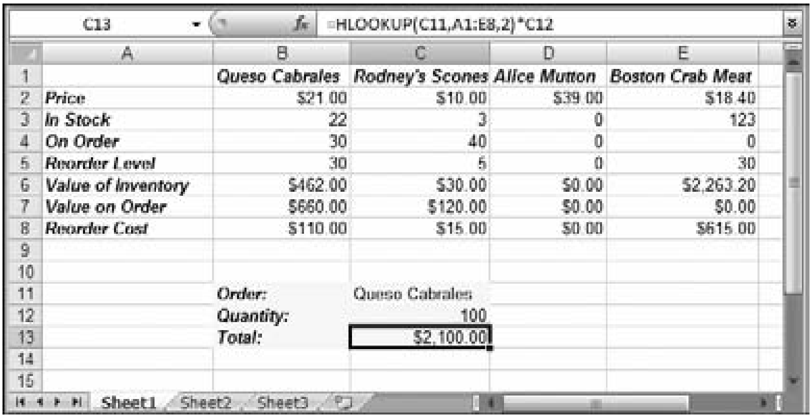

C13 fx =HLOOKUP(C11,A1:E8,2)*C12

	A	B	C	D	E
1		Queso Cabrales	Rodney's Scones	Alice Mutton	Boston Crab Meat
2	Price	$21.00	$10.00	$39.00	$18.40
3	In Stock	22	3	0	123
4	On Order	30	40	0	0
5	Reorder Level	30	5	0	30
6	Value of Inventory	$462.00	$30.00	$0.00	$2,263.20
7	Value on Order	$660.00	$120.00	$0.00	$0.00
8	Reorder Cost	$110.00	$15.00	$0.00	$615.00
9					
10					
11		Order:	Queso Cabrales		
12		Quantity:	100		
13		Total:	$2,100.00		
14					
15					

Sheet1 Sheet2 Sheet3

Figure 12-2: *Horizontal lookups are just as straightforward as vertical lookups; the direction is the only thing that changes. Here, Excel uses the function HLOOKUP to scan the top row of product names.*

This formula tells Excel to search the range of cells from A1:E8. Because this is a horizontal lookup, Excel starts out by inspecting the first row (cells A1 to E1), where it attempts to find a cell that matches the content in C11 *(Queso Cabrales)*. When Excel finds a match, it uses the offset of 2 to step down to the second row and retrieve the cell value with the product price *(21)*. Finally, the formula multiplies the retrieved price with the number of units *(100)*, which are in cell C12.

Advanced Lookups

VLOOKUP() and HLOOKUP() work well for linking together tables in a worksheet. They also impose a few restrictions, however, including:

- The lookup column has to be the first column or row in the range. That means you can't use VLOOKUP() to retrieve columns to the left of the lookup column, and you can't use HLOOKUP() to retrieve rows above the lookup row.
- You must choose between horizontal and vertical lookups. You can't perform a lookup that finds cells based on a lookup column and a lookup row.

In this section, you'll learn how to skirt these restrictions with the help of other functions.

Tip: You simply can't get around certain lookup rules. The lookup functions aren't much use if you have potentially multiple matches. That means you can't use a lookup function to retrieve your top-10 selling products, for example. If you want to use this sort of logic, then you should probably opt for Excel's list feature (Chapter 14), which provides filtering capabilities.

MATCH(): Finding the Position of Items in a Range

The MATCH() function lets you find the position of an item in a range. On its own, MATCH() doesn't accomplish a whole lot, but when used in conjunction with some of the functions described later in this section, you'll find it's really handy. Here are some MATCH() fundamentals.

To use MATCH(), you simply specify the search value (either a number or text) and the range you want to search:

```
MATCH(search_for, range, [match_type])
```

The range you use must be one-dimensional. That means you can search through the column of cells A1:A10 or the row of cells A1:E1, but you can't search the *grid* of cells A1:E10.

The *match_type* argument's optional, but highly recommended. It can take one of these three values:

- **0**. Gives you the position of the first item that matches exactly.
- **1**. Gives you the position of the largest value that's equal to or less than the search argument. In order for this argument to work correctly, you must perform a lookup in a range of numbers or text values that are ordered in ascending order. If you use the formula *=MATCH(100, {12,23,48,101,321}, 1)*, the result is 3 because the third value in the range *(48)* is the largest number less than 100.
- **–1**. Gives you the position of the smallest value that's equal to or greater than the search argument. In order for this argument to work correctly, you must perform a lookup in a range of numbers or text values that are ordered in descending order. If you use the formula *=MATCH("b", {"dough-nut","cracker", "banana","apple"},–1)*, the result is 3 because the third value in the range (banana) is the smallest text value that falls after b in dictionary order.

Note: You'll almost always want to use a match_type of 0. However, the MATCH() function automatically uses a match_type of 1 if you omit the match_type parameter. Consider yourself warned.

If MATCH() finds the value you're searching for, it gives you a number indicating its position. If you're searching the range of cells A1:A10, and the search item's in cell A3, then the MATCH() function returns 3 (because it's the third cell in this range). If it finds no match, MATCH() returns #N/A.

INDEX(): Retrieving the Value from a Cell

You may have already noticed that the MATCH() function really does half the work of the VLOOKUP() and HLOOKUP() functions. It can find the position of the search item, but it gives you the position as an *index* number, instead of the actual cell content. (Index 1 is the first item in a range, 2 is the second, and so on.) The MATCH() function becomes useful only when you combine it with another function, like INDEX().

INDEX() is powerful not only because it can retrieve the actual cell content, but also because it lets you move to any row or column you want. The next section shows you how useful that feature is. But first, here's a quick look at how INDEX() works.

INDEX() gives you a value from a range of cells, using the index number you specify:

```
INDEX(range, row_number, [column_number]
```

Consider the following formula:

```
=INDEX(A2:A10, 3)
```

This formula returns the content of the third cell in the range: A4. If you supply a two-dimensional range, you need to give a row and column index. Here's an example:

```
=INDEX(A2:B10, 3, 2)
```

This formula gives you the cell in the third row and the second column, which is B4.

Note: Remember that the index numbers used by MATCH() and INDEX() are offsets, not actual cell references. In other words, if you specify a range that starts on cell E10, Excel designates E10 as index number 1, E11 as index number 2, and so on. The actual row number in the worksheet doesn't matter.

Performing a "Left Lookup"

On their own, MATCH() and INDEX() are just curiosities. But when you combine them, you have the ability to overcome many of the limitations inherent in VLOOKUP() and HLOOKUP().

Consider the worksheet in Figure 12-3. Here, Excel performs the lookup using column B (the product name), and the formula retrieves information from column A (the product ID). This example's a *left lookup*, something that isn't possible using the VLOOKUP() function.

To solve this problem, you first use MATCH() to find the position of the product name you're looking for:

```
=MATCH(B2, B8:B74, 0)
```

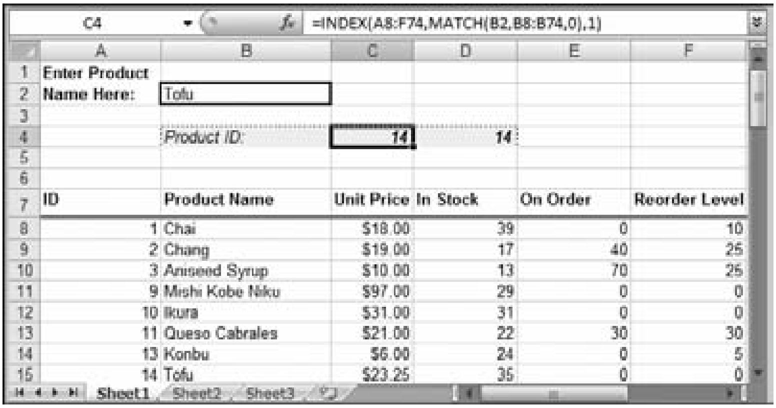

C4 =INDEX(A8:F74,MATCH(B2,B8:B74,0),1)

	A	B	C	D	E	F
1	Enter Product					
2	Name Here:	Tofu				
3						
4		Product ID:	14	14		
5						
6						
7	ID	Product Name	Unit Price	In Stock	On Order	Reorder Level
8	1	Chai	$18.00	39	0	10
9	2	Chang	$19.00	17	40	25
10	3	Aniseed Syrup	$10.00	13	70	25
11	9	Mishi Kobe Niku	$97.00	29	0	0
12	10	Ikura	$31.00	31	0	0
13	11	Queso Cabrales	$21.00	22	30	30
14	13	Konbu	$6.00	24	0	5
15	14	Tofu	$23.25	35	0	0

Sheet1 Sheet2 Sheet3

Figure 12-3: *This worksheet is designed to let you search for a product by name, and then retrieve its corresponding ID information. Because the ID info is to the left of the lookup column, VLOOKUP() doesn't help. The solution is to combine MATCH() and INDEX(), so that MATCH() gives you the offset row number from the Product Name column. INDEX() then uses that value in the ID column. Excel displays the final result–the product ID–in cell C4, which contains the formula that's doing all the work. In essence, this formula breaks the problem down into two questions: "How far down is this product name?" and "What product ID is just as far down?" MATCH() answers the first question, while INDEX() answers the second.*

This formula gives you a value of 8 because the searched for value *(Tofu)* is located in the eighth row of the cell range you're searching. Next, you can use the value from the MATCH() function to retrieve the cell content by employing the INDEX() function. The trick is that the INDEX() function uses a *range* that covers all the rows and columns in the table of data:

```
=INDEX(A8:F74, match_result, 1)
```

Here's how you'd combine these functions into a single formula:

```
=INDEX(A8:F74, MATCH(B2,B8:B74,0), 1)
```

Tip: Remember, the INDEX() function lets you perform a lookup in a column or a row.

Performing a Double Lookup

Another shortcoming with the VLOOKUP() and HLOOKUP() functions is that you can't use them simultaneously. These functions don't help if you want to write a formula that finds a cell at the intersection of a specific column and row heading. Imagine you want to write a formula to find the number of sales recorded in January 2004 at the London location, in the spreadsheet shown in Figure 12-4. You can solve this problem with the INDEX() function, which is much more flexible than either the VLOOKUP() or the HLOOKUP() function.

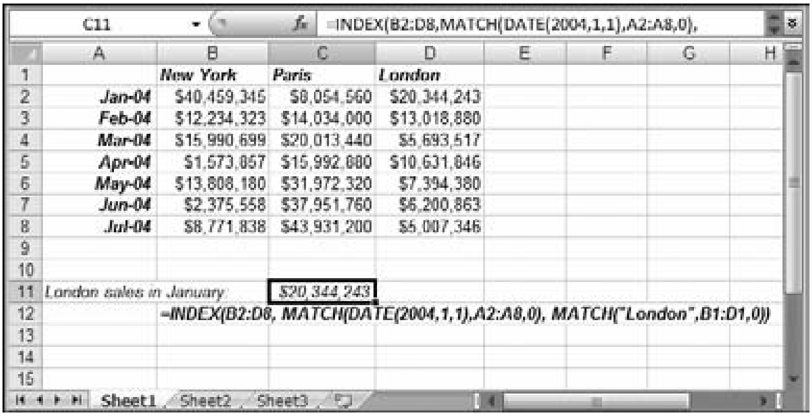

C11 =INDEX(B2:D8,MATCH(DATE(2004,1,1),A2:A8,0),

	A	B	C	D
1		New York	Paris	London
2	Jan-04	$40,459,345	$8,054,560	$20,344,243
3	Feb-04	$12,234,323	$14,034,000	$13,018,880
4	Mar-04	$15,990,699	$20,013,440	$5,693,517
5	Apr-04	$1,573,857	$15,992,880	$10,631,846
6	May-04	$13,808,180	$31,972,320	$7,394,380
7	Jun-04	$2,375,558	$37,951,760	$6,200,863
8	Jul-04	$8,771,838	$43,931,200	$5,007,346
11	London sales in January		$20,344,243	
12		=INDEX(B2:D8, MATCH(DATE(2004,1,1),A2:A8,0), MATCH("London",B1:D1,0))		

Figure 12-4: *The formula in cell C11 performs a double lookup by using MATCH() twice to get the offsets, and INDEX() once to retrieve the value. This method works because the INDEX() function accepts both a row and a column offset.*

Note: With just a tad more work, you can write a formula that lets readers of your spreadsheet indicate which city and which month they're interested in seeing sales figures for. The next section covers these steps.

Because you need to look up the rows *and* columns, you'll need to use the MATCH() function twice.

First, you'd use the MATCH() function to find the row that has the January sales figures:

```
=MATCH(DATE(2004,1,1), A2:A8, 0)
```

Note: Remember that the DATE() function, as explained on page 327, creates a date without your needing to know its underlying serial number.

Next, you'd use MATCH() to find the column with the sales figures for the London office:

```
=MATCH("London", B1:D1, 0)
```

With these two pieces of information, you can build an INDEX() function that searches the whole range for the appropriate sales figure:

```
=INDEX(B2:D8, MATCH(DATE(2004,1,1),A2:A8,0), MATCH("London",B1:D1,0))
```

In this example, the search values (the date and the store location) are hard-coded in the formula, but you can just as easily retrieve this information from another cell.

OFFSET(): Shifting Cell References

Not surprisingly, Excel gives you another way to solve lookup problems. The OFFSET() function lets you take a cell reference and move it to a new location. If you take the cell reference A1, and use OFFSET() to move it down two rows, then Excel would refer to it as cell A3.

For the most part, everything you can accomplish with the OFFSET() function you can also achieve by combining INDEX() and MATCH(). It's really a matter of preference. Some Excel gurus prefer to use INDEX() and MATCH(), while others find that OFFSET() makes for clearer formulas. Once you've finished this section, you'll be able to make your own choice.

The OFFSET() function requires three arguments: the cell reference you want to move, the number of rows you want to move, and the number of columns you want to move:

```
OFFSET(reference, rows_to_offset, cols_to_offset, [height], [width])
```

The following formula shifts the reference A1 to A3, and then returns whatever content is in that cell:

```
=OFFSET(A1, 2, 0)
```

To move a row upwards or to a column that's to the left, specify a negative number for the *rows_to_offset* or *cols_to_offset* parameter, respectively. You can also use the height and width parameters to expand the cells your reference covers, converting it from a single-cell reference into a full-fledged range. Here's one example:

```
=SUM(OFFSET(D2, 0, 0, 7, 1))
```

This formula uses OFFSET() to expand a reference (D2) to a range containing seven cells (D2:D8). The original cell serves as the "anchor point" that becomes a corner in the range. Then Excel passes on that range to the SUM() function. The end result is equivalent to this formula:

```
=SUM(D2:D8)
```

Now that you understand OFFSET(), you can use it in place of the INDEX() function in most situations. You just need to start at the top-left corner of your table, and use the OFFSET() function to get to the cell you want.

Consider the formula in Figure 12-3, which performs a left lookup on the product list:

```
=INDEX(A8:F74, MATCH(B2,B8:B74,0), 1)
```

You can easily rewrite this using the OFFSET() function, as shown here. Cell A8, at the top-left corner of the table, is the starting point:

```
=OFFSET(A8, MATCH(B2,B8:B74,0)-1, 0)
```

Like those in the INDEX() section (page 346), this formula still lets you search through the entire table of data to find the Product Name (using the Product ID value you typed in). In fact, there's really no meaningful difference between the two functions. In most cases, you can use OFFSET() or INDEX() interchangeably. One exception is that only the OFFSET() function has the ability to convert a single-cell reference into a range.

Other Reference and Lookup Functions

MATCH(), INDEX(), and OFFSET() are definitely the three most useful reference and lookup functions. However, Excel also provides a few more similar functions that occasionally come in handy. Table 12-1 gives you a quick tour.

Table 12-1. *Miscellaneous reference and lookup functions*

Function	Syntax	Description
AREAS()	AREAS(reference)	Returns the number of areas in a reference. An area is a single cell or a range of contiguous cells. If you have more than one area, all the areas need to be enclosed within double parentheses. For example, *=AREAS(A1)* returns a value of 1; *=AREAS((A1, B3:D5, E8))* returns a value of 3.
COLUMN()	COLUMN(reference)	Returns the column number for a reference. *=COLUMN(B3)* returns the number 2, which represents column B. If you leave out the reference, then you get the column number of the current cell (the cell that holds your formula).
COLUMNS()	COLUMNS(range)	Counts the number of columns in a range.
ROW()	ROW(reference)	Returns the row number for a reference. *=ROW(B3)* returns the number 3. If you leave out the reference, then you get the row number of the current cell (the cell that holds your formula).
ROWS()	ROWS(range)	Counts the number of rows in a range.
CHOOSE()	CHOOSE (index_num, value1, value2,…)	Returns a value from a list depending on the index number. For example, *=CHOOSE(1, A1, A2, B1)* returns the content in cell A1 because that's the first position in the list of values that follows.
LOOKUP()	LOOKUP(search_for, lookup_range, result_ range)	LOOKUP() is a slimmed-down version of HLOOKUP() and VLOOKUP(). It searches for a value in one range, and returns the value in the same index position in the second range. Unfortunately, you've no way to enforce strict matching, which reduces the usefulness of this function.

INDIRECT() and ADDRESS(): Working with Cell References Stored As Text

INDIRECT() and ADDRESS() are two of Excel's strangest reference functions. They let you work with cell references stored as text.

The INDIRECT() function retrieves the content from any cell you specify. The twist is that you specify the cell using literal text (a string) to describe the location of the cell. So, where normally you'd add two cells this way:

```
=A1+A2
```

If you were to use INDIRECT() to refer to cell A1, you'd write the formula this way:

```
=INDIRECT("A1")+A2
```

Note the quotation marks around the cell name; they indicate that A1 is just a piece of text, not a cell reference. The INDIRECT() function examines this text and gives you the corresponding cell reference.

The obvious question is why would anyone bother to specify cell references as ordinary text? The most common reason is that you want to create an extremely flexible formula. So, for example, instead of hard coding your cell references (as in =A1+A2), you may want your formula to read *another* cell to find out which cell reference it should use. The end result is that the person reading the spreadsheet can change a cell reference used in a calculation, without needing to edit a complex formula by hand. Figure 12-5 shows a worksheet that demonstrates this trick. It lets the person using the spreadsheet choose the numbers to add, without forcing him to write a SUM() function by hand.

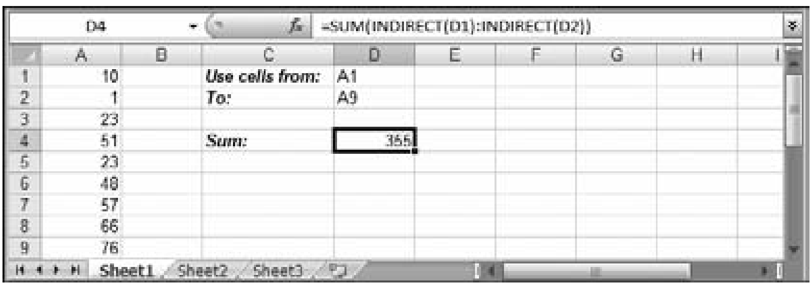

Figure 12-5: *This worksheet shows the quirky INDIRECT() function at work. Here, a dynamic SUM() formula is used to add a range of cells. The SUM() formula uses the INDIRECT() function to read the cell references stored in cells D1 and D2. It then uses these to create the range of cells that are added together.*

Here's an example of this technique:

```
=INDIRECT(A1)+INDIRECT(A2)
```

This formula does *not* add the content of cells A1 and A2. Instead, it looks in the cells A1 and A2 to determine which two cells it *should* add together. If the cell A1 contains the text "B2" and the cell A2 contains the text "D8", the formula becomes the following:

```
=INDIRECT("B2")+INDIRECT("D8")
```

which is the same as:

```
=B2+D8
```

This formula assumes that there are two cell references, entered as text, in cells A1 and A2. If there's something else in these cells, the INDIRECT() function doesn't work. If cell A1 contains a number, for instance, the formula INDIRECT(A1) returns the error code #REF! (which is Excel shorthand for, "I expected a cell reference, but all you gave me was a lousy number"). Figure 12-6 shows what you can—and can't—do with the INDIRECT() function.

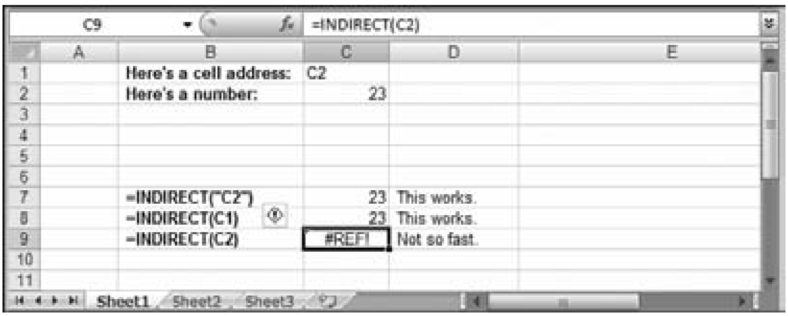

Figure 12-6: *This worksheet shows three attempts at using INDIRECT(). In C7 (whose formula's written out in B7), the cell reference (C2) is a hard-coded piece of text. The formula displays the value contained in cell C2. C8 has an extra layer of redirection. The formula grabs the content in C1, which is the piece of text "C2," and then supplies that to the INDIRECT() function. The INDIRECT() function then displays the content from cell C2, as in the previous example. The final =INDIRECT("C2") formula grabs the content from C2, which isn't a cell reference (it's the number 23). As a result, the INDIRECT() function accomplishes nothing.*

The ADDRESS() function performs a similarly strange operation, but in reverse. You supply ADDRESS() with row and column index numbers, and Excel gives you a piece of text that contains the cell reference. The following formula returns $B$1:

```
=ADDRESS(1, 2)
```

Remember, all the ADDRESS() function does is return a cell reference as a string. This piece of text isn't too useful; all you can really do with it is display it in a cell—or, if you're really warped, pass it to the INDIRECT() function to convert it back to a real cell reference.

Consider the example shown in Figure 12-4, which uses a lookup function to find the sales in a particular city on a particular month. The formula looked like this:

```
=INDEX(B2:D8, MATCH(DATE(2004,1,1),A2:A8,0), MATCH("London",B1:D1,0))
```

Instead of displaying the actual sales figures, you could use the ADDRESS() function to display the cell reference for the spreadsheet reader to see, like so:

```
=ADDRESS(MATCH(DATE(2004,1,1),A2:A8,0), MATCH("London",B1:D1,0))
```

This formula displays the text $C$4. If you want to get craftier, then you could add some additional text by modifying the formula like this:

```
="The number you are looking for is in cell " &
ADDRESS(MATCH(DATE(2004,1,1),A2:A8,0), MATCH("London",B1:D1,0))
```

This formula displays the text *The number you are looking for is in cell $C$4*. Clearly, you won't use a formula like this often, but it does raise some interesting possibilities.

If you don't want the ADDRESS() function to return a string with an absolute reference (page 254), then you can supply a third parameter, called *abs_number*. It takes one of the following values:

1 (Returns an absolute address, as in $B$1)

2 (Fixes the row number, as in B$1)

3 (Fixes the column, as in $B1)

4 (Returns a relative reference, as in B1)

TRANSPOSE(): Changing Rows into Columns and Vice Versa

TRANSPOSE() is an interesting function that you can use to change the structure of a table of data. Using TRANSPOSE(), you can invert the data so that all the rows become columns, and the columns become rows (Figure 12-7). In that respect, it works the same as the Transpose pasting option (page 107). The Paste Special approach, however, creates a distinct copy of the data. The TRANSPOSE() function, on the other hand, creates a linked table that's bound to the original data, which means that if you change the original table, the transposed table also changes. The TRANSPOSE() function is, therefore, ideal for showing more than one representation of the same data.

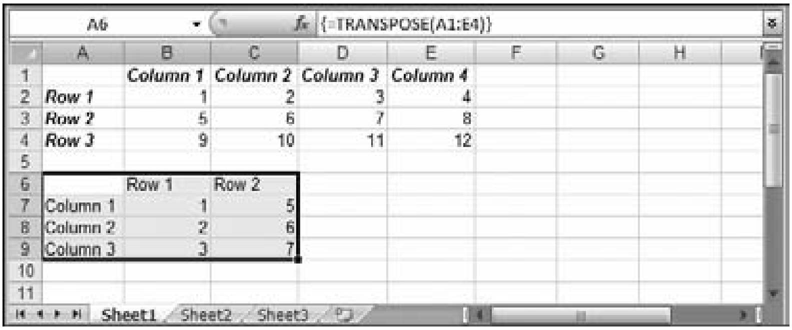

Figure 12-7: *In this worksheet, the TRANSPOSE() function has inverted most of the data in a table, turning columns into rows and rows into columns.*

The best way to get started with the TRANSPOSE() function is to try a simple example. Just follow these steps:

1. **Find a table of cells you want to transpose. Make a note of its size in rows and columns.**

   To determine the size of the transposed table, simply switch the number of rows and columns. That means a table made up of three rows and four columns becomes a table made up of four rows and three columns.

2. **Move to the area where you want to insert the transposed cells, and select the appropriately sized area.**

   If you want to transpose a table with three rows and two columns, then you need to select a grid of cells that's two rows high and three columns wide. You need to select all the cells before you begin because you'll be creating an array formula (page 253). (You need an array formula because the TRANSPOSE() function returns a whole table's worth of results.)

3. **Press the equal sign (=) and begin typing the TRANSPOSE() formula.**

   The TRANPOSE() formula requires one argument, which is the range of cells you want to transpose. Here's an example: *=TRANSPOSE(A1:E4)*. Don't press Enter yet.

4. **Commit the function by pressing Ctrl+Shift+Enter.**

   This step inserts the formula into all the selected cells as an array formula.

Tip: Editing array formulas can be a challenge. To remove or replace an array formula, you need to first select ***all*** the cells that use it. If you have a large array formula, this can be awkward. One handy-dandy shortcut is to use the Go To command to select the array formula cells. First, move to one of the cells that has the array formula. (You'll know you're in the right place when you see the curly brackets around the formula in the formula bar.) Press Ctrl+G to show the Go To dialog box. Next, click the Special button, choose "Current array," and then click OK. Presto: Excel selects the entire group of cells that uses this array formula.

The HYPERLINK() Function: Creating a Dynamic Link

The lookup and reference category holds an odd collection of functions. In theory, all the functions are unified by a common theme: They extract values or information from another cell. But every once in a while, a function turns up in the lookup and reference category because it isn't accepted anywhere else. The HYPERLINK() function is one such example.

The HYPERLINK() function has no real similarity to the other information and reference functions. It has a single purpose in life: to create an underlined hyperlink that someone can click. When clicked, this link can open another workbook, jump to somewhere else in the current workbook, or open a website URL in a browser window.

To use the HYPERLINK() function, you supply two function arguments: the target (like a web page, the path to a file, or a cell reference) and the text that you want to appear in the cell. Here's what the function looks like:

```
HYPERLINK(link_location, [cell_text])
```

Both the link_location and the cell_text arguments are text-based, so you need to use double quotes. Here's a formula that creates a hyperlink to a web page:

```
=HYPERLINK("http://www.prosetech.com", "Click to see company site")
```

The links that the HYPERLINK() function creates behave the same as the links you can create with Excel's Insert→Hyperlink command (page 90). This fact raises a reasonable question—why would you use a function to accomplish what you can polish off just as easily using Excel's ribbon? In most cases, the HYPERLINK() function doesn't add any benefit. However, the HYPERLINK() function really makes sense when you need to create *dynamic* links whose contents vary depending on what's in other cells.

Consider this formula:

```
=HYPERLINK(A20, "Click to see company site")
```

When you click this cell, the hyperlink sends you to whatever web address you've entered in cell A20. Imagine you repeat this formula in several different places. All you need to do is change one cell—A20—to update all the links.

The HYPERLINK() function also gives you the chance to create some more complex conditional formulas (for more on conditional formulas, see page 364). You may want to send the clicker to a different place depending on the content in other cells. Here's an example:

```
=IF(B2="Acme Company", HYPERLINK("http://www.acme.com", "Click to
see company site"), "")
```

This formula actually generates a quirky—but potentially useful—disappearing link trick. If cell B2 contains the text "Acme Company," the IF() function calls the HYPERLINK() function and inserts a new link to the company website. But if B2 contains something else—perhaps representing an invoice for another company website—a blank value appears in its place. It's not too often that you'll want to combine hyperlinks with these types of special effects, but when you do, the HYPERLINK() function is ready and waiting.

Information Functions

Excel features yet another group of functions closely related to the lookup and reference functions. Called *information functions*, they let you retrieve information about the type of content found in any particular cell you want to examine. You can find the full list of information functions in the Information group (choose Formulas→Function Library→More Functions→Information).

The "IS" Functions: Checking the Value Inside a Cell

The most important information functions are those that start with the word IS. These functions let you test whether a cell's blank, has numeric content, and so on.

The IS functions are Boolean functions, which means they give you a result of either TRUE or FALSE. On their own, the IS functions aren't too impressive. However, you can combine them with other conditional functions to make simple decisions. (The IF() function tests a condition, and then inserts one of two values based on whether the condition's true or false. For a refresher on the IF() function, see page 364.)

You could use the IF() function in combination with the ISNUMBER() function to avoid performing a calculation if a cell doesn't contain numeric content. Here's the formula you'd use:

```
=IF(ISNUMBER(A1), 10/A1, NA())
```

In this example, if the ISNUMBER() test returns TRUE (that is, if A1 contains a number), then Excel uses the first argument. That means the program performs the calculation *10/A1*, and displays the result in the cell. If cell A1 doesn't contain numeric content, the ISNUMBER() function returns FALSE, and Excel uses the second argument. In this case, Excel executes the NA() function, which displays the error code #N/A. Figure 12-8 shows several IS functions at work.

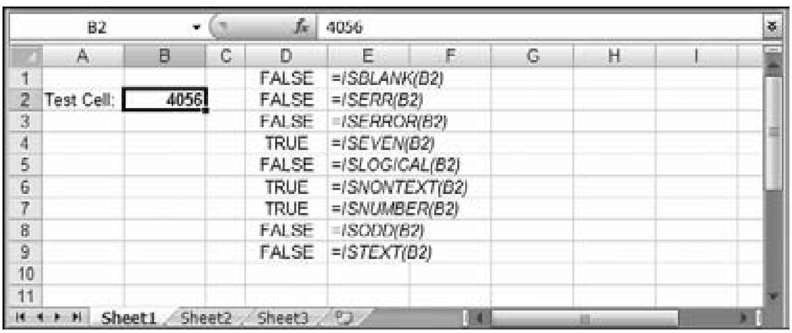

***Figure 12-8:** This worksheet shows the result of a series of tests on a single cell, B2. On their own, these TRUE and FALSE results may not impress you, but they represent building blocks you can use to create powerful conditional expressions.*

Here's another example that displays an error message when an error exists in a cell:

```
=IF(ISERROR(D10), "The calculation could not be completed due to an error.", "")
```

You could insert this formula in cell E10, so it appears right next to the cell that may or may not have an error (cell D10). If an error doesn't exist, E10 stays blank. If an error does happen, then E10 offers a little bit of extra information.

Excel includes adds a handy shortcut for error checking: the IFERROR() function. IFERROR() checks the indicated cell for an error. If there's no error, IFERROR() displays the value from that cell. But if the cell contains an error, IFERRROR() displays the alternate value that you supply as an argument. Here's an example that tests cell D10:

```
=IFERROR(D10, "The calculation could not be completed due to an error.")
```

IFERROR() may not save much typing, but it does make for more readable formulas.

Table 12-2 lists the IS functions. Usually, you'll use the IS functions with a literal value or a cell reference, although you can also use a range. When you use a range, the value is true only if the condition is satisfied in *all* the cells in the range. That means that *=ISBLANK(A1:A10)* returns TRUE only if all 10 cells in the range are empty.

Table 12-2. *The IS functions*

Function	Returns TRUE if…
ISBLANK()	The cell is empty, meaning it doesn't contain text, numbers, or any other content.
ISERR()	The cell contains an error other than #N/A.
ISERROR()	The cell contains any error, including #N/A.
ISEVEN()	The value is an even number.
ISLOGICAL()	The value is a Boolean (TRUE or FALSE) value.
ISNA()	The cell contains the #N/A error.
ISNONTEXT()	The value isn't text; it could be blank, a number, or a Boolean value.
ISNUMBER()	The value is a number, and not blank.
ISODD()	The value is an odd number.
ISREF()	The value is a cell reference. Thus, ISREF(A1) is true, but ISREF("A1") is not.
ISTEXT()	The value is text, and not blank.

TYPE() and ERROR.TYPE(): Finding a Value's Data Type or Error Type

Both the TYPE() and ERROR.TYPE() functions examine a cell and return a number that describes its content. You can use these functions to build conditional formulas (page 364).

The TYPE() function returns a number that represents the type of data in a cell. Possible numbers include:

1 (Number)

2 (Text)

4 (Logical Value)

16 (Error)

64 (Range)

The ERROR.TYPE() returns a number that represents the type of error that's happened in a cell. If the cell doesn't contain an error, the ERROR.TYPE() function returns #N/A. Possible numbers include:

1 (#NULL!)

2 (#DIV/0!)

3 (#VALUE!)

4 (#REF!)

5 (#NAME?)

6 (#NUM!)

7 (#N/A)

To learn how to use these kinds of functions to build conditional logic, see Chapter 13.

INFO() and CELL(): Gathering Info About Your Computer and Your Worksheet's Cells

INFO() and CELL() are two of the strangest information functions out there. While they're useful for stumping coworkers with Excel trivia, they don't solve many practical problems.

The INFO() function provides information related to the computer and operating system that's running Excel. To kick it into gear, you specify a single text argument, which tells Excel what information to retrieve. If you specify "numfile," Excel returns the number of all the worksheets in all the workbook files that are currently open:

```
=INFO("numfile")
```

Note: As you can see here, you need to enclose the argument value in quotation marks.

Table 12-3 lists all the possible arguments you can use with the INFO() function. Possible, of course, doesn't mean any of these make it into heavy rotation.

Table 12-3. *Argument values for the INFO() function*

Text	Gives you
directory	The current directory path. This path is usually the last directory you browsed to in the Open or Save dialog boxes, which means that it's not necessarily the path where Excel's stored the current spreadsheet.
numfile	The number of all the worksheets in all the workbook files that are currently open.
origin	The reference, as text, of the cell in the top-left corner of the currently visible area. The cell reference is prepended with the fixed text $A. This trick is useful for backward compatibility with ancient versions of the Lotus 1-2-3 spreadsheet program.
osversion	The version of your computer's operating system.
recalc	The current recalculation mode, which determines when formula results are refreshed. Chapter 13 (page 386) describes the different recalculation modes and how to set them.
release	The version of Excel.
system	The type of operating system—"pcdos" for a PC, or "mac" for a Macintosh.

The CELL() function performs a similar trick, except that it returns information about a cell. It requires two arguments: a string indicating the type of information you want, and a reference pointing to the cell you want to examine. Here's an example that retrieves the number format of the cell A2:

```
=CELL("format", A2)
```

If you omit the cell reference, the CELL() function operates on the cell where the formula exists. The CELL() function is useful primarily for maintaining compatibility with other software. People occasionally use it in add-ins or macros that automate tasks in a worksheet, although most of its features are duplicated by other, more straightforward information and reference functions like ROW(), COLUMN(), INDIRECT(), and ADDRESS(). One interesting piece of information that you can retrieve from CELL(), but not from these other functions, is the file name (including the full file path) of the current spreadsheet:

```
=CELL("filename")
```

For a full list of values you can use with the CELL() function, refer to the Excel help reference—or just start typing. As soon as you add the CELL() function to a formula and type the opening parenthesis, Excel pops up a list of possible values that you can choose from.

Tip: The CELL() and INFO() functions don't refresh automatically. If you need to recalculate these formulas, then press F9 to update the worksheet.

Tutorial: Generating Invoices from a Product Catalog

The reference, lookup, and information functions really shine when you want to build worksheets that automatically carry out a series of boredom-inducing tasks.

And they let you tie your information together in interesting ways. Maybe you want to determine monthly wages based on billing rates, build a parts list for an order, or create an invoice based on items that have been ordered from a product catalog. This last task, which is shown in the next example, is sometimes called a *dynamic invoicer*. It lets sales representatives quickly create an invoice by choosing from a catalog of company products.

The dynamic invoicer employs the following techniques:

- Looking up information from a table with VLOOKUP().
- Avoiding calculations when data isn't present with IF() and ISBLANK().
- Referencing data from another worksheet.
- Using absolute references to make it easier to copy formulas.

First, you build the table filled with your company's products. Figure 12-9 shows you the list for this example.

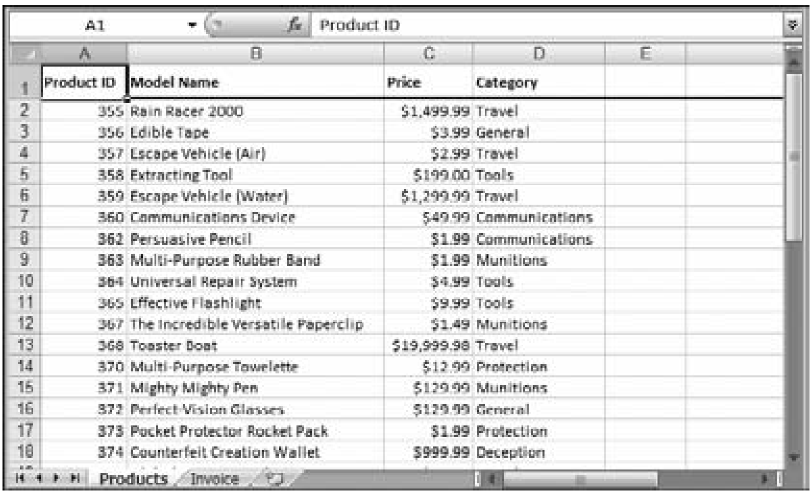

Product ID	Model Name	Price	Category
355	Rain Racer 2000	$1,499.99	Travel
356	Edible Tape	$3.99	General
357	Escape Vehicle (Air)	$2.99	Travel
358	Extracting Tool	$199.00	Tools
359	Escape Vehicle (Water)	$1,299.99	Travel
360	Communications Device	$49.99	Communications
362	Persuasive Pencil	$1.99	Communications
363	Multi-Purpose Rubber Band	$1.99	Munitions
364	Universal Repair System	$4.99	Tools
365	Effective Flashlight	$9.99	Tools
367	The Incredible Versatile Paperclip	$1.49	Munitions
368	Toaster Boat	$19,999.98	Travel
370	Multi-Purpose Towelette	$12.99	Protection
371	Mighty Mighty Pen	$129.99	Munitions
372	Perfect-Vision Glasses	$129.99	General
373	Pocket Protector Rocket Pack	$1.99	Protection
374	Counterfeit Creation Wallet	$999.99	Deception

***Figure 12-9:** Here's a product catalog that a sales rep might use to build an invoice. It includes the Product ID (used for the lookup), the Model Name (used to identify the selected products), and the Price (used to calculate the total cost of an order).*

Next, you build a worksheet that actually *creates* the invoices. You should probably include space for entering customer information, followed by a large area for the list of items that have been ordered, as shown in Figure 12-10.

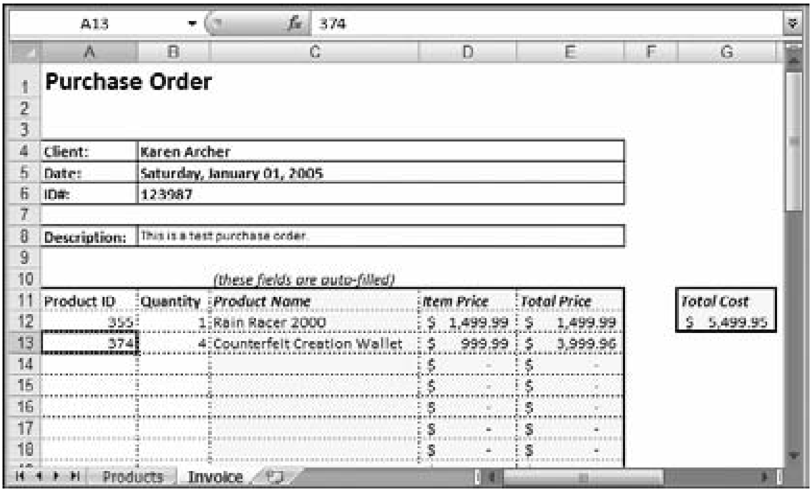

Figure 12-10: *To add items to this invoice, the sales rep simply types in the product ID and quantity. The lookup functions retrieve the model name and unit price, and then determine the total price.*

Here, your best approach is to retrieve an item from the product catalog based on the product ID; this prevents people using spreadsheet having to type out full product names. Therefore, you need to write a lookup function to scan the product table for a matching ID. And since the invoice needs to record the product name and its price, you need two lookup functions. Here's the lookup that retrieves the price for the first item in the invoice:

```
=VLOOKUP(A12, Products!$A$2:$D$42, 3, FALSE)
```

The lookup for the model name is the same, but the column offset is 2 instead of 3:

```
=VLOOKUP(A12, Products!$A$2:$D$42, 2, FALSE)
```

The reference to the table of products is preceded by the word *Products* and an exclamation mark. That's because the products table isn't on the invoice worksheet. Instead, it's in a different worksheet in the same workbook named Products. (Remember, in order to perform a calculation with data in another worksheet, you just need to preface your cell reference with the name of the worksheet, followed by the exclamation mark. For a refresher on how functions interact across worksheets and workbooks, see page 257.)

You'll also notice that this function uses a mix of absolute and relative references. The reference for the product ID is relative (A12) because each line item in the invoice has its own product ID, which is what the sales rep enters to start filling up the purchase order. The product table is a range made up of two absolute references *($A$2:$D$42)*. This range ensures that as you copy the formula down to subsequent rows on the invoice table, each copy refers to the same product catalog.

As written, the function still suffers from one problem. If you copy the formula to all the available rows in the invoice, you'll notice that the value #N/A appears in each row where you haven't entered a product ID. This error message happens because the VLOOKUP() function attempts to perform a lookup for a product with a product ID of 0, which doesn't exist. The #N/A error message is a significant problem because it prevents you from calculating the total cost of the order with the SUM() function. If the SUM() function attempts to add together a range of cells that includes a #N/A value, it simply returns the #N/A error code.

To solve this problem, you must use conditional logic to test if a product ID has been entered. If it hasn't, you can simply put a blank value in the price column. Here's the corrected formula:

```
=IF(ISBLANK(A12), 0, VLOOKUP(A12, Products!$A$2:$D$42, 3, FALSE))
```

You can now build an invoice in seconds, just by typing in a few product ID numbers in the leftmost column. The products appear automatically, and you can calculate the overall total with a simple formula that uses the SUM() function:

```
=SUM(E12:E100)
```

Note: This dynamic invoicer's only limitation is that the sales rep must know which product ID codes to use (or must look them up on the first worksheet). An even more impressive worksheet could automate this process with some sort of list box control. You'll see how to take this step on page 703.

CHAPTER
13

Advanced Formula Writing and Troubleshooting

Over the last five chapters, you've learned how to use Excel's impressive function toolkit to calculate everything from statistical trend lines to a payment schedule for a home mortgage. Now that you've had a close look at these functions, it's time to consider a few techniques to help you get the most from your formulas.

In this chapter, you'll tackle four new topics that'll make you a formula master:

- How to use conditional logic with functions like SUMIF() and COUNTIF().
- How to make formulas more readable by using named ranges.
- How to control when Excel recalculates your worksheets.
- How to solve mysterious errors by using Excel's formula auditing tools.

Conditions in Formulas

Chapter 8 gave you a first look at how to use conditional logic when writing Excel formulas. The basic principles are easy: you construct a condition using the logical operators like <, >, =, and <>, and then use this condition with a *conditional function*. So far, you've considered only one conditional function—IF()—which performs different actions depending on the result of a calculation.

Here's an example that uses the IF() function. The following formula carries out the operation in either the second or third argument, depending on the value of cell A20:

```
=IF(A20>10000, A20*5%, A20*3%)
```

Translation: For values greater than 10,000, Excel executes the formula *A20*5%*; otherwise, it carries out the second formula. If A20 contains the dollar amount of a sales invoice, you can use this formula to determine the commission for a sales person. If the sale exceeds the magic $10,000 amount, a larger 5-percent commission kicks in.

IF(): Building Conditional Formulas

IF() is one of the most useful conditional functions, but it's not the only one. In fact, you'll find several more logical functions in the Formulas→Function Library→Logical category. These formulas include the IFERROR() function you tackled on page 357, and the following:

- **TRUE() and FALSE()**. These functions don't take any arguments. The TRUE() function simply returns the logical value *true*, and the FALSE() function returns *false*. You can use these functions to enter logical values directly into a cell or formula, but you won't need them very often.
- **AND()**. This function accepts two (or more) conditions, and then returns *true* if all of them are true. If any condition's false, the AND() function returns *false*.
- **OR()**. This function accepts two (or more) conditions, and then returns *true* if any one of them is true. The OR() function returns *false* only if all conditions are false.
- **NOT()**. This function accepts a condition and reverses it. Consider the formula *=NOT(ISBLANK(A10))*. Assuming A10 is empty, the ISBLANK() function returns *true;* by contrast, the NOT() formula returns the reverse—*false*. You rarely need the NOT() function, but sometimes it helps make your logic a little clearer.

On their own, these logical functions don't do much. However, you can combine them in some interesting ways with the IF() function. Imagine you want to make sure the 5-percent rate (from the sales commission example) kicks in only if the sale exceeds $10,000 *and* the sale happens after the year 2006, when the new commission rules came into effect. You can rewrite the earlier formula to take this fact into account using the AND() function. If the invoice date's in cell B4, here's what the formula would look like:

```
=IF(AND(A20>10000,YEAR(B4)>2006), A20*5%, A20*3%)
```

Similarly, you may encounter a situation where you want to alter your logic so that the higher commission rate kicks in if at least one of two different criteria is met. In this case, you'd use the OR() function.

Tip: The information functions described in Chapter 12 are also quite useful in conditional statements. The "IS" functions, like ISERROR() and ISBLANK(), especially lend themselves to these statements.

You can also choose between more than two options by *nesting* multiple IF() statements in a single formula. (Nesting is a technique that lets you put one function inside another.) The following formula uses a commission of 2 percent if the sale is under $500, a commission of five percent if the total is above $10,000, and a 3 percent commission for all other cases.

```
=IF(A20<500, A20*2%, IF(A20>10000, A20*5%, A20*3%))
```

The formula begins by checking the first condition (whether A20 is less than $500). If the condition is met, then Excel carries out the first expression, *A20*2%*. If A20 is *not* less than $500, Excel moves on to the second expression, which is actually another IF() function. It then checks the second condition and chooses between the two remaining commission levels.

Excel allows up to a staggering 64 nested IF() statements in one formula. (This limitation actually applies to all Excel functions. You can't nest functions more than 64 levels deep.) However, it's unlikely that anyone could actually understand a formula with 64 IF() statements. Nesting IF() statements may get the job done, but it can lead to some extremely complicated formulas, so tread carefully.

Note: In some situations, nesting multiple IF() statements may cause the formula to become too complex and error-prone. In these cases, you may want to consider simplifying your logic by breaking it into several linked formulas, or building a custom function using full-fledged VBA (Visual Basic for Applications) code. You'll see an example of that in Chapter 29.

Along with the functions in the Logical category, Excel also includes a few more functions that use conditions. These functions include the COUNTIF() and SUMIF() functions, as described in the next two sections.

COUNTIF(): Counting Only the Cells You Specify

To understand the purpose of COUNTIF() and SUMIF(), you need to remember that the COUNT() and SUM() functions devour everything in their path. But what if you want to pick out *specific* cells in a range and count or sum only these cells? You might try to use COUNT() or SUM() in conjunction with an IF() statement, but that doesn't solve the problem. Consider the following formula:

```
=IF(ISBLANK(A1), 0, COUNT(A1:A10))
```

This formula checks to see if cell A1 is blank. If it's blank, the formula returns 0. If it isn't blank, the formula returns the count of all the numeric values from A1 to A10. As you can see, the actual counting operation's an all-or-nothing affair. Excel either counts all of the cells or ignores them all. There's no way to count just some of the cells (or, similarly, add just some of the cells). The COUNTIF() and SUMIF() functions address this issue by letting you specify a condition for *each cell*.

The COUNTIF() function is the more straightforward of the two. It takes two parameters: the range of cells you want to count, and the criteria that a cell needs to satisfy in order to be counted:

```
COUNTIF(range, criteria)
```

The criteria argument is the key to unlocking the real power of COUNTIF(). The formula tests every cell in your range to see if it meets your criteria, and counts it *only* if it does. With the criteria you can:

- Test if a cell matches a specific value.
- Test if a cell is greater or less than a specific number.
- Test if a cell matches, or is greater or less than, a number in another cell.
- Test if a cell contains text that matches a simple pattern.

Consider the list of products shown in Figure 13-1, which stretches from row 2 to 42. Counting the number of products is easy—you just use the plain-vanilla count function:

```
=COUNT(A2:A42)
```

This formula returns *41*, which is the total number of nonblank cells in the range. Now, what if you want to count the number of products with a price over $500? This challenge calls for the COUNTIF() function. Here's the formula:

```
=COUNTIF(C2:C42, ">500")
```

Note that, in this case, the formula counts the cells in column C. That's because column C contains the price information for each product, which you need in order to evaluate the condition. (When using the COUNTIF() function, the condition in the second argument's always a string, which means you need to make sure to place it inside quotation marks.)

Figure 13-1: *This worksheet shows a list of products. The COUNT() function makes it easy to count the products, but COUNTIF() gives you the additional ability to count only those products that reach a set price threshold or are in a certain category.*

To understand how the criteria argument works, you need to realize that it's *not* a logical condition like the ones used with the IF() statement. Instead, it's a snippet of text that contains *part* of a condition. When COUNTIF() springs into action, it creates a full condition for each cell in its assigned range. In the formula shown earlier, the criteria is *>500*. Each time the COUNTIF() function tests a cell, it uses this criteria to generate a full-fledged condition on the fly. The first cell in the range is C2, so the condition becomes *C2>500*. If the condition is true (and for C2, in Figure 13-1, it is), COUNTIF() counts the cell.

Using this logic, you can easily construct other conditional counting formulas. Here's how you'd count all the products in the Travel category (column D) by examining the cells in the Category column:

```
=COUNTIF(D2:D42, "=Travel")
```

Tip: If your condition uses the equal sign, you can omit it. For example, COUNTIF() assumes that the condition *"Travel"* is equivalent to *"=Travel"*.

You can even draw the information you want to use in your condition from another cell. In this case, you simply need to use the text concatenation operator (&) to join the cell value with the conditional operator you want to use. (See page 316 for an explanation of how concatenation works.) If the reader of your spreadsheet enters the category name in cell G1, you could count matching products using the following formula:

```
=COUNTIF(D2:D42, "=" & G1)
```

This formula joins the equal sign to whatever text is in cell G1. Thus, if G1 has the text Tools, the criteria becomes *"=Tools"*.

You can use a similar technique to use a function in the criteria argument. Here's a formula that counts the number of products that are above the average price:

```
=COUNTIF(C2:C42, ">" & AVERAGE(C2:C42))
```

SUMIF(): Adding Only the Cells You Specify

The SUMIF() function follows the same principle as COUNTIF(). The only difference is that it accepts an optional third argument:

```
SUMIF(test_range, criteria, [sum_range])
```

The first argument is the range of cells you want the criteria to test, the second is the criteria itself, and the third is the range of cells you want to sum. So if the first cell in the *test_range* passes the test (that is, causes the criteria to evaluate to *true*), the function adds whatever the first cell is in the *sum_range* to the total sum. Both the *test_ range* and *sum_range* must have the same number of cells—usually they'll be different columns in the same table.

SIDEBARTYPE

How Excel Compares Different Types of Data

The type of data that's in a cell—ordinary numbers, text, or dates—influences how Excel compares it in a condition. If you use the greater than (>) and less than (<) symbols with numbers, Excel performs the expected numeric comparison. However, if you use these symbols with text, Excel performs a letter-by-letter alphabetic comparison, which means Excel considers the word *apple* less than *banana* because the first letter, a, comes before b in the alphabet. Comparisons are always case-insensitive, so *tRaVeL* matches *Travel*.

Date and time comparisons work the same as numeric comparisons, although it helps to remember that Excel stores dates as a number of days (see page 70 for an explanation) and times as a fractional number between 0 and 1. In other words, Excel deems that dates or times in the past are smaller than more recent dates in a comparison.

Finally, Excel lets you use the asterisk (*) wildcard when you're making text comparisons. The asterisk stands for any sequence of characters. Thus, the criteria *"=T*"* matches any cell that contains text that starts with the letter T. Here's how you'd use it with the COUNTIF() function:

```
=COUNTIF(D2:D42,"=T*")
```

This formula matches any category that starts with the letter T, which includes Travel and Tools.

Imagine you want to calculate the sum of all the products in the Travel category, as shown in Figure 13-1. In this case, you'd want to *test* the cells from the Category column (D2:D42) to see whether they contain the word Travel; these cells constitute your test range. But then you'd want to *add* the cells from the Price column (C2:C42), and those cells would be your sum range. Here's the formula:

```
=SUMIF(D2:D42, "=Travel", C2:C42)
```

The two ranges (D2:D42 and C2:C42 in this example) need to have the same number of cells. For each cell Excel tests, that's one cell it may or may not add to the total.

If you omit the *sum_range* argument, the SUMIF() function uses the same range for testing and for summing. You could use this approach to calculate the total value of all products with a price above $500:

```
=SUMIF(C2:C42, ">500")
```

Along with COUNTIF() and SUMIF(), Excel has one more IF-formula: AVERAGEIF(), which calculates the average of cells that fit set criteria. Figure 13-2 shows an example.

COUNTIFS() and SUMIFS(): Counting and Summing Using Multiple Criteria

The COUNTIF(), SUMIF(), and AVERAGEIF() functions suffer from one limitation: They can evaluate cells using only one criterion. Fortunately, Excel has a set of matching functions that don't have this restriction. They have the same names as the functions you've already learned about, with an *S* added at the end:

- **SUMIFS()** is the same as SUMIF(), except it accepts multiple conditions.

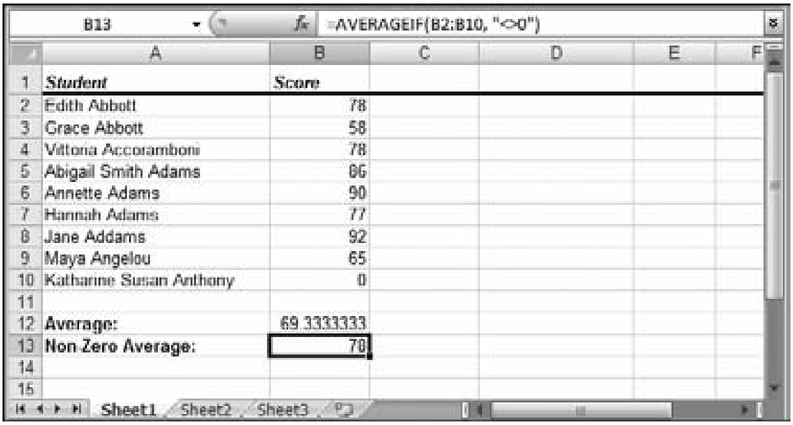

Figure 13-2: *Here are two approaches to calculate an average. The formula in B13 ignores all scores of 0, which represent students who didn't take the test. The result of 78 percent (returned by the AVERAGEIF() function) is clearly quite a bit higher than the value of 69 percent and change (returned by the AVERAGE() function), but it's a more realistic representation of student performance.*

- **COUNTIFS()** is the same as COUNTIF(), except it accepts multiple conditions.
- **AVERAGEIFS()** is the same as AVERAGEIF(), except it accepts multiple conditions.

If you're truly demanding, you'll be happy to know that all three of these functions accept over 100 separate conditions.

Earlier in this chapter you saw how you could use COUNTIF() either to count the number of products over $500 or to count the number of travel products in the list. But what if you want to use *both* conditions at once—in other words, you want to zero in on just those travel products that pass the $500 threshold? In this situation, you need COUNTIFS() to evaluate your two conditions. The following explanation builds the formula you need one piece at a time. (See Figure 13-1 if you need help visualizing the spreadsheet that this formula is searching.)

The first argument for COUNTIFS() identifies the range you want to use to evaluate your first condition. In this example, this first condition tests whether a product's in the Travel category. In order to test this condition, you need to grab the entire Category column, like so:

```
=COUNTIFS(D2:D42, ...)
```

Now, you need to fill in the condition that Excel uses to test each value. In this case, you need a condition that checks that the value matches the text "Travel":

```
=COUNTIFS(D2:D42, "=Travel", ...)
```

The fun doesn't stop here. You can use the same technique to fill in the second condition, which looks for prices that exceed $500 dollars. Once again, you fill in the range you want to use (this time it's the Price column) and the condition:

```
=COUNTIFS(D2:D42, "=Travel", C2:C42, ">500")
```

This COUNTIFS() function counts only products that meet both conditions. However, you can continue this process by tacking on more and more range and condition arguments to make ridiculously stringent conditions.

Note: When you use the COUNTIFS() function, all your ranges need to be exactly the same size. (In other words, they need to have the same number of cells.) If you break this rule, Excel becomes terribly confused and shows you the *#VALUE!* error. After all, the idea is that you're looking at different parts of the same list, and it wouldn't make sense for one column to be longer than another one in the same table.

The SUMIFS() function works in a similar way. It's different because the range you want to sum up may not match the ranges you want to use to evaluate your conditions. To clear up any confusion, you add an extra argument right at the beginning of the formula, which identifies the cells you want to add up.

The following formula calculates the total of all products between $500 and $1,000—a feat that's impossible with SUMIF().

```
=SUMIFS(C2:C42, C2:C42, ">500", C2:C42, "<1000")
```

Notice that, in this example, the ranges you're using for summing and testing the two conditions are all the same.

The COUNTIFS() and SUMIFS() functions offer a nifty way to create reports. See the two tables in Figure 13-3 for an example. The first table (rows 1 through 8) has the source data, which is a list of dates that various contractors worked, and the total number of hours they logged. The second table (rows 12 through 14) uses the SUMIFS() function to figure out things like how many hours each programmer's spent programming or testing.

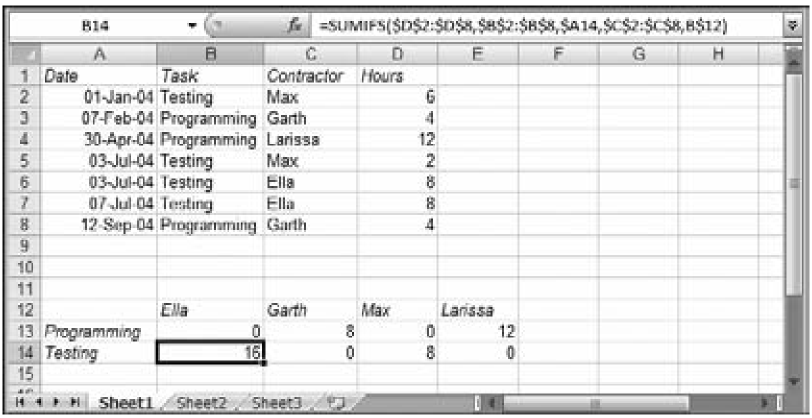

Figure 13-3: *SUMIFS() can be a godsend when boiling down large amounts of information. In the highlighted cell shown here, the formula returns the total number of hours from rows that match two conditions, indicating that Ella was testing for a total of 16 hours.*

Cell B14 answers the question: How many hours has Ella spent on testing? The following SUMIFS() formula solves this question using two conditions. First, it scans column C in order to find all rows that have Ella's name. Then, it scans column B to determine what type of work Ella did. Here's the final formula:

```
=SUMIFS(D2:D8, B2:B8, A14, C2:C8, B12)
```

Note that the two conditions are drawn from other cells. The first condition (which acts on cells B2:B8) matches the value in cell A14, which has the task type (in this case, "Testing"). The second condition (which acts on cells C2:C8) matches the value in B12, which has Ella's name.

Descriptive Names for Cell References

One of the obvious problems with Excel formulas is that they don't make for the easiest reading in the world. Consider this formula:

```
=(A1+A2)*B1
```

It's immediately obvious that this formula adds together two numbers (the numbers in cells A1 and A2), and then multiplies the result by a third number (the number in cell B1). However, the formula gives absolutely no clue about the *purpose* of this calculation. There's no way to know whether it's converting currencies, calculating a discount, or measuring the square footage of your llama day-care center. To answer these questions, you need to look at the worksheet and track down the cells this formula references.

On the other hand, consider the next formula, which uses descriptive names in place of cryptic cell references. Although it performs the same calculation, this formula provides much more information about what it's actually trying to do—calculate the retail price of a product:

```
=(ProductCost + ShippingCost) * MarkupPercentage
```

Excel lets you build formulas using such descriptive names, or *named ranges*. All you have to do is define the ranges as you create your worksheet, and then you can use these names instead of cell references.

Named ranges provide other benefits besides conveying the meaning of a formula:

- They make complex, nested formulas more understandable.
- They make it easy to quickly find a cell or select a group of cells. That makes them ideal for navigating large worksheets or for applying formatting to cell ranges that frequently change.
- They reduce the likelihood of some types of errors. For example, you're unlikely to notice if you've used A42 instead of A41, but you can spot the wrong name when you use TotalWithTax instead of TaxRate.

- They use absolute references. That way, you don't need to worry about having a formula change when you copy a formula from one cell to another. (Although 99.9 percent of all Excel fans use absolute references for named ranges, this Excel rule is one you can break for some unusual tricks. To learn about this technique and other unusual tricks, check out *Excel Hacks* by David and Raina Hawley.)
- They add an extra layer between your formulas and your worksheet. If you change the structure of your worksheet, then you don't need to modify your formulas. Instead, you simply need to edit the named ranges so they point to the new cell locations. This ability is particularly useful when writing macros (Chapter 28), because it lets you avoid writing direct cell references into your code.

In the next few sections, you'll learn how to create named ranges. You'll also pick up a few tricks for defining and applying names automatically.

Creating and Using a Named Range

Creating a named range is easy. Just follow these steps:

1. **Select the cells you want to name.**

   You can name a single cell or an entire range of cells.

2. **Look for the box at the left end of the formula bar, which indicates the address of the current cell (for example, C5). This part of the formula bar is called the *name box*. Click the name box once.**

   The text inside (the cell address) is now selected.

3. **Start typing the name you want to use.**

   What you type replaces the cell address (see Figure 13-4).

4. **Press Enter to confirm the new name.**

   You can now select the name at any time (see Figure 13-5). To use a name, click the drop-down arrow to the right side of the name box to show a list of all the names defined in the workbook. When you click a name, Excel jumps to the appropriate position and selects the corresponding cell or range of cells.

Note: When you name a range, you have to follow certain rules. All names have to start with a letter, and they can't contain spaces or special characters–except the underscore (_). Also, all names have to be unique. Finally, you can't create a range name that matches a valid cell address. For example, BB10 isn't a valid range name, because every worksheet has a cell with the address BB10.

Once you've created a name, you can use the name in any formula, just as you would any cell reference or cell range. As a shortcut, you can pick the name from a handy list. To do so, start entering your formula, and then select Formulas→Defined Names→Use in Formula. This selection pops open a menu with all the names you've defined in your workbook. Select the name you want to use, and then Excel inserts it into the formula.

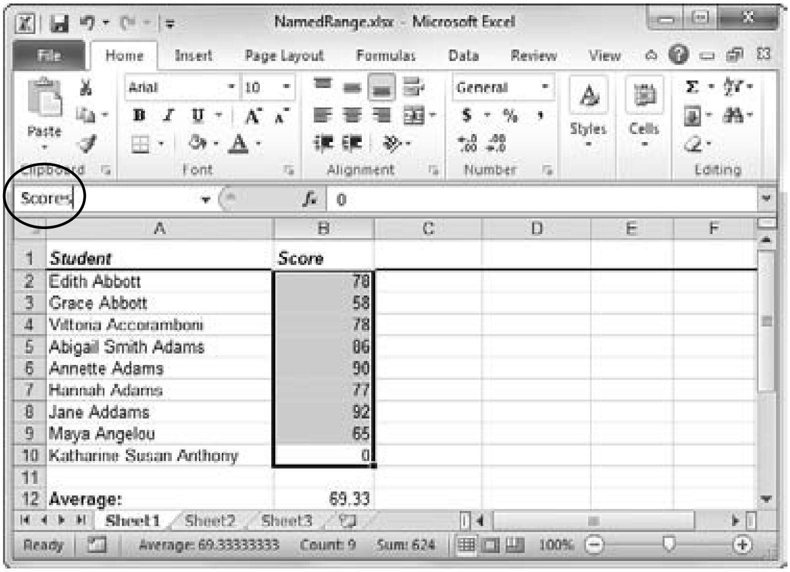

Figure 13-4: *This worksheet has a named range so you can easily use the student grades in your formulas without remembering specifically where the cells are. (It also lets you change the range and update all your formulas in one step, if the worksheet changes.) To define the range for student scores, simply select the cells with the numeric scores (B2 to B10), and then enter the new name using the name box in the formula bar (circled).*

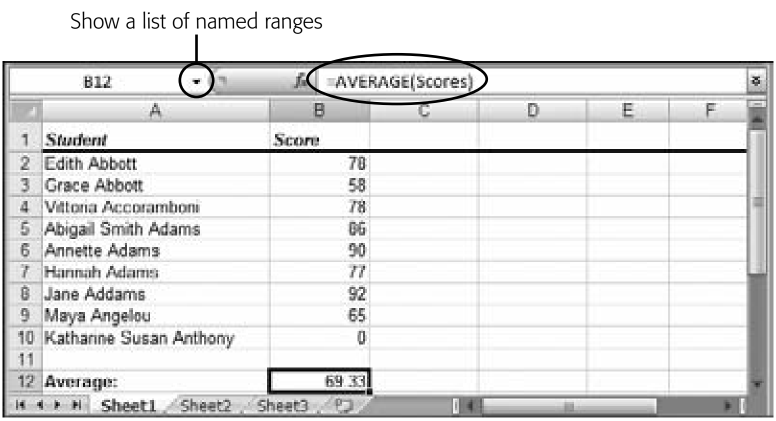

Figure 13-5: *Once you've defined and named a cell range, you use it in formulas (as shown with the AVERAGE() function here). To see a list of all the named ranges in your worksheet, click the drop-down arrow in the box on the left side of the formula bar. (If you pick a name from this list, Excel selects the corresponding cells.)*

Creating Slightly Smarter Named Ranges

The name box gives you a quick way to define a name. However, you have other ways to supply a few optional pieces of name-related information, like a description. To create a name with this optional information:

1. **Select the cells you want to name.**

2. **Choose Formulas→Defined Names→Define Name.**

   Or, without using the ribbon, just right-click the selected cells, and then choose Name a Range. Either way, the New Name dialog box appears (Figure 13-6).

Figure 13-6:
In the New Name dialog box, you fill in the standard name and cell reference information, along with two extra details: the ever-important Scope setting (which lets you organize names in workbooks with multiple worksheets) and a Comment setting (which lets you add descriptive text that explains your name).

3. **Fill in the information for your named range.**

   The only information you need to add is the name, but you can tweak several settings:

   - **Name** is the name you're giving to the range (like TaxRate).
   - **Scope** lets you control where you can use the name. Ordinarily, names have Workbook scope, which means you can use them on any worksheet in your workbook. However, you may choose to limit the scope to a single, specific worksheet instead. This way, your name lists aren't as cluttered, because you see only the names for the current worksheet. Also, you don't need to worry about using the same name in different worksheets. Worksheet-scoped names are almost always the best way to go.
   - **Comment** is a description for your named range. This description appears in a tooltip when you enter the name into a formula. You can use any text you want for the comment, but Excel experts use it to describe what type of information's in the cells, when the data was last updated, and whether the range depends on information in other workbooks.
   - **Refers to** is a cell reference that indicates the cells you selected in step 1. If you didn't get it quite right, then you can modify the range now.

4. **Click OK to create the name.**

   You can now use your name in formulas anywhere. In fact, Excel even helps you out with its Formula AutoComplete feature, as shown in Figure 13-7.

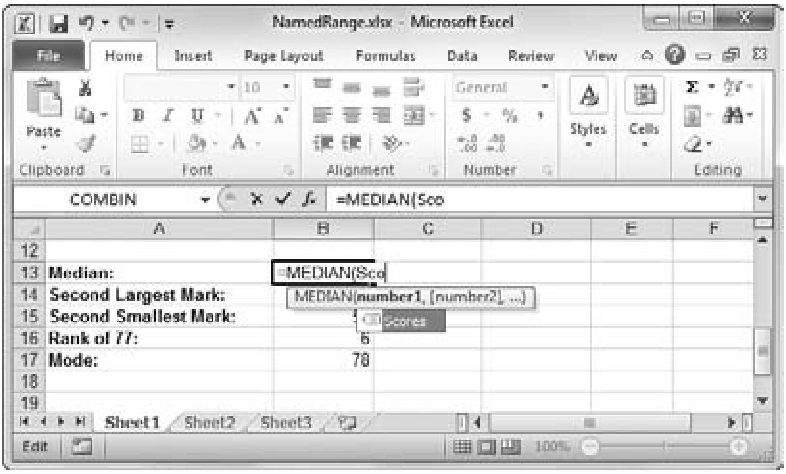

***Figure 13-7:** Excel's AutoComplete feature doesn't just suggest possible function names–it also shows you the named ranges that match the text you've typed so far. And if your name has a comment, you see that text appear in a tooltip to help you out.*

Naming Formulas and Constants

The New Name dialog box also provides the key to unlock additional naming features. Using this dialog box, you can create names that point to things other than cell ranges. You can create nicknames for frequently used formulas (like "My_ Net_ Worth"). This section describes how.

Internally, Excel treats all names as formulas. That is, when you create a named range, Excel simply generates a formula that points to that range, like *=$A$1:$A$10*. To check this out, select a few cells, and then choose Formulas→Defined Names→Define Name. In the "Refers to" text box at the bottom of the window, you see a formula with the corresponding cell reference.

Although most names refer to cell ranges, there's no reason why you can't create names with different formulas. You can use this approach to define a fixed constant value that you want to use in several formulas. Here's an example:

```
=4.35%
```

To create this constant, enter the name in the text box at the top of the New Name dialog box (for example, SalesCommission), enter the formula in the "Refers to" text box, click Add, and then click OK. You can now use this name in place of the constant in your worksheet calculations:

```
=A10*SalesComission
```

People often use named constants to make it easier to insert frequently used text. You may want to declare a company name constant:

```
="Acme Enterprises, Incorporated."
```

You can then use that text in a number of different cells. And best of all, if the company name changes, you simply need to update the named formula, no matter how many times it's used on the worksheet!

Similarly, you can create names that reference more complex formulas. These formulas can even use a combination of functions. Here's an example that automatically displays the name of the current day:

```
=TEXT(TODAY(), "dddd")
```

Tip: Overall, you'll probably find that the most useful type of name is one that refers to a range of cells. However, don't hesitate to use named constants and named formulas, which are particularly useful if you have static text or numbers that you need to use in multiple places.

Managing Named Ranges

It's all well and good to create named ranges, but sooner or later you'll need to tweak your handiwork by doing things like deleting names you don't need anymore, or editing the ones you regularly use (so that they designate a different area, for example). For these tasks, you need to choose Formulas→Defined Names→Name Manager to show the Name Manager (Figure 13-8) or press Ctrl+F3.

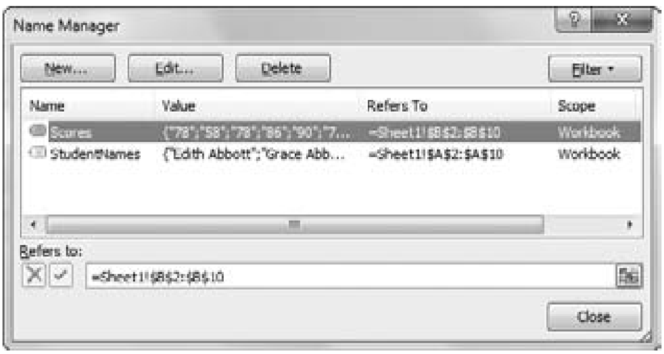

Figure 13-8: *The Name Manager shows you a list of all the names in your workbook. For better organization, you can click a column to re-sort the list. You can also resize the window to see more names at once.*

The Name Manager is the starting point when you need to add, delete, or edit existing named ranges. Here's what you can do:

- **To add a name**, click New. You get to the familiar New Name dialog box.
- **To remove a name**, select it in the list, and then click Delete. Any formulas using that name display the error code *#NAME? or #REF!*, indicating that Excel can't find the named range.

- **To edit a name**, select it in the list, and then click Edit. You see an Edit Name dialog box, which looks exactly the same as the New Name dialog box. Excel's nice enough to keep track of names. If you rename TaxRate to TaxPercentage, Excel adjusts every formula that used TaxRate to now use TaxPercentage, thereby avoiding an error.

Tip: If you need to change the cell reference for an existing name, but you don't need to change anything else (like its name or comment), Excel's got a quicker way. Just select the name in the list, and then modify the cell reference in the "Refers to" box at the bottom of the window. (Click-lovers can change the cell reference using the mouse. Just click inside the "Refers to" box to highlight the cell reference on the worksheet. Then, drag on the worksheet to draw the new cell reference.)

The Name Manager is a great place to review all the names you're using, check for errors, and make adjustments. Many Excel experts create workbooks that have hundreds of names. If you find yourself in this situation, then you may have a hard time finding the names you want in the list of names. Fortunately, the Name Manager includes a filtering feature that can help.

Filtering cuts down the name list so that instead of showing all the names in your workbook, it shows only the names you're interested in. The Name Manager lets you use several types of filters:

- You can find names that are limited to the current worksheet.
- You can find names that point to cells with errors (and those that don't).
- You can find ordinary names, or names used in tables (Chapter 14).

Figure 13-9 shows the filter settings you can choose from.

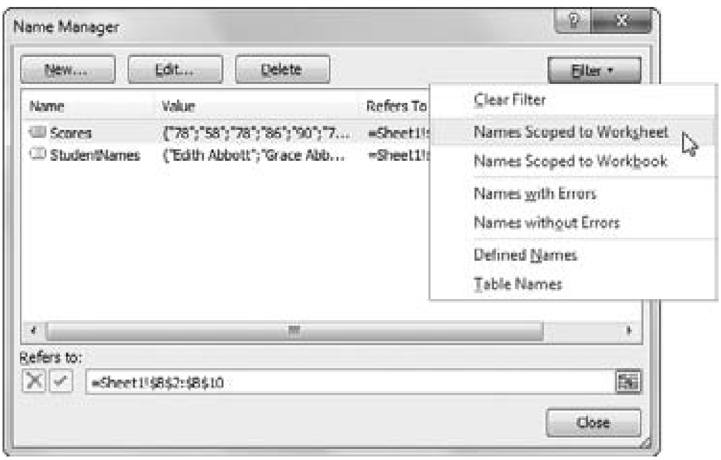

Figure 13-9: *To apply a filter, click the filter button, and then make a choice from the drop-down menu. You'll notice that this menu has several sections, separated by horizontal lines. Each section represents a separate type of filter. If you want, you can apply a filter from more than one section at once. You can look for names that are limited to the current worksheet and point to cells with errors, as in this example. To get back to normal and show all names, choose Clear Filter.*

GEM IN THE ROUGH

Getting a List of Names

After defining dozens of different names on a worksheet, you may appreciate an easy way to review them. The Name Manager is a great tool while you're using Excel, but there's another choice that lets you look over your names from the comfort of your armchair. You can paste a list of all your names into a worksheet, complete with their cell references. (Typically, you don't keep this list in your worksheet. Instead, you create it, print it out, and then remove it from the worksheet.)

To generate the list of names, click the cell where you want to start the list. Remember, you need two free columns—one for the range names and one for the range addresses—and you don't want to overwrite any existing data. Once you're at the right place, select Formulas→Defined Names→Use in Formula→Paste Names. When the Paste Name dialog box appears, click the Paste List button.

This name list is static, which means Excel doesn't update it if you add more names to your worksheet. Instead, you need to generate the list again.

Automatically Creating Named Ranges

Excel also has the built-in smarts to automatically generate named ranges for you. To perform this trick, Excel searches a group of cells you select and, with a little help from you, identifies which cell or cells would serve as appropriate names. If you have a column title in cell A1, Excel can use the text in that cell to label the range of cells underneath it.

To use automatic naming, follow these steps:

1. **Select the cells for which you want Excel to create one or more named ranges.**

   Excel lets you quickly create one or more named ranges. Your selection must include the cells that you want to be part of the named range, plus the cell or cells that contain the descriptive text. Figures 13-10 and 13-11 show examples.

2. **Choose Formulas→Defined Names→Create from Selection.**

   The Create Names from Selection dialog box appears.

3. **Specify which part of your selection has the text label (or labels) you want to use by turning on the appropriate checkbox.**

   In order to create a name, Excel has to find some text to use as a label. You use the Create Names from Selection dialog box to tell Excel where that text is. If the text is contained in column headings, then select "Top row." If there are row and column headings, then select "Top row" and "Left column." Excel selects the options that it thinks make sense based on your selection.

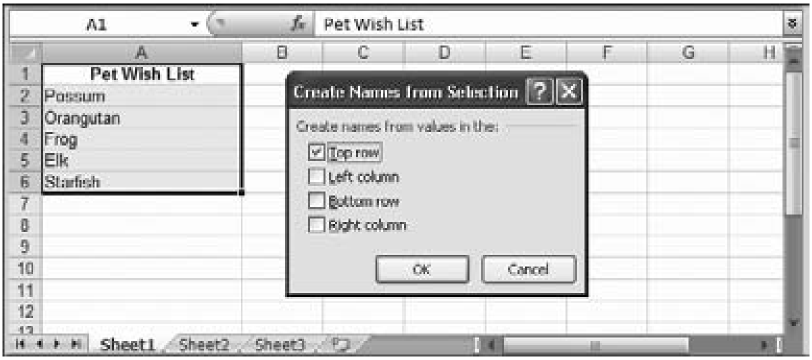

Figure 13-10: *Excel helps you quickly create named ranges. Use the Create Names from Selection dialog box to indicate which cells Excel should use to generate the names. If you just want to create a single named range (from a vertical list, for example), then you'd pick "Top row" in the Create Names from Selection dialog box. Here, the named range is Pet_Wish_List (Excel automatically inserts underscores in place of spaces), and the included cells in this range are A2 to A6.*

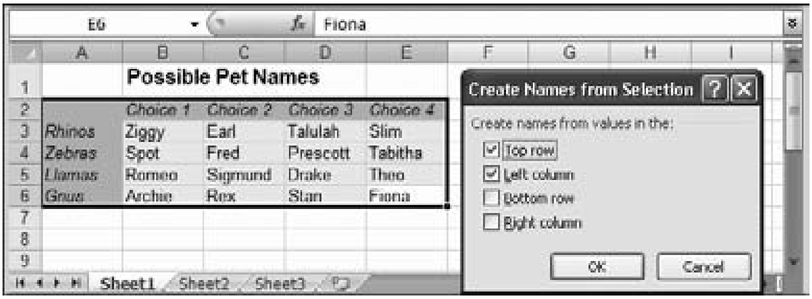

Figure 13-11: *You can also use Create Names from Selection with a table of data. In this case, by turning on Left Column, Excel uses each of the animals in column A as the title for a different named range. The Rhinos range, for example, contains cells B3 to E3. If you turn on the "Top row" checkbox, then Excel also creates named ranges for each column, like Choice_1 (cells B3:B6), Choice_2 (cells C3:C6), and so on. It's up to you whether you want to create ranges for row headings, column headings, or both.*

4. **Click OK to generate the names.**

   Excel's automatic naming ability seems like a real timesaver, but it can introduce as many problems as it solves. Here are some of the quirks and annoyances to watch out for:

   - The named ranges may not be the ones you want, especially if you use long, descriptive column titles. The name *Length* is more manageable than the name *Length_in_Feet*—which is what Excel sticks you with if your name-containing cell happens to contain the text *Length in Feet*. You could rename your named range after the fact (page 377), but it's extra work.
   - If you frequently use the Create Names from Selection tool, you may suddenly find yourself with dozens of names, depending on your worksheet's complexity. Since you haven't created these names yourself, you won't necessarily know what each one references. One way to figure it out is to use the technique described in the box on page 378.
   - Excel doesn't generate names if your column or row labels contain only numbers. If you have a table with a list of part numbers or customer IDs, then Excel assumes the numbers are a part of the data, and not labels it can use for naming. And if your column or row names are text values that start with a number (like "401K"), then Excel adds an underscore at the beginning to make a valid name ("_401K").
   - If you don't have *any* row or column labels, Excel can't generate any names.

Applying Names to Existing Formulas

Named ranges are extremely helpful, but unless you've planned out your spreadsheet writing in a careful, orderly manner, you may find that you spend more time managing these names than actually using them. Suppose you've written a bunch of formulas that work their magic on a table of data that doesn't use any named ranges. Then, *after* you've written the formulas, you catch the named range bug and apply names to, say, all the columns in your table. Your formulas still work, but they don't use the names you've just created.

Note: There's one thing to remember before you get started with the Apply Names feature. Excel can use named ranges only if you've defined them. If your worksheet doesn't include any named ranges, or if they don't match the ranges used in your formulas, this feature doesn't do anything.

So you've either got to go through and manually revise each formula—so that, for example, *=SUM(A2:A85)* becomes *=SUM(My_Stocks)*—or, if you're not a fan of carpal tunnel syndrome, you can use one of Excel's most valuable shortcuts: the ability to automatically replace old-school alphanumeric cell references with reader-friendly named range labels. To use this timesaving technique, follow these steps:

1. **Select the cells that contain the formulas you want to change.**

   Of course, you'll want to use this shortcut only if you've got a bunch of formulas that don't yet use one or more existing named ranges. Select the cells that contain the formulas you want to change. This shortcut can work on more than one formula at a time (even if each formula references a different named range), so it's OK to select a group of formulas.

2. **Select Formulas→Defined Names→Define Name→Apply Names.**

   The Apply Names dialog box appears (Figure 13-12), with a list of all the names in your workbook. Excel searches for and highlights the appropriate names. If your formulas have ranges you don't want to use, then deselect them.

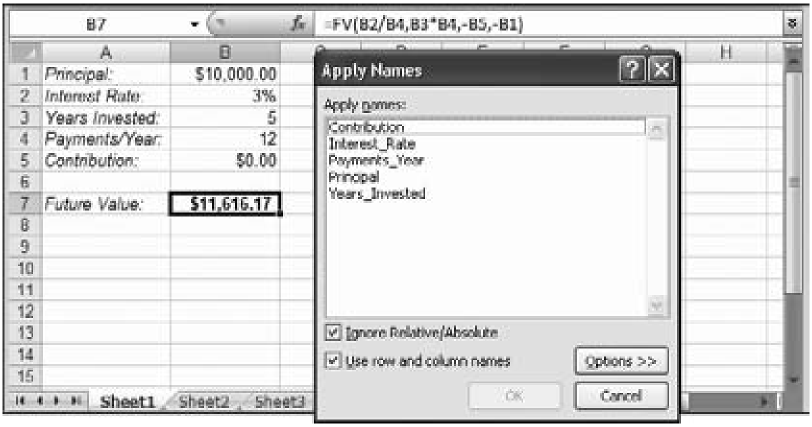

***Figure 13-12:** When you use the Formulas→Defined Names→Define Name→Apply Names command, the formula in B7 changes from =FV(B2/B4,B3*B4,-B5,-B1) to the more understandable (but longer) formula =FV(Rate/Payments, Years*Payments, -Contribution, -Principal).*

If the Ignore Relative/Absolute checkbox is turned on, Excel changes any matching reference, whether it's absolute or relative. This behavior's usually what you want, because most people writing formulas use relative references (like A2 instead of $A$2)—unless they explicitly need absolute references to make copying and pasting easier. If you don't turn this option on, then Excel replaces only absolute references.

You'll also notice the "Use row and column names" checkbox. This setting applies only if you generated names using the Formulas→Named Cells→Create from Selection technique described in the previous section. If you used this technique to create row and column names, then you can turn off the "Use row and column names" checkbox to prevent Excel from using these names in your formulas. You may skip this step because Excel can create horribly complex formulas with the row and column names. That's because in order to get a specific cell in a range, Excel uses *range intersection*. Consider the grid of pet names in Figure 13-11. If you apply names with the "Use row and column names" checkbox turned on, then Excel converts a simple formula like *=B3* into *=RHINOS*

Choice_1. In other words, Excel notices that the intersection of the RHINOS range (B3:E3) and the Choice_1 range (B3:B6) is the single cell B3. You may consider this either an elegant adjustment or an unnecessary way to complicate your life.

3. **Click OK to apply the names.**

   Every time Excel finds a range that matches one of the highlighted range names, it replaces the alphanumeric cell reference with the range name.

Variable Data Tables

When you create a formula, you generate a single result based on the operators, functions, and data that you used. This situation makes sense when you're nailing down last year's profit-loss report, but it's not always as handy when you're making projections for the *future*. In these cases, it's often helpful to compare several possibilities. One tool you can use is a *variable data table*.

A variable data table is simply a table that shows multiple results, based on different source data. You could use a variable data table to see how the return of an investment varies based on different interest rates. Because Excel shows all the results side by side, you can quickly compare them.

You could create your own comparison table by hand without too much trouble. In fact, using the power of relative references, you could create one formula and copy it into several cells to quickly create a table without needing any formula tweaking. Some Excel fans prefer this approach. However, when you use variable data tables you save more than a few keystrokes.

Note: Variable data tables have nothing in common with the similarly named *table feature*, which you'll learn about in Chapter 14.

Creating a One-Variable Data Table

A *one-variable* data table provides a single column of results. It's called one-variable because there's only one input value that changes. For example, if you want to compare how an investment payoff changes based on interest rates, you can create a one-variable data table that uses the interest rate as the changing variable.

Here's how you create a one-variable data table:

1. **Create a sample formula that performs the calculation you plan to use.**

   When you create this formula, make sure that you use a cell reference for the value you want to change, not a literal value. Otherwise, Excel can't *change* the value to calculate the different possibilities.

In the example in Figure 13-13, the formula uses the FV() financial function (page 297) to calculate the future value of an investment. The variable is the rate argument, which is stored in cell B3.

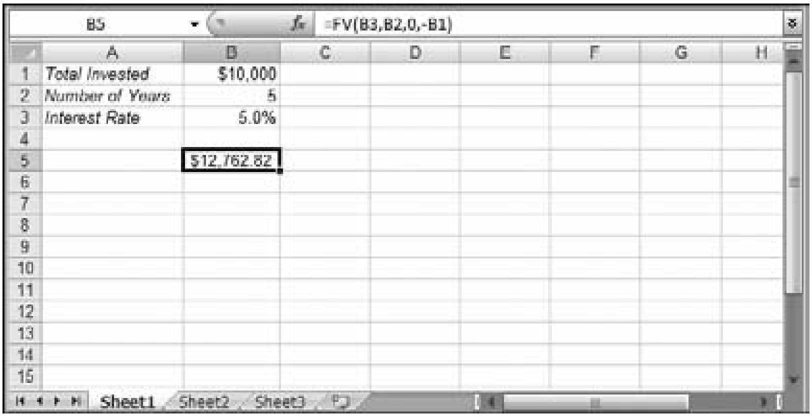

Figure 13-13: *Before you can create a one-variable data table, you need to start with a formula that calculates the result, using cell references. In this example, the formula is in cell B5, and it calculates the value of a $10,000 investment after five years at 5 percent interest.*

2. **Set up the table. To do so, create a column that has each value you want to test.**

   In this example, you need to create a column with the different interest rate possibilities. The calculated results appear in the column on the right. You may find getting the correct arrangement a little tricky. You need to make sure the formula you want to modify is just above the column where the results will appear, as shown in Figure 13-14.

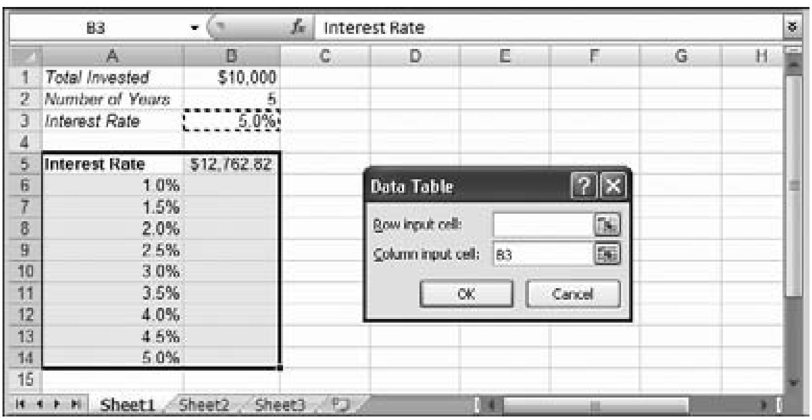

Figure 13-14: *In the Data Table dialog box, specify the cell that has the changing value, and then click OK.*

3. **Generate the table by selecting the whole table, and then choosing Data→Data Tools→What-If Analysis→Data Table.**

   The Data Table dialog box appears (Figure 13-14).

4. **In the "Column input cell" box, type the address of the cell that has the value you want to change.**

   This part is tricky. Excel needs to figure out which part of the formula it's supposed to modify to generate the different results that appear in the table. In this example, the formula's *=FV(B3,B2,0,–B1)*. The cell B3 supplies the rate (currently 5 percent). This rate is the piece of information that needs to change, so you should enter B3 in the "Column input cell" text box. (When Excel generates the data table, it looks at the first column, gets the new rate value, and substitutes this value into the formula where the reference B3 was previously.)

5. **Click OK to create the table.**

   Excel fills in all the different results—in this example, the value of the investment based on different interest rates (Figure 13-15). If you actually look at the content of each cell, you'll find that it uses an array formula with the TABLE() function to work its magic.

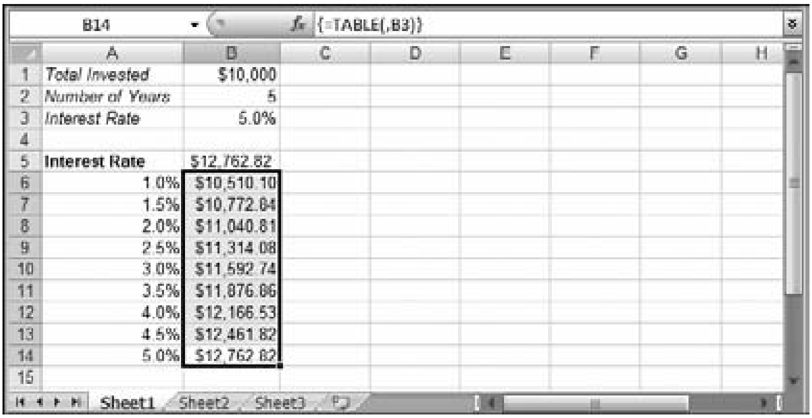

B14 {=TABLE(,B3)}

	A	B
1	Total Invested	$10,000
2	Number of Years	5
3	Interest Rate	5.0%
4		
5	Interest Rate	$12,762.82
6	1.0%	$10,510.10
7	1.5%	$10,772.84
8	2.0%	$11,040.81
9	2.5%	$11,314.08
10	3.0%	$11,592.74
11	3.5%	$11,876.86
12	4.0%	$12,166.53
13	4.5%	$12,461.82
14	5.0%	$12,762.82
15		

Figure 13-15: *The end result of a one-variable data table: a table of investment profits based on different interest rates.*

Creating a Two-Variable Data Table

It's almost as easy to create a *two-variable* data table, where two different values change. In this case, you place one set of changing values in the column on the left. You place the other set of changing values in the row on the top. Figure 13-16 shows an example.

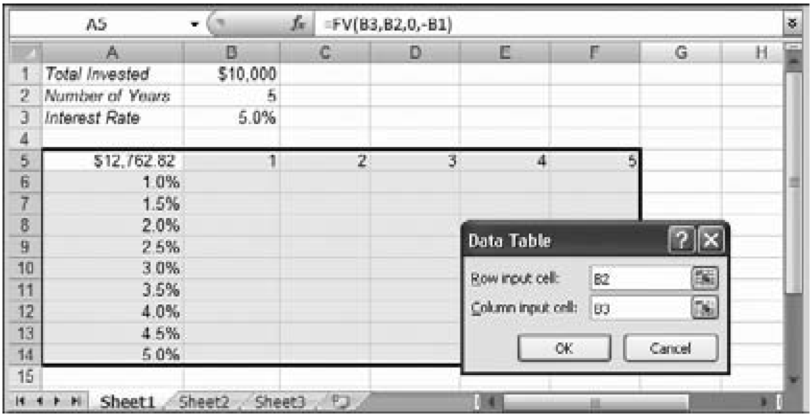

Figure 13-16: *This two-variable data table studies how an investment performs for different interest rates and different investment periods (in years). The row of values B5 to F5 has the number of years, while the column of values A6 to A14 has the interest rate. The formula needs to be in the top-left corner of the table.*

Next, select the whole table and choose Data→Data Tools→What-If Analysis→Data Table. This time, fill out both text boxes (the top row provides an investment period value in place of cell B2, and the leftmost cell provides the interest rate in place of cell B3. When you click OK, Excel generates the complete table (Figure 13-17).

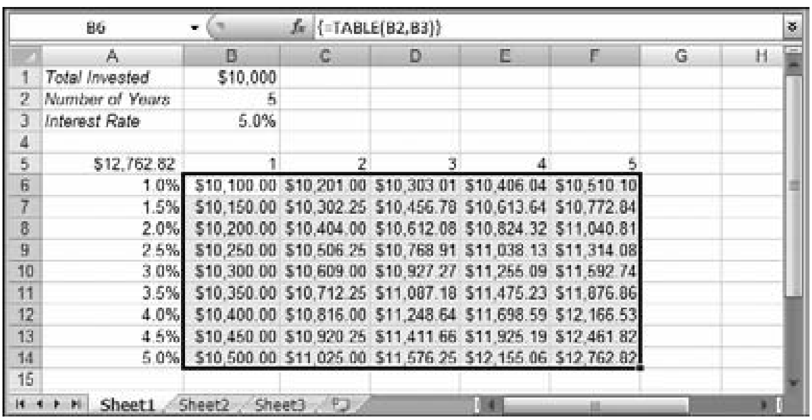

	A	B	C	D	E	F
1	Total Invested	$10,000				
2	Number of Years	5				
3	Interest Rate	5.0%				
4						
5	$12,762.82	1	2	3	4	5
6	1.0%	$10,100.00	$10,201.00	$10,303.01	$10,406.04	$10,510.10
7	1.5%	$10,150.00	$10,302.25	$10,456.78	$10,613.64	$10,772.84
8	2.0%	$10,200.00	$10,404.00	$10,612.08	$10,824.32	$11,040.81
9	2.5%	$10,250.00	$10,506.25	$10,768.91	$11,038.13	$11,314.08
10	3.0%	$10,300.00	$10,609.00	$10,927.27	$11,255.09	$11,592.74
11	3.5%	$10,350.00	$10,712.25	$11,087.18	$11,475.23	$11,876.86
12	4.0%	$10,400.00	$10,816.00	$11,248.64	$11,698.59	$12,166.53
13	4.5%	$10,450.00	$10,920.25	$11,411.66	$11,925.19	$12,461.82
14	5.0%	$10,500.00	$11,025.00	$11,576.25	$12,155.06	$12,762.82
15						

Figure 13-17: *The final table shows how the money you make changes, based on the time you wait and the interest rate.*

Tip: What-if analysis is the art of exploring possibilities with your numbers. The goals of what-if analysis are many—you might want to discover hidden relationships, set new targets, or prepare for future outcomes. For some more sophisticated ways to perform what-if analysis, you can read about scenarios and goal seeking in Chapter 21.

Controlling Recalculation

Ordinarily, Excel recalculates a formula whenever you change any of the cells the formula uses, and whenever you open the workbook containing the formula. This behavior is extremely convenient because it ensures that your information never gets out of date. However, it can cause trouble if your worksheet contains time-consuming calculations or extremely large tables of data. In this situation, life can slow to a crawl while Excel recalculates and updates all its formulas.

Tip: You'll know the recalculation is at work when the mouse pointer changes to an hourglass and the word Recalc (followed by the number of cells left to be recalculated) appears on the left side of the formula bar. Even though your cursor looks like an hourglass, you can still enter data and use the menu, but every time you perform any of these activities, the recalculation process is briefly paused.

With the exponential increase in computer speed (and Microsoft's continuous tweaking of Excel's performance), large worksheets aren't the problem they used to be. However, some Excel aficionados still report problems when crunching data that's thick with formulas. If you're using Excel to plot the path of orbital satellites or to analyze census data, you may be interested in switching off Excel's automatic calculation to make your worksheet more responsive.

To select the calculation mode, choose Formulas→Calculation→Calculation Options, and then pick an option from the menu.

- **Manual** turns off all automatic calculation.
- **Automatic** is the standard way of life for Excel spreadsheets. Excel recalculates formulas whenever cell changes affect them.
- **Automatic Except Data Tables** uses automatic recalculation for all formulas except those in data tables (Chapter 14).

In manual recalculation mode, Excel calculates most simple formulas when you enter them. However, if you modify cells that the formula uses, then the program doesn't recalculate the formula—instead, it keeps displaying the old result. When you see the word *Calculate* in the Status bar, that's your clue that Excel would have ordinarily performed a recalculation, but didn't because of your settings. It means that some data's changed, and a recalculation's a good idea to refresh your formula results.

When you decide you do want to perform your recalculation, you can choose Formulas→Calculation→Calculate Now (or hit F9) to recalculate all the formulas in the workbook. Alternatively, you can choose Formulas→Calculation→Calculate Sheet (or the shortcut Shift+F9) to recalculate just those formulas on the current worksheet.

Note: Unfortunately, the recalculation settings are program-wide. When you switch off automatic calculation for one worksheet, it affects any other worksheets you're using. If you don't remember that you've switched off automatic recalculation, you may end up using and printing stale data.

Solving Formula Errors

Errors … they happen in the most unexpected places, transforming rows of calculations into unhelpful error codes like *#NAME?, #VALUE!*, and *#MORON!* (OK, that last one doesn't actually appear in Excel, but it might as well, given the sense of defeat and frustration these error codes can give you.) In some cases, you can see how to fix an error just by looking at the formula. However, sometimes the problem isn't so easy to solve, especially if your formulas perform calculations using the results of *other* formulas. In such cases, you can have a tough time tracking down where the original error happened.

Excel provides some interesting *formula auditing tools*—a handful of features that you can use to inspect broken formulas or figure out what's going on in really complex ones. These tools make it much easier to fix errors.

With any error, your first step is to identify the error code by using the information listed in Table 8-2 on page 242. If the problem isn't immediately obvious, then you can use the Formula Auditing tools to perform the following tasks:

- Evaluate an expression step-by-step, until you hit the error. That way, you know exactly what part of the formula's causing the error.
- Trace the *precedents* of a formula that's causing an error. Precedents are the cells that a particular formula references. In the formula *=A1+B1*, both A1 and B1 are precedents. If either one of these cells contains an error, then the error gets fed into—and trips up—the formula.
- Trace the *dependents* of a cell. Dependents are other cells that use the current cell. If one cell has the formula *=A1+B1*, and another cell contains *=A1*10*, both these cells are dependents of cell A1. If A1 has an error, it infects both formulas.
- Perform an error check on the entire worksheet. Excel's error check is like a spell check. One by one, it takes you to each cell that has an unresolved problem.

To perform any of these tasks, you'll need the Formulas tab's Formula Auditing section. The following sections explain how you use it to find errors, evaluate formulas piece-by-piece, and trace relationships.

Step-by-Step Evaluation

Complex formulas usually include multiple *sub-expressions*. Each sub-expression is a piece of any formula that's evaluated separately. It may be an arithmetic operation in parentheses, a nested function, or even just a cell reference. In order to understand what's causing an error in your formula, you need to know which sub-expression caused the problem.

Excel's solution is to provide a feature—called the Evaluate Formula tool—that evaluates your formula one sub-expression at a time. Using this tool, you can watch as Excel computes your formula, up until the point where the error appears.

To watch the step-by-step execution of a formula that contains an error, follow these steps:

1. **Move to the cell that contains the formula that's producing the error.**

   You don't need to highlight the formula, you just need to be in the offending cell.

2. **Choose Formulas→Formula Auditing→Evaluate Formula.**

   The Evaluate Formula dialog box appears (Figure 13-18), with the formula in a large, multiline textbox.

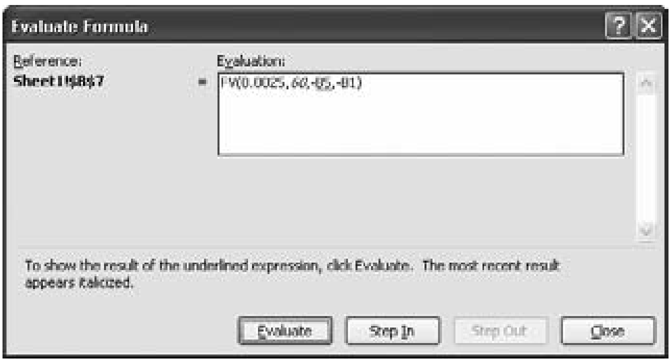

Figure 13-18: *The first two arguments in this formula have been evaluated. (The second argument [60] is italicized, indicating that Excel calculated it in the last step.) The next time you click Evaluate, Excel evaluates the third argument, which is underlined. If you want to show the contents of this cell before evaluating it, you can click Step In.*

3. **Excel underlines the part of the formula that it's about to evaluate. Click the Evaluate button.**

   Excel evaluates the sub-expression and replaces it with the calculated value. It might replace a cell reference with the cell's actual value, evaluate an arithmetic operator, or execute a function. The value appears in italic, indicating that it's the most recent value that Excel has calculated.

4. **Repeat step 3 until the sub-expression that generates the error occurs.**

   When the error occurs, you'll see the error code appear in your formula. When you click Evaluate again, the error code spreads, encompassing the whole expression or the function that uses it. Consider the ill-fated formula =*1+5/0*. The first step (the division) creates a divide-by-zero error, and the formula appears as =*1+#DIV/0!*. But you can't add 1 to an error, so, in the next step, the error spreads to the whole formula, which becomes =#DIV/0! in the end.

5. **When the calculation process ends, you can click Restart to repeat the calculation from the beginning.**

   You can also click Close to stop evaluating the formula and return to your worksheet at any time.

Tip: Step-by-step evaluation isn't just for solving errors. It can also help you understand *why* a formula doesn't produce the result you expect. You can use the Evaluate Formula dialog box with a formula that doesn't cause an error in exactly the same way as one that does. By watching the calculation proceed step-by-step, you may realize that the order of operations Excel follows is subtly different from the order you expected. You can then edit the formula accordingly.

Tracing Precedents and Dependents

The Evaluate Formula dialog box is one way you can examine complex formulas' anatomy. However, depending on the complexity of your formulas, you can end up having to move through a long series of steps before you find the problem. In this case, you may be interested in using a different approach, one that uses Excel's ability to graphically trace linked cells. This feature isn't any better or worse than the Evaluate Formula dialog box—it's just another tool that you can use to resolve problems, depending on the situation and your own preference.

GEM IN THE ROUGH

Digging Deeper into Linked Formulas

Even when you've discovered the sub-expression that's causing the trouble, you still may not have found the root of the problem. If the sub-expression that's causing the error is a reference, it may point to another cell that contains another formula. If it does, then you need to evaluate *that* formula step-by-step in order to find the real mistake.

To evaluate the second formula, you can move to the appropriate cell and start the step-by-step evaluation process by clicking the Evaluate Formula button. However, Excel also provides a useful shortcut that lets you jump from one formula into another. The secret is the Step In and Step Out buttons in the Evaluate Formula dialog box (Figure 13-19).

When you're using the Evaluate Formula dialog box, the Step In button becomes available just before you evaluate a sub-expression containing a cell reference. If you click the Step In button at this point, Excel adds a new text box to the dialog box underneath the first one. This new text box shows the contents of the referenced cell. Excel also informs you if the cell contains a formula or a constant (just read the label at the bottom of the Evaluate Formula dialog box). If the cell contains a constant or there's no calculation left to perform, you need to click Step Out to return to the original formula. If the cell does contain a formula, you can click the Evaluate button to start evaluating it–one sub-expression at a time–and then click Step Out once you're finished.

In fact, Excel lets you dig even deeper into chains of linked formulas. Every time you find a cell reference that points to another formula-holding cell, you can click Step In to show the formula in a new text box. You can continue this process with no practical limit. If you exceed the space available in the Evaluate Formula dialog box, Excel just adds a scroll bar to help you out.

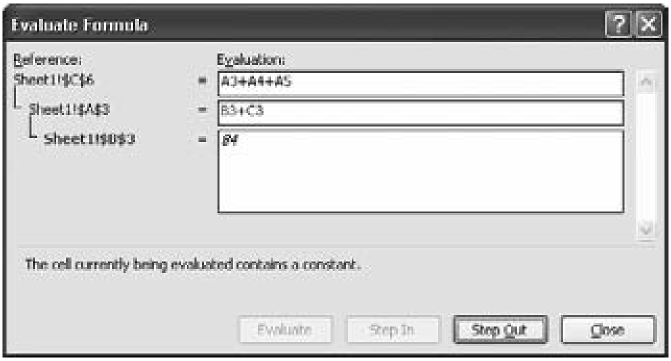

Figure 13-19: *In this example, the Step In button has taken you three levels deep into a formula. The formula it's evaluating is in the first box; it's A3+A4+A5. However, clicking Step In adds a second box, which reveals that cell A3 itself contains a formula (B3+C3). Finally, another click of Step In shows a third box, which zooms in on the first part of the second formula (B3), and shows that the cell it points to holds the number 84.*

First, here's a quick review of how Excel thinks about precedents and dependents. Consider the following formula:

```
=A1+B1
```

If this formula is in cell C1, that makes A1 and B1 *precedents* of C1. In other words, C1 relies on the values in A1 and B1 in order to do its work. If either of these cells contains an error value, the problem spreads into C1. You can say that C1 is the *dependent* of both A1 and B1.

Excel's tracing features let you see a graphical representation of these relationships—in the form of blue arrows—right on your worksheet without needing to look in another window or dialog box.

To see tracing in action, move to a cell that contains one or more cell references, and then choose Formulas→Formula Auditing→Trace Precedents. Excel displays solid blue arrows that link the cells together. If you click Trace Precedents in the cell C1 that contains that formula *=A1+B1*, you see two arrows. One points from A1 to C1, and the other points from B1 to C1. Figures 13-20 and 13-21 show examples.

Note: If a formula references a cell in another worksheet or workbook, Excel draws a dotted line linking your cell to a small grid icon. This icon represents the other worksheet or workbook and can't see the actual cell that the formula links to.

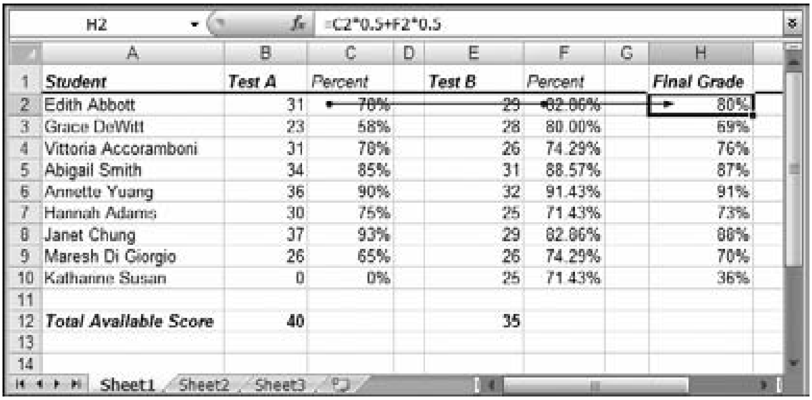

	A	B	C	D	E	F	G	H
1	Student	Test A	Percent		Test B	Percent		Final Grade
2	Edith Abbott	31	78%		29	82.86%		80%
3	Grace DeWitt	23	58%		28	80.00%		69%
4	Vittoria Accoramboni	31	78%		26	74.29%		76%
5	Abigail Smith	34	85%		31	88.57%		87%
6	Annette Yuang	36	90%		32	91.43%		91%
7	Hannah Adams	30	75%		25	71.43%		73%
8	Janet Chung	37	93%		29	82.86%		88%
9	Maresh Di Giorgio	26	65%		26	74.29%		70%
10	Kathanne Susan	0	0%		25	71.43%		36%
11								
12	Total Available Score	40			35			
13								
14								

Figure 13-20: *This example shows the direct precedents of cell H2. As you can see, H2 calculates the student's final grade based on the test results that are stored in cells C2 and F2. Because these two arrows overlap, they appear as one arrow, but you can clearly see two circles, each of which represents the starting point of an arrow (one each on cells C2 and F2).*

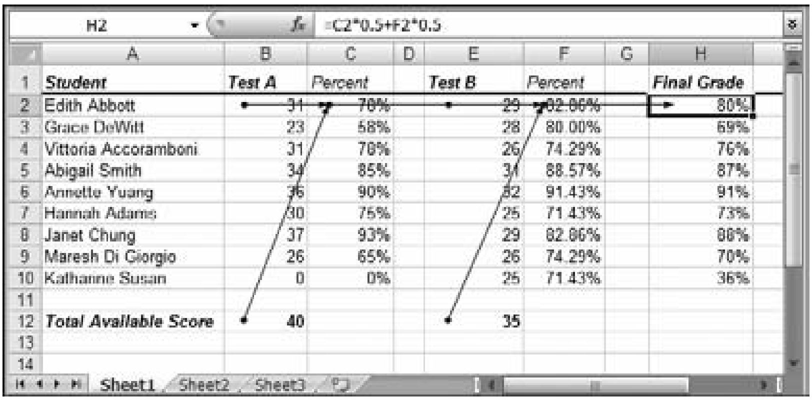

	A	B	C	D	E	F	G	H
1	Student	Test A	Percent		Test B	Percent		Final Grade
2	Edith Abbott	31	78%		29	82.86%		80%
3	Grace DeWitt	23	58%		28	80.00%		69%
4	Vittoria Accoramboni	31	78%		26	74.29%		76%
5	Abigail Smith	34	85%		31	88.57%		87%
6	Annette Yuang	36	90%		32	91.43%		91%
7	Hannah Adams	30	75%		25	71.43%		73%
8	Janet Chung	37	93%		29	82.86%		88%
9	Maresh Di Giorgio	26	65%		26	74.29%		70%
10	Kathanne Susan	0	0%		25	71.43%		36%
11								
12	Total Available Score	40			35			
13								
14								

Figure 13-21: *Excel also lets you trace multiple levels of relationships. Just click the Trace Precedents button again to see whether the precedent cells have other precedents. Here you can see that the test result cells are themselves calculations that rely on other cells. C2 makes its calculations using cells B2 and B12.*

The first time you click Trace Precedents, you see the *direct precedents*. These cells are the ones directly referenced by the current formula. However, these precedents may themselves refer to other cells. To see *these* cells, click Trace Precedents again.

There's no limit to how many times you can click Trace Precedents. As long as there are more indirect precedents, Excel continues adding arrows. At any point, you can remove a single level of arrows by clicking the Remove Precedent Arrows button, or you can clear everything by choosing Formulas→Formula Auditing→Remove Arrows.

Tip: Nothing prevents you from tracing the precedents for a bunch of different cells: just move to another cell and repeat the process for each cell you want to trace, one after the other. You can see all the arrows at once, which can make for a tangled worksheet. When you click Remove Arrows, Excel removes all the precedent arrows and any dependent arrows for every cell you've traced. You can remove the arrows for just one cell by moving to it, and then choosing Formulas®Formula Auditing®Remove Precedent Arrows.

You can trace dependents in the same way that you trace precedents—just choose Formulas→Formula Auditing→Trace Dependents (see Figure 13-22). If you click Trace Dependents and cell A1 is selected, Excel adds an arrow connecting A1 to any other cells that *refer* to A1.

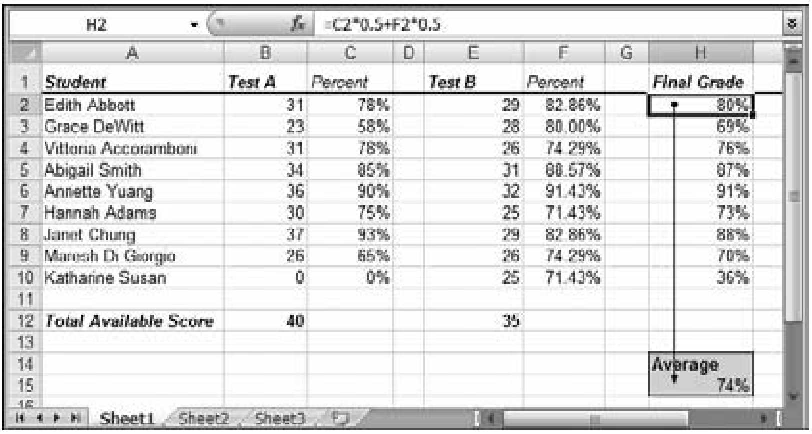

	A	B	C	D	E	F	G	H
1	Student	Test A	Percent		Test B	Percent		Final Grade
2	Edith Abbott	31	78%		29	82.86%		80%
3	Grace DeWitt	23	58%		28	80.00%		69%
4	Vittoria Accoramboni	31	78%		26	74.29%		76%
5	Abigail Smith	34	85%		31	88.57%		87%
6	Annette Yuang	36	90%		32	91.43%		91%
7	Hannah Adams	30	75%		25	71.43%		73%
8	Janet Chung	37	93%		29	82.86%		88%
9	Maresh Di Giorgio	26	65%		26	74.29%		70%
10	Katharine Susan	0	0%		25	71.43%		36%
11								
12	Total Available Score	40			35			
13								
14								Average
15								74%

Figure 13-22: *If you click Trace Dependents on cell H2, Excel indicates that this cell is used in the average calculation in cell H15. However, it isn't the only value that cell H15 uses. To see all the precedents, you'd need to move to H15, and then click Trace Precedents.*

Tip: Check out the Missing CD page at *www.missingmanuals.com/cds* to see a screencast (an animated online tutorial) that shows how the precedent and dependent arrows work. It was created using Excel 2007, but the process is the same for Excel 2010. Surf to the Missing CD page for *Excel 2007: The Missing Manual* to check it out.

There really isn't a difference between precedent and dependent arrows—they're just two different ways of looking at the same idea. In fact, every arrow Excel draws connects one precedent to one dependent. Finally, Excel's tracing tools also work with formulas that *aren't* working (which is important, after all, when it comes to troubleshooting). Figure 13-23 shows how the tool works when your formulas are generating error codes.

Error Checking

Sometimes, you may have a large worksheet containing a number of errors that are widely distributed. Rather than hunt for these errors by scrolling endlessly, you can jump straight to the offending cells using Excel's error-checking feature.

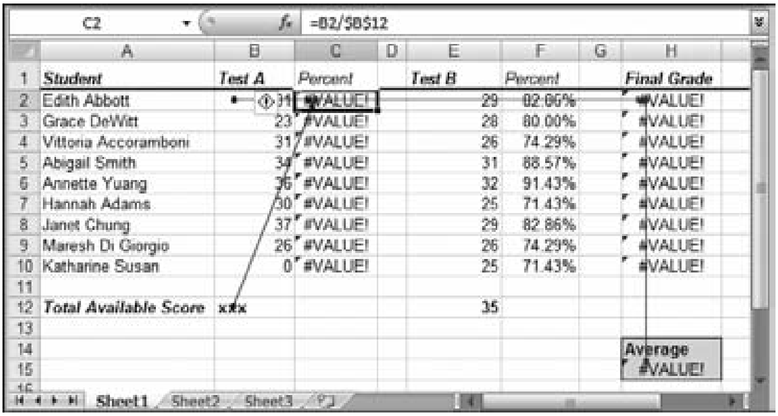

Figure 13-23: *Excel's tracing features work with any formulas—whether or not they contain an error. But Excel also includes a related feature, Trace Error, which works only with formulas that result in error values. When you select a cell with an error code, and then choose Trace Error, Excel traces all the precedents that lead back to the error by using blue arrows. Then, Excel uses red arrows to indicate how the error spread. In this example, two blue arrows show the precedents of cell C2, where the error occurred. The error then spread to cell H2 and, finally, to the current cell, H15.*

To perform an error check, follow these steps:

1. **Move to the position where you want to start the error check.**

   If you want to check the entire worksheet from start to finish, click the first cell. Otherwise, go to the location where you want to start checking. As with a spell check, Excel moves from column to column first, and then from row to row. However, the error checker automatically loops back to the beginning of your worksheet, making sure to check every cell before it stops.

2. **Choose Formulas→Formula Auditing→Error Checking.**

   If Excel doesn't find any errors in your worksheet, it displays a message indicating that its work's complete. Otherwise, you see the Error Checking dialog box, as shown in Figure 13-24, which indicates the offending cell and formula. This box also provides a number of options.

The Error Checking dialog box contains the following options:

- **Next or Previous**. Use these buttons to move from one error to the next.
- **Help on this Error**. Click this button to jump to Excel's online help, which lists common causes of specific errors. It may give you some insight into your own troubles.

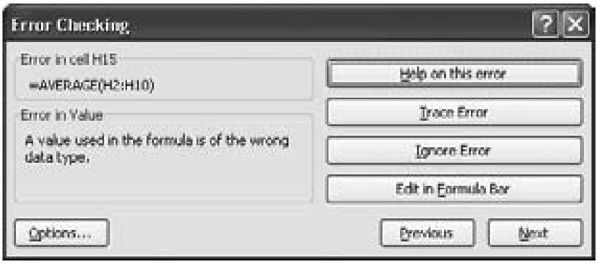

Figure 13-24:
Excel's error checker helps you scan through a worksheet and quickly jump to the cells that contain errors. You can click the Trace Error button to quickly jump to the Evaluate Formula dialog box and start analyzing the problem.

- **Trace Error**. Use this button to open the Evaluate Formula dialog box, where you can move one step at a time through the evaluation of the formula.
- **Ignore Error**. Click this button to skip the error, and ignore the cell from this point onward. If you want the error checker to pay attention to a cell you've previously decided to ignore, click the Options button, and then click Reset Ignored Errors.
- **Edit in Formula Bar**. Use this button to start editing the formula. This choice doesn't close the error checker—instead, you can click Resume to get back to checking other errors once you've made your change.
- **Options**. Click this button to open the Excel Options dialog box, with the Formulas section selected. The Formulas section includes a small set of error-checking options under the headings "Error checking" and "Error checking rules" (page 28). Ordinarily, you don't need to change any of these error checking options, as the factory settings are stringent enough to ensure that Excel catches all problems.

GEM IN THE ROUGH

Follow the Arrow

If you have a complex number-laden spreadsheet with formulas that pull values from all over the place, you may need to scroll around to find the precedents or dependents that interest you. Excel has a trick to help you out–just double-click the appropriate arrow to follow it back to its source cell, no matter where it lies.

If you've got a particularly tricky worksheet, it may contain formulas that draw upon values in other worksheets or workbooks. If you double-click the arrow in this situation, then Excel pops up the Go To dialog box (page 28), with the information about the source cell already filled in at the top of the list. If you want to follow the arrow to the new worksheet or file, just select the reference (something like, *[SuperSecretValues.xlsx]Sheet1!$A$3*), and then click OK.

CHAPTER 14

Tables: List Management Made Easy

Excel's grid-like main window gives you lots of freedom to organize your information. As you've seen in the chapters so far, tables of data can assume a variety of shapes and sizes—from complex worksheets that track expenses, to a simple list of dishes your guests are bringing to a potluck dinner.

Some tables are quite sophisticated, with multiple levels, subtotals, and summary information. (You'll learn about how to manage these multi-tiered creations in the next chapter.) But in many cases, your table consists of nothing more than a long list of data, with a single row at the top that provides descriptive column headings. These types of tables are so common that Excel provides a set of features designed exclusively for managing them. These tools let you control your tables in style—sorting, searching, and filtering your information with just a couple of mouse clicks. Excel even includes a group of functions expressly designed to analyze the information in tables. But before you can use any of these tools, you have to convert your garden-variety table into a *structured table*.

In this chapter, you'll learn more about what, exactly, a structured table is, how to create one, and how to make use of all its features and frills.

Note: Don't confuse structured tables with the variable data tables you used for what-if analysis. These tables have a similar moniker but nothing else in common.

The Basics of Tables

An Excel table is really nothing more than a way to store a bunch of information about a group of items. Each item occupies a separate row, and different kinds of information about the item reside side by side in adjacent columns. In database terminology, the rows are *records*, and the columns of information are *fields*. For example, the records could represent customers, and the fields could contain things like name, address, purchase history, and so on.

Excel tables have a number of advantages over ordinary worksheet data:

- **They grow and shrink dynamically**. As you fill data into adjacent rows and columns, the table grows to include the new cells. And as a table changes size, any formulas that use the table adjust themselves accordingly. In other words, if you have a formula that calculates the sum of a column in a table, the range that the SUM() function uses expands when you add a new record to the table.
- **They have built-in smarts**. You can quickly select rows and columns, apply a custom sort order, and search for important records.
- **They excel (ahem) at dealing with large amounts of information**. If you need to manage vast amounts of information, you may find ordinary worksheet data a little cumbersome. If you put the same information in a table, you can simply apply *custom filtering*, which means you see only the records that interest you.
- **They can link to databases**. Tables are perfectly useful in standalone worksheets. However, they can also double as indispensable tools for navigating information contained in a database. In Chapter 23, you'll learn how to get information *out of* a database and *into* an Excel table.

Creating a Table

Creating a table is easy. Here's how:

1. **Choose the row where you want your table to start.**

   If you're creating a new table, the worksheet's first row is a good place to begin. (You can always shift the table down later by putting your cursor in the top row, and then choosing Home→Cells→Insert→Insert Sheet Rows.) This first row of the table is where you enter any column titles you want to use, as explained in the next step.

Note: Be careful when placing content in the cells directly *beneath* your table. If your table expands too far down, you'll run up against these filled-up cells. Although you can use commands like Home→Cells→Insert→Insert Sheet Rows to add some extra space when things get crowded, it's always better to start off with plenty of breathing room.

2. **Enter the column titles for your table, one column title for each category you want to create.**

   To create the perfect table, you need to divide your data into categories. For example, if you're building a table of names and addresses, you probably want your columns to hold the standard info you see on every form ever created: First Name, Last Name, Street, City, and so on. The columns you create are the basis for all the searching, sorting, and filtering you do. For instance, if you have First Name and City columns, you can sort your contacts by first name or by city.

   If you want, you can start to add entries underneath the column headings now (in the row directly below the column titles). Or just jump straight to the next step to create the table.

3. **Make sure you're currently positioned somewhere inside the table (anywhere in the column title row works well), and then choose Insert→Tables→Table.**

   Excel scans the nearby cells, and then selects all the cells that it thinks are part of your table. Once Excel determines the bounds of your table, the Create Table dialog box appears, as shown in Figure 14-1.

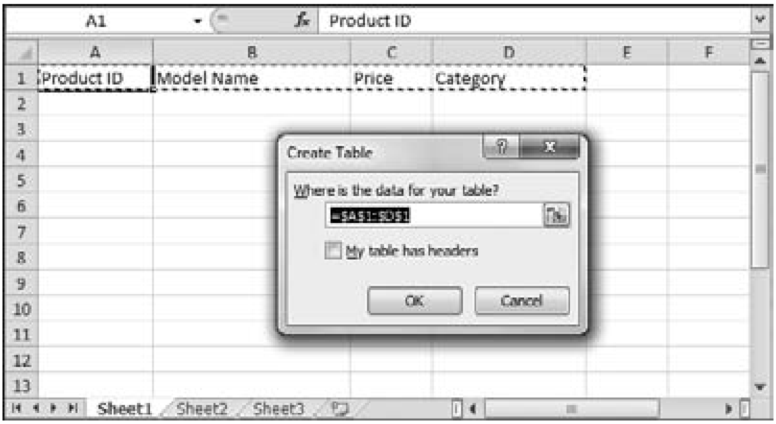

Figure 14-1: *The Create Table dialog box displays the cell references for the currently selected range. In this example, the selection includes only the headings (there's no data yet). You can change the range by typing in new information or by clicking the mini worksheet icon at the right end of the cell range box, which lets you select the range by dragging the appropriate cells in the worksheet.*

4. **Make sure the "My table has headers" checkbox is turned on. This option tells Excel you're using the first row just for column headers. Then click OK.**

   Excel transforms your cells into a table, like the one shown in Figure 14-2. You can tell that your ordinary range of cells has become a genuine table by the presence of a few telltale signs. First, tables start out with automatic formatting that gives each row a shaded background (alternating between blue and gray). Second, the column headings appear in bold white letters on a dark background, and each one includes a drop-down arrow that you can use for quick filtering (a feature you'll explore on page 411).

If you create a table from a group of cells that don't include column titles, don't turn on the "My table has headers" checkbox. When you create the table, Excel adds a row of columns at the top with generic names like Column1, Column2, and so on. You can click these cells, and then edit the column titles, to be more descriptive.

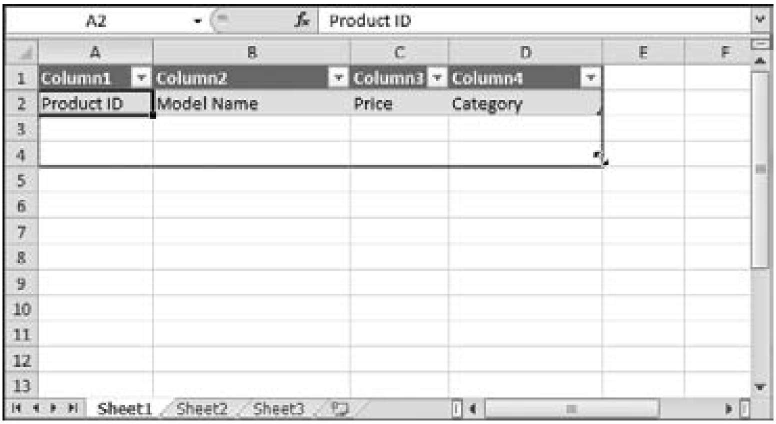

Figure 14-2: *To quickly resize your table, look for the tiny triangle icon at the bottom-right corner (under the two-headed arrow in this figure), and then drag it to encompass more (or fewer) rows and columns.*

Keep in mind that tables consist of exactly two elements: column headers (Figure 14-3) and rows. Tables don't support row headers (although there's no reason why you can't create a separate column and use that as a row title). Tables also have a fixed structure, which means that every row has exactly the same number of columns. You can create multiple tables on the same worksheet, but you're often better off placing them on separate worksheets so you can more easily manage them.

Figure 14-3: *Here's one unsung frill in every table. When you can't see the column headers any longer (because you've scrolled down the page), the column buttons atop the worksheet grid change from letters (like A, B, C) to your custom headers (like Product ID, Model Name, and Price). This way, you never forget what column you're in, and you don't need to scroll up to use the sorting and filtering features that you can access through the drop-down column menus (page 412).*

Formatting a Table

When you move to a cell in a table, Excel adds a new tab to the ribbon—the Table Tools | Design tab. This tab has five sections:

- **Properties**, which lets you rename your table and resize it. (The table name is important only if you choose to use it in a formula, as described on page 421.)
- **Tools**, which lets you use pivot tables (described in Chapter 22), find duplicate rows (page 415), and convert your table back to ordinary cells.
- **External Table Data**, which lets you work with external data (like records drawn from a database) using an Excel table. (You'll learn more in Chapter 23 about how to perform maneuvers like that.)
- **Table Style Options** and **Table Styles**, which help you make your table look pretty.

Every table starts out with some basic formatting, and you can use the ribbon and the Format Cells dialog box (as discussed in Chapter 5) to further change its appearance. However, Excel gives you an even better option—you can use *table styles*.

A table style is a collection of formatting settings that apply to an entire table. The nice part about table styles is that Excel remembers your style settings. If you add new rows to a table, Excel automatically adds the right cell formatting. Or, if you delete a row, Excel adjusts the formatting of all the cells underneath to make sure the *banding* (the alternating pattern of cell shading that makes each row easier to read) stays consistent.

When you first create a table, you start out with a fairly ordinary set of colors: a gray–blue combination that makes your table stand out from the rest of the worksheet. By choosing another table style, you can apply a different set of colors and borders to your table.

Note: Excel's standard table styles don't change the fonts in a table. To change fonts, you can change the theme (page 174), or select some cells, and then, from the ribbon's Home→Font section, pick the font you want.

To choose a new table style, head to the ribbon's Table Tools | Design→Table Styles section. You'll see a gallery of options, as shown in Figure 14-4. As you move over a table style, Excel uses its live preview feature to change the table, giving you a sneak peak at how your table would look with that style.

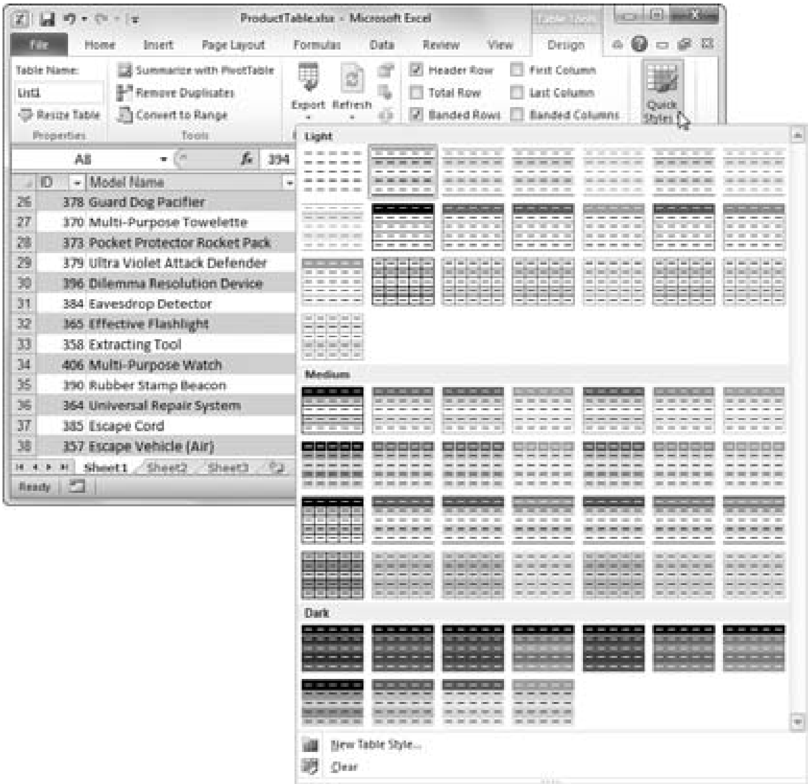

Figure 14-4: *Depending on your Excel window's width, in the ribbon, you may see the table style gallery. Or, if there's not enough room available, you see a Quick Styles button that you need to click to display a drop-down style gallery (as shown here).*

Note: Notice that some table styles use banding, while others don't.

Table styles work like cell styles, which you learned about on page 176. Like cell styles, they let you standardize and reuse formatting. Table styles, however, include a whole package of settings that tell Excel how to format different portions of the table, including the headers, first and last columns, the summary row (page 422), and so on. To get a better feel for the variety of options in a table style, you can create your own by clicking the drop-down arrow in the Table Tools | Design→Table Styles section, and then choosing New Table Style. You'll see a New Table Quick Style dialog box that lets you go through the somewhat painstaking process of tweaking the formatting for each part of the table (Figure 14-5).

Note: You can't edit the built-in table styles. However, you can change the table styles you create. In the table gallery, just right-click a style, and then choose Modify.

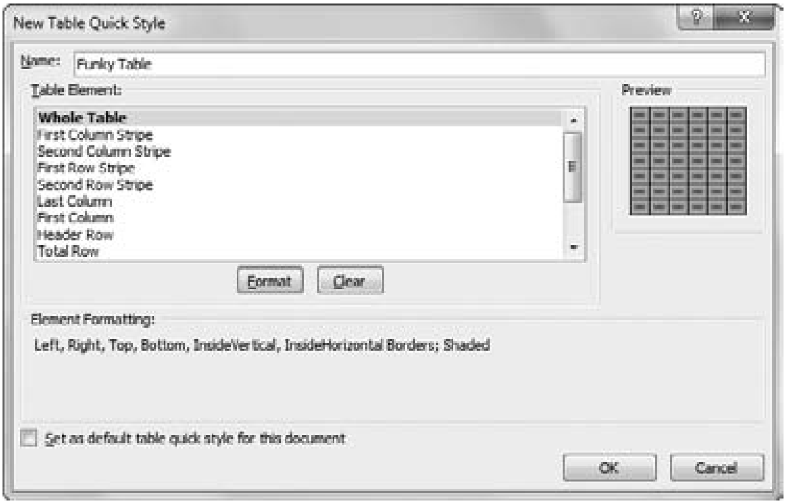

Figure 14-5: *To set the formatting for a part of the table, select it in the Table Element list, and then click Format. Excel opens a slimmed-down Format Cells dialog box that includes only the Font, Border, and Fill tabs. For example, to change the font for the whole table, select Whole Table, click Format, and then make your changes. If you think you've created the perfect table style, you can turn on the "Set as default table quick style" checkbox so Excel uses it whenever you create a new table in this workbook.*

Most of the time, it's not worth creating your own table styles because it's simply too much work (and the prebuilt table styles give you a good selection of formatting choices). However, you'll notice that the built-in table styles have a limited set of colors. Excel limits them because table styles use colors from the current theme, which ensures that your table meshes well with the rest of your worksheet (assuming you've been sticking to theme colors elsewhere). To get different colors for your tables, you can change the theme by choosing from the Page Layout→Themes→Themes gallery. Page 174 has more about themes.

Along with the table style and theme settings, you have a few more options to fine-tune your table's appearance. Head over to the ribbon's Table Tools | Design→Table Style Options section, where you see a group of checkboxes, each of which lets you toggle on or off different table elements:

- **Header Row** lets you show or hide the row with column titles at the top of the table. You'll rarely want to remove this option. Not only are the column headers informative, but they also include drop-down lists for quick filtering (page 411).
- **Total Row** lets you show or hide the row with summary calculations at the bottom of your table. You'll learn how to configure this row on page 422.
- **First Column** applies different formatting to the first column in your table, if it's defined in the table style.
- **Last Column** applies different formatting to the last column in your table, if it's defined in the table style.

- **Banded Rows** applies different formatting to each second row, if it's defined in the table style. Usually, the banded row appears with a background fill. Large-table lovers like to use banding because it makes it easier to scan a full row from right to left without losing your place.
- **Banded Columns** applies different formatting to each second column, if it's defined in the table style. Folks use banded less than banded rows, because people usually read tables from side to side (not top to bottom).

UP TO SPEED

The Difference Between Excel Worksheets and Databases

An Excel table uses some of the same concepts as a database—namely, the idea of records and fields. However, databases and Excel worksheets are two very different entities.

For starters, databases—which programs like Microsoft Access and SQL Server let you create—have much stricter rules than Excel worksheets. Before you can add any data to a table in a database, you must carefully define the table. You need to specify not only the name of each field, but also the type of information the field can contain. Although Excel provides some of these so-called data validation features (which you'll explore in Chapter 24), the program isn't nearly as strict about it—validation is completely optional. Also, unlike Excel, most modern databases are *relational*, which means they contain multiple tables that have specific links to one another. For example, a relational database might tie together customers in one table and the orders they've made in another. In Excel, a worksheet can hold multiple tables of data, but there's no way to tie them together.

Most importantly, databases play a dramatically different role in the world of business. Typically, Excel is an *end user* program, which means ordinary mortals who generally know how to create an Excel file, design what it's going to look like, and then fill it up with data to use it. Ex-math majors, on the other hand, usually create databases, and they store information, behind-the-scenes, that non-programmer types end up using. For example, every time you use Google or search on Amazon for something to buy, you're actually seeing answers that have been stored in, and generated by, massive and powerful databases. In Chapter 23, you'll see how you can use Excel to retrieve information from a database, and then analyze it in a worksheet.

Editing a Table

Once you've created a table, you can perform three basic editing tasks:

- **Edit a record**. This part's easy. Just modify cell values as you would in any ordinary worksheet.
- **Delete a record**. First, go to the row you want to delete (you can be in any column). Then choose Home→Cells→Delete→Delete Table Rows. Excel removes the row and shrinks the table automatically. For faster access that bypasses the ribbon altogether, just right-click a cell in the appropriate row, and then choose Delete→Table Rows.
- **Add a new record**. To add a record, head to the bottom of the table, and then type a new set of values just underneath the last row in the table. Once you finish typing the first value, Excel expands the table automatically, as shown in Figure 14-6.

If you want to insert a row but don't want it to be at the bottom of the table, you can head to your chosen spot, and then choose Home→Cells→Insert→Insert Table Rows Above (or right-click, and then choose Insert→Table Rows Above). Excel inserts a new blank row immediately *above* the current row.

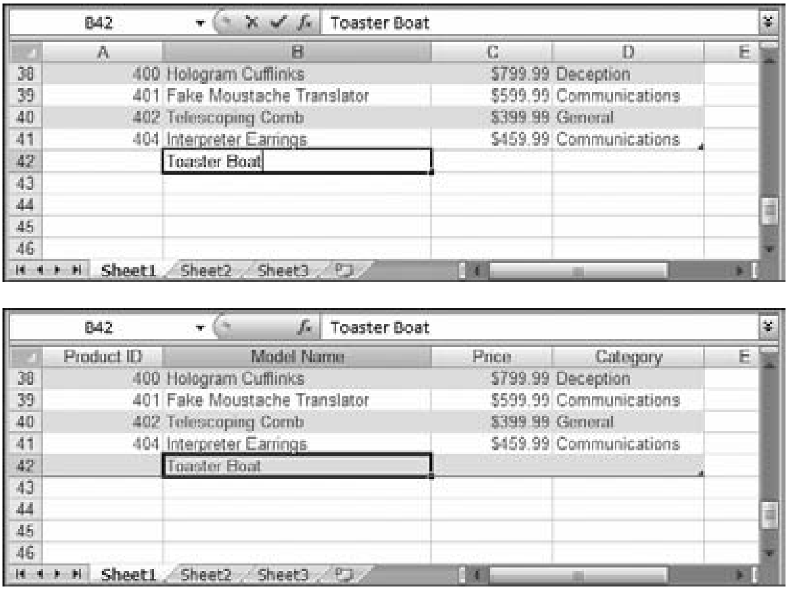

Figure 14-6: *Top: Here, a new record is being added just under the current table.*

Bottom: Once you enter at least one column of information and move to another cell, Excel adds the new row to the table and formats it. Don't worry—if Excel expands a table against your wishes, you aren't completely powerless. To correct Excel's mistake, look for the lightning bolt icon that appears immediately next to the newly added column. This doohickey is a smart tag that lets you reverse the expansion. Click it once, and Excel displays a pop-up menu giving you two choices: Undo Table AutoExpansion (to return your table to its previous size) and Stop Automatically Expanding Tables (to turn off this behavior altogether).

Note: Notice that when you insert or remove rows, you're inserting or removing *table* rows, not *worksheet* rows. The operation affects only the cells in that table. For example, if you have a table with three columns and you delete a row, Excel removes three cells, and then shifts up any table row underneath. Any information in the same row that exists *outside* the table is unaffected.

You may also decide to change the structure of your table by adding or removing columns. Once again, you'll find this task is like inserting or removing columns in an ordinary worksheet. (The big difference, as shown in Figure 14-7, is that any rows or columns *outside* your table remain unaffected when you add new rows or columns.)

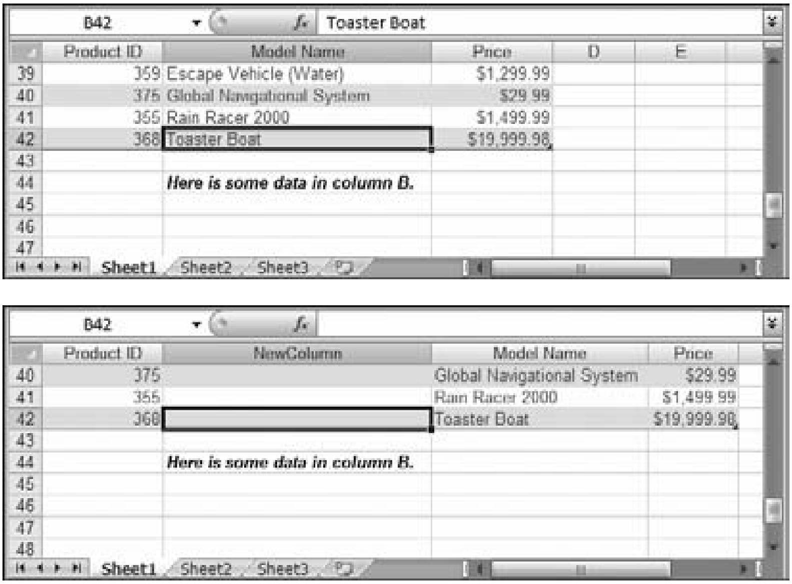

Figure 14-7: *Excel makes an effort to leave the rest of your worksheet alone when you change your table's structure. For example, when expanding a table vertically or horizontally, Excel moves cells out of the way only when it absolutely needs more space. The example here shows the point. Compare the before (top) and after (bottom) pictures: Even though the table in the bottom figure has a new column, it hasn't affected the data underneath the table, which still occupies the same column. The same holds true when deleting columns.*

To add a column to the left of a column you're currently in, select Home→Cells→Insert→Insert Table Columns to the Left. Excel automatically assigns a generic column title, like Column1, which you can then edit. If you want to add a column to the right side of the table, just start typing in the blank column immediately to the right of the table. When you've finished your entry, Excel automatically merges that column into the table, in the same way that it expands to include new rows.

To delete a column, move to one of its cells, and then choose Home→Cells→Delete→Delete Table Column.

Finally, you can always convert your snazzy table back to an ordinary collection of cells. Just click anywhere in the table, and then choose Table Tools | Design→Tools→Convert to Range. But then, of course, you don't get to play with your table toys anymore.

Selecting Parts of a Table

Once you've created a table, Excel provides you with some nice time-saving tools. For example, Excel makes it easy to select a portion of a table, like an individual row or column. Here's how it works:

- **To select a column**, position your mouse cursor over the column header. When it changes to a down-pointing arrow, click once to select all the values in the column. Click a second time to select all the values plus the column header.

- **To select a row**, position your mouse cursor over the left edge of the row until it turns to a right-pointing arrow, and then click once.
- **To select the entire table**, position your mouse at the top-left corner until it turns into an arrow that points down and to the right. Click once to select all the values in the table, and click twice to select all the values plus the column headers.

Figure 14-8 shows an example.

	Product ID	Model Name	Price	Category
2	360	Communications Device	$49.99	Communications
3	401	Fake Moustache Translator	$599.99	Communications
4	404	Interpreter Earrings	$459.99	Communications
5	397	Nonexplosive Cigar	$29.99	Communications
6	362	Persuasive Pencil	$1.99	Communications
7	376	Cloaking Device	$9,999.99	Deception
8	394	Correction Fluid	$1.99	Deception
9	374	Counterfeit Creation Wallet	$999.99	Deception
10	400	Hologram Cufflinks	$799.99	Deception
11	377	Indentity Confusion Device	$6.99	Deception
12	388	Contact Lenses	$59.99	General

Figure 14-8:
You can easily select an entire column in a table. Just position the mouse as shown here, and then click once.

Once you've selected a row, column, or the entire table, you can apply extra formatting, or use another Excel feature like validation (page 694) or charting (Chapter 17). However, changing a part of a table isn't exactly like changing a bunch of cells. For example, if you give 10 cells a hot pink background fill, that's all you get—10 hot pink cells. But if you give a column a hot pink background fill, your formatting change may initially affect 10 cells, but every time you add a new value in that column, it also gets the hot pink background. This behavior, in which Excel recognizes that you're changing parts of a table, and applies your change to new rows and columns automatically, is called *stickiness*.

Sorting and Filtering a Table

As you've seen, Excel tables make it easier to enter, edit, and manage large collections of information. Now it's time to meet two of the most useful table features:

- **Sorting** lets you order the items in your table alphabetically or numerically according to the information in a column. By using the correct criteria, you can make sure the information you're interested in appears at the top of the column, and you can make it easier to find an item anywhere in your table.
- **Filtering** lets you display only certain records in your table based on specific criteria you enter. Filtering lets you work with part of your data and temporarily hide the information you aren't interested in.

You can quickly apply sorting and filtering using the drop-down column headers that Excel adds to every table.

Note: Don't see a drop-down list at the top of your columns? A wrong ribbon click can inadvertently hide them. If you just see ordinary column headings (and you know you have a bona fide table), choose Data→Sort & Filter→Filter to get the drop-down lists back.

Applying a Simple Sort Order

Before you can sort your data, you need to choose a *sorting key*—the piece of information Excel uses to order your records. For example, if you want to sort a table of products so the cheapest (or most expensive) products appear at the top of the table, the Price column would be the sorting key to use.

In addition to choosing a sorting key, you also need to decide whether you want to use ascending or descending order. Ascending order, which is most common, organizes numbers from smallest to largest, dates from oldest to most recent, and text in alphabetical order. (If you have more than one type of data in the same column—which is rarely a good idea—text appears first, followed by numbers and dates, then true or false values, and finally error values.) In descending order, the order is reversed.

Note: Remember, it's technically possible to have numbers in Excel that are stored as text, as described in Chapter 2. Simply prefix these values with an apostrophe ('). For example, you might may store a number as text when you're entering numeric content that doesn't convey a numeric value, like a Social Security number. Excel sorts these values alphabetically, which means it looks at the text string one character at a time until it finds a difference. Thus, even though 42 is less than 102 in a numeric sort, the text *42* is greater than *102* in an alphabetic sort, because the first character *4* comes after *1.*

To apply a new sort order, choose the column you want to use for your sort key. Click the drop-down box at the right side of the column header, and then choose one of the menu commands that starts with the word "Sort." The exact wording depends on the type of data in the column, as follows:

- **If your column contains numbers**, you see "Sort Smallest to Largest" and "Sort Largest to Smallest".
- **If your column contains text**, you see "Sort A to Z" and "Sort Z to A" (see Figure 14-9).
- **If your column contains dates**, you see "Sort Oldest to Newest" and "Sort Newest to Oldest".

When you choose an option, Excel immediately reorders the records, and then places a tiny arrow in the column header to indicate that you used this column for your sort. However, Excel doesn't keep re-sorting your data when you make changes or add new records (after all, it would be pretty distracting to have your records jump around unexpectedly). If you make some changes and want to reapply the sort, just go to the column header menu and choose the same sort option again.

If you click a second column, and then choose Sort Ascending or Sort Descending, the new sort order replaces your previous sort order. In other words, the column headers let you sort your records quickly, but you can't sort by more than one column at a time.

Sorting with Multiple Criteria

Simple table sorting runs into trouble when you have duplicate values. Take the product table sorted by category in Figure 14-9, for example. All the products in the Communications category appear first, followed by products in the Deception category, and so on. However, Excel doesn't make any effort to sort products that are in the *same* category. For example, if you have a bunch of products in the Communications category, then they appear in whatever order they were in on your worksheet, which may not be what you want. In this case, you're better off using *multiple sort criteria*.

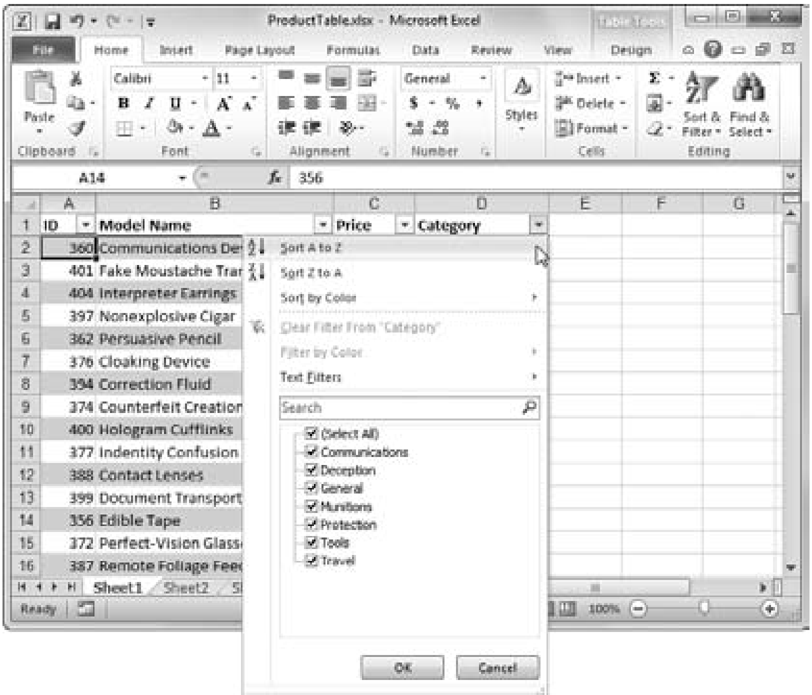

Figure 14-9: *A quick click is all it takes to order records in ascending order by their category names. You don't need to take any action to create these handy drop-down lists–Excel automatically provides them for every table.*

With multiple sort criteria, Excel orders the table using more than one sorting key. The second sorting key springs into action only if there are duplicate values in the first sorting key. For example, if you sort by Category and Model Name, Excel first separates the records into alphabetically ordered category groups. It then sorts the products in each category in order of their model name.

To use multiple sort criteria, follow these steps.

1. **Move to any one of the cells inside your table, and then choose Home→Editing→Sort & Filter→Custom Sort.**

   Excel selects all the data in your table, and then displays the Sort dialog box (see Figure 14-10) where you can specify the sorting keys you want to use.

Tip: You can use the Home→Editing→Sort & Filter→Custom Sort command with any row-based data, including information that's not in a table. When you use it with nontable data, Excel automatically selects the range of cells it believes constitutes your table.

---

2. **Fill in the information for the first sort key in the Column, Sort On, and Order columns.**

   Figure 14-10 shows how it works.

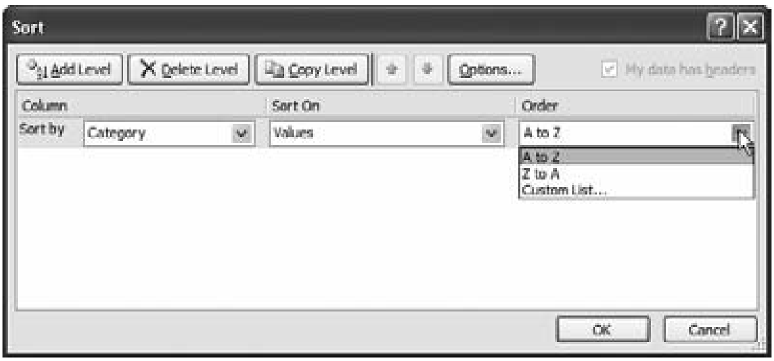

Figure 14-10: *To define a sorting key, you need to fill in the column you want to use (in this example, Category). Next, pick the information you want to use from that column, which is almost always the actual cell values (Values). Finally, you need to choose the order for arranging values, which depends on the type of data. For text values, as in this example, you can pick A to Z, Z to A, or Custom List (page 76).*

3. **If you want to add another level of sorting, click Add Level, and then follow the instructions in step 2 to configure it.**

   You can repeat this step to add as many sorting levels as you want (Figure 14-11). Remember, it makes sense to add more levels of sorting only if there's a possibility of duplicate value in the levels you've added so far. For example, if you've sorted a bunch of names by last name, you want to sort by first name, because some people may share the same last name. However, it's probably not worth it to add a third sort on the middle initial, because very few people share the same first and last name.

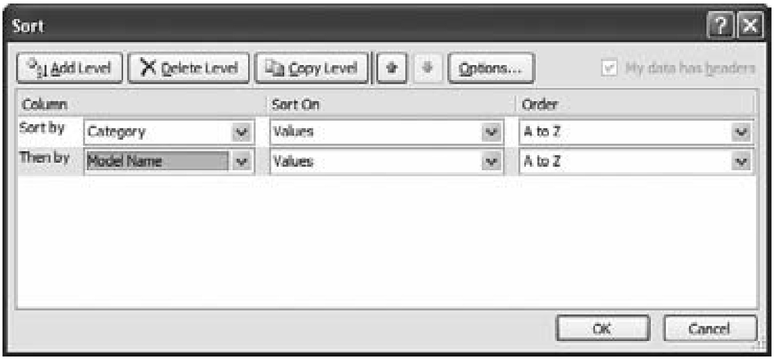

Figure 14-11: *This example shows two sorting keys: the Category column and the Model Name column. The Category column may contain duplicate entries, which Excel sorts in turn according to the text in the Model Name column. When you're adding multiple sort keys, make sure they're in the right order. If you need to rearrange your sorting, select a sort key, and then click the arrow buttons to move it up the list (so it's applied first) or down the list (so it's applied later).*

4. **Optionally, click the Options button to configure a few finer points about how your data is sorted.**

   For example, you can turn on case-sensitive sorting, which is ordinarily switched off. If you switch it on, *travel* appears before *Travel*.

5. **Click OK.**

   Excel sorts your entire table based on the criteria you've so carefully specified (Figure 14-12).

POWER USERS' CLINIC

Sorting with a Custom List

Most of the time, you'll want to stick with the standard sorting orders. For example, you'll put numbers in numeric order, dates in chronological order, and text in alphabetical order. But not always. For example, you may have good reason to arrange the categories in Figure 14-12 in a different order that puts more important categories at the top of the table. Or you may have text values that have special meaning and are almost always used in a specific nonalphabetical order, like the days of the week (Sunday, Monday, Tuesday, and so on) or calendar months (January, February, March, April, and so on).

You can deal with these scenarios with a custom list that specifies your sort order. In the Order column, choose Custom List. This choice opens the Custom List dialog box, where you can choose an existing list or create a new one by selecting NEW LIST, and then typing in your values. (Page 76 has more on creating specialized lists.) Figure 14-13 shows an example.

Custom list sorting works best when you have a relatively small number of values that never change. If you have dozens of different values, it's probably too tedious to type them all into a custom list.

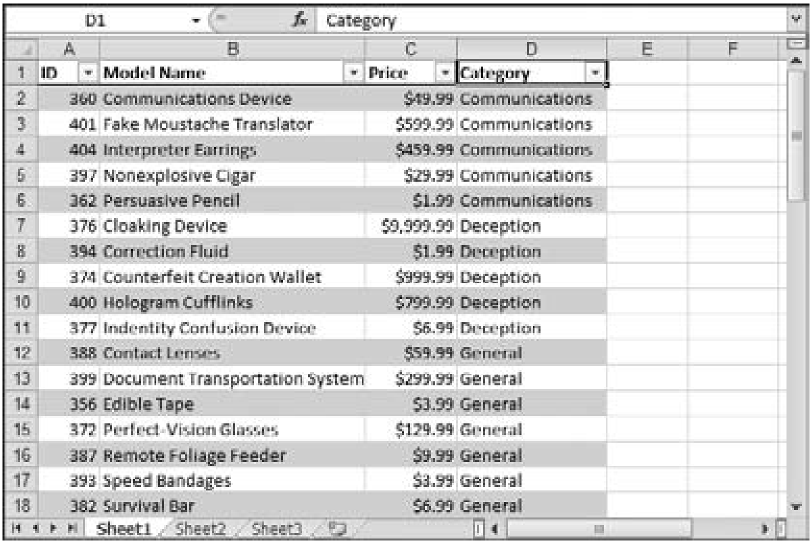

	A	B	C	D
1	ID	Model Name	Price	Category
2	360	Communications Device	$49.99	Communications
3	401	Fake Moustache Translator	$599.99	Communications
4	404	Interpreter Earrings	$459.99	Communications
5	397	Nonexplosive Cigar	$29.99	Communications
6	362	Persuasive Pencil	$1.99	Communications
7	376	Cloaking Device	$9,999.99	Deception
8	394	Correction Fluid	$1.99	Deception
9	374	Counterfeit Creation Wallet	$999.99	Deception
10	400	Hologram Cufflinks	$799.99	Deception
11	377	Indentity Confusion Device	$6.99	Deception
12	388	Contact Lenses	$59.99	General
13	399	Document Transportation System	$299.99	General
14	356	Edible Tape	$3.99	General
15	372	Perfect-Vision Glasses	$129.99	General
16	387	Remote Foliage Feeder	$9.99	General
17	393	Speed Bandages	$3.99	General
18	382	Survival Bar	$6.99	General

Figure 14-12: *The worksheet shows the following sort's result: alphabetically ordered categories, each of which contains a subgroup of products that are themselves in alphabetical order.*

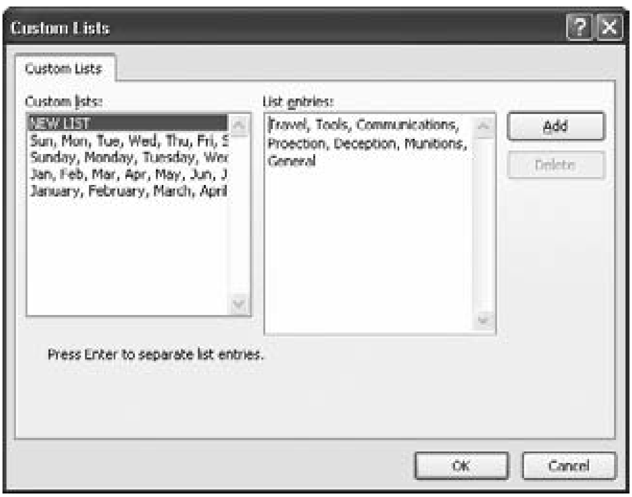

Figure 14-13: *Using a custom list for your sort order, you can arrange your categories so that Travel always appears at the top, as shown here. Once you've finished entering a custom list, click Add to store the list for future use.*

Sorting by Color

One of Excel's weirdest sorting options is sort-by-color, which lets you arrange your records based on the *color* of each cell. This oddball feature may come in handy if you've specifically highlighted certain cells that you're interested in (for example, you've given them a hot pink background fill). But this feature is really designed to work in conjunction with conditional formatting (page 186). You can create a conditional format that changes the foreground or background color of the cells you're interested in. Then, you use color sorting to make sure those cells rise to the top.

You can do color-based sorting in two ways. The quickest option is picking out cells with one background color, and then sending them to the top. To sort this way, open the drop-down color list, choose Sort By Color, and then choose your color. Excel helps out by letting you pick from a list that has all the background colors that are applied to cells in the current column. (This list doesn't include the background color that's applied from the table style.)

If you want to pick out multiple colors, you need to create a custom sort, and create a separate sort key to pick out each color. So if you want to sort cells in the order red, green, yellow, and then no color, you need to add three sorting keys. Choose Data→Sort & Filter→Sort, and then add three rules, as shown in Figure 14-14. Using the Sort On column, you can choose to sort using the background color (choose Cell Color) or the font color (Font Color).

The custom sort feature also works with icon sets (page 579) if you choose Cell Icon from the Sort On column. Same idea, but instead of picking a color that should rise to the top of the list, you pick one of the icons from your icon set.

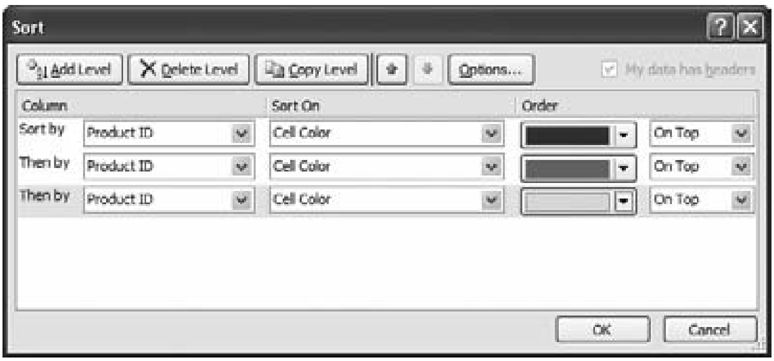

Figure 14-14: *In this table, rows with a red background appear at the top, followed by green and then yellow rows (followed by any rows that don't have a background color). If you want to send a color to the bottom of a table instead of the top, change On Top to On Bottom.*

Filtering with the List of Values

Sorting is great for ordering your data, but it may not be enough to tame large piles of data. You can try another useful technique, *filtering*, which lets you limit the table so it displays only the data that you want to see. Filtering may seem like a small convenience, but if your table contains hundreds or thousands of rows, filtering is vital for your day-to-day worksheet sanity.

Here are some situations where filtering becomes especially useful:

- To pluck out important information, like the number of accounts that currently have a balance due. Filtering lets you see just the information you need, saving you hours of headaches.
- To print a report that shows only the customers who live in a specific city.
- To calculate information like sums and averages for products in a specific group. You can use a function like SUBTOTAL(), described on page 423, to perform calculations using only the cells that are currently visible.

In theory, you could use the search feature in the data form window to find the records that interest you. However, while the search feature's perfect for finding one or two matches in a sea of data, it's much less suited to finding a large number of rows, because there's no way you can see them all at once, or print them out on a single sheet of paper.

Automatic filtering, like sorting, uses the drop-down column headings. When you click the drop-down arrow, Excel shows a list of all the distinct values in that column. Figures 14-15 and 14-16 show how filtering works on the Category column.

To remove a filter, open the drop-down column menu, and then choose Clear Filter.

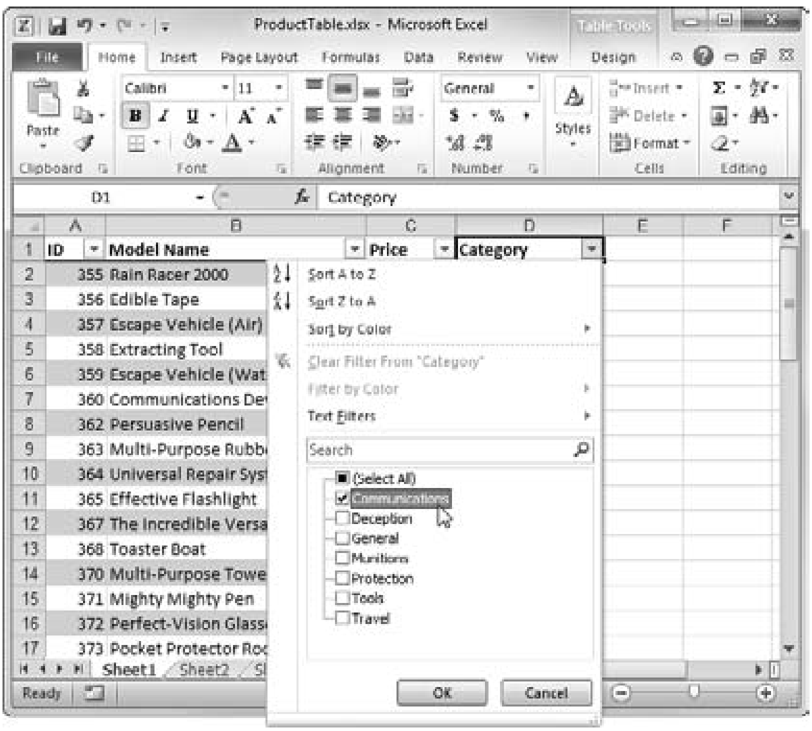

Figure 14-15: *Initially, each value has a checkmark next to it. Clear the checkmark to hide rows with that value. Or, if you want to home in on just a few items, clear the Select All checkmark to remove all the checkmarks, and then choose just the ones you want to see in your table. Either way, remember to click OK to make your choices take effect.*

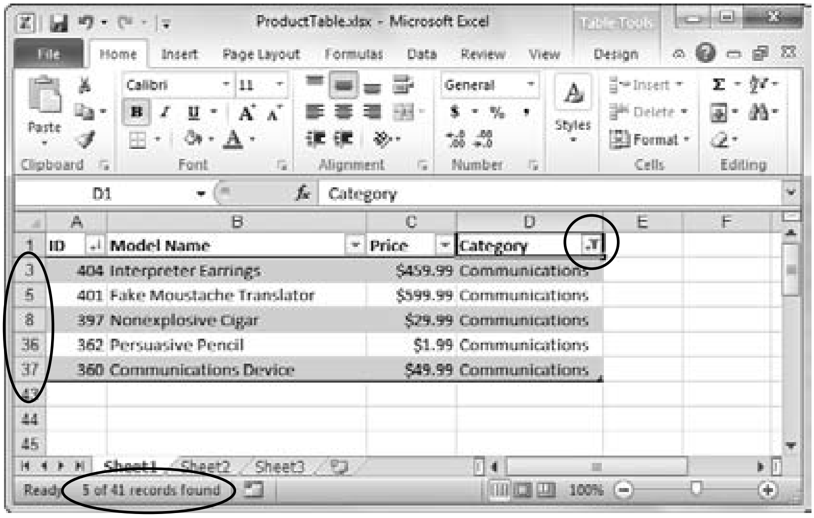

Figure 14-16: *If you select Communications and nothing else from the Category list in the product table example, the table displays only the five products in the Communications category. Excel takes several steps to make the situation more apparent. First, it adds a strange funnel icon to the Category header. Second, it switches the row number color from black to blue for the filtered records (in this case, rows 2 through 6). Third, Excel adds a status bar message that indicates the number of filtered matches and the total number of records in the list.*

Creating Smarter Filters

The drop-down column lists give you an easy way to filter out specific rows. However, in many situations you'll want a little more intelligence in your filtering. For example, imagine you're filtering a list of products to focus on all those that top $100. You could scroll through the list of values, and remove the checkmark next to every price that's lower than $100. What a pain in the neck that would be.

Thankfully, Excel has more filtering features that can really help you out here. Based on the type of data in your column (text, a number, or date values), Excel adds a wide range of useful filter options to the drop-down column lists. You'll see how this all works in the following sections.

Filtering dates

You can filter dates that fall before or after another date, or you can use preset periods like last week, last month, next month, year-to-date, and so on.

To use date filtering, open the drop-down column list, and then choose Date Filters. Figure 14-17 shows what you see.

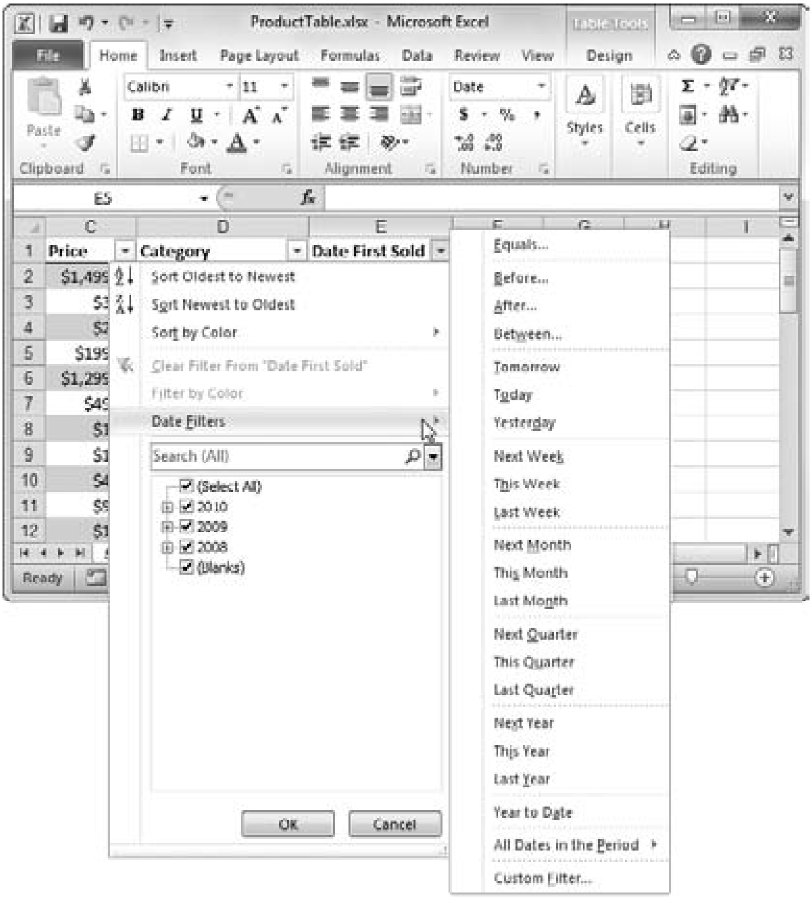

***Figure 14-17:** Shown here is the mind-boggling array of ready-made date filtering options you can apply to a column that contains dates. For example, choose Last Week to see just those dates that fall in the period Sunday to Saturday in the previous week.*

Filtering numbers

For numbers, you can filter values that match exactly, numbers that are smaller or larger than a specified number, or numbers that are above or below average.

To use number filtering, open the drop-down column list, choose Number Filters, and then pick one of the filter options. For example, imagine you're trying to limit the product list to show expensive products. You can accomplish this quite quickly with a number filter. Just open the drop-down column list for the Price column, and then choose Number Filters→Greater Than Or Equal To. A dialog box appears where you can supply the $100 minimum (Figure 14-18).

WORKAROUND WORKSHOP

The Disappearing Cells

Table filtering's got one quirk. When you filter a table, Excel hides the rows that contain the filtered records. For example, in Figure 14-17 you'll notice that the row numbers jump straight from 8 to 36, indicating that all the rows in between are hidden. In fact, all Excel really does is shrink each of these rows to have a height of 0 so they're neatly out of sight. The problem? When Excel hides a row, it hides all the data in that row, *even if the data isn't a part of the table.*

That property means that if you place a formula in one of the cells to the right of the table, then this formula may disappear from your worksheet temporarily when you filter the table! This behavior is quite a bit different from what happens if you delete a row, in which case cells outside the table aren't affected.

If you frequently use filtering, you may want to circumvent this problem by putting your formulas underneath or above the table. Generally, putting the formulas above the table is the most convenient choice, because the cells don't move as the table expands or contracts.

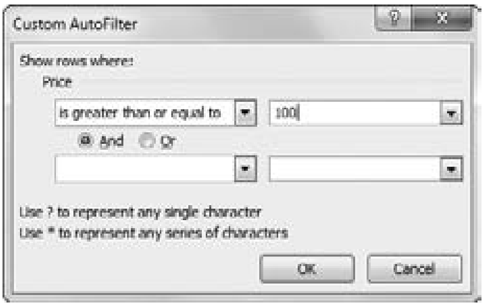

Figure 14-18:
This dialog box lets you complete the Greater Than Or Equal To filter. It matches all products that are $100 or more. You can use the bottom portion of the window (left blank in this example) to supply a second filter condition that either further restricts (choose And) or supplements your matches (choose Or).

Filtering text

For text, you can filter values that match exactly, or values that contain a piece of text. To apply text filtering, open the drop-down column list, and then choose Text Filters.

If you're performing filtering with text fields, you can gain even more precise control using wildcards. The asterisk (*) matches any series of characters, while the question mark (?) matches a single character. So the filter expression *Category equals T** matches any category that starts with the letter T. The filter expression *Category equals T????* matches any five-letter category that starts with T.

Dealing with Duplicate Rows

Hard-core table types know that every once in a while, despite the utmost caution, a duplicate value slips into a table. Fortunately, Excel has tools that let you find duplicates, wherever they're hiding, and remove them.

Highlighting Duplicates

It's not too hard to fish out these duplicates—one option is to use sorting (described earlier) on the column where you suspect a duplicate exists. Then, if you spot two identical values, you can delete one of the table rows (page 402). Of course, in order for this technique to work, you have to be ready to scroll through all the records and check each one. In a supremely long list, that job could take some time.

Fortunately, Excel has another solution—conditional formatting. You've already used it to highlight important data (in Chapter 6), but you can also use it to make repeating values stand out like sore thumbs. Here's how:

1. **Select the table column (as described on page 404) that you want to check for duplicate values.**

   For example, you could select the Product ID column to look for products that have the same ID value.

Note: You can highlight more than one column, but if you do, Excel highlights identical values that appear in more than one column. For example, if the same number appears in the Product ID column and in the Price column, Excel highlights it even though it isn't really a duplicate.

---

2. **Choose Home→Styles→Conditional Formatting→Highlight Cells Rules→Duplicate Values.**

   When the Duplicate Values dialog box appears, choose the type of formatting you'd like to use to highlight repeated values. People often choose to change the background color.

3. **Click OK.**

   Excel changes the background color of all values that appear more than once in the selected column (or columns), as shown in Figure 14-19. Conditional formatting keeps working even *after* you've applied it. So, if you add a new record that duplicates the value of an existing record in the column you're checking, Excel immediately highlights it. It's like having a duplicate value cop around at all times.

Tip: If you have an extremely large table, you may like to use color-based sorting (page 410) to bring the duplicate records to the top of the table. For example, if you highlighted duplicate Product ID values with a light red background fill, you click the Product ID column's drop-down arrow, choose Sort By Color, and then pick the same color (which automatically appears in the menu).

D43 Travel

	A	B	C	D	E
1	Product ID	Model Name	Price	Category	
2	355	Rain Racer 2000	$1,499.99	Travel	
3	356	Edible Tape	$3.99	General	
4	357	Escape Vehicle (Air)	$2.99	Travel	
5	358	Extracting Tool	$199.00	Tools	
6	359	Escape Vehicle (Water)	$1,299.99	Travel	
7	360	Communications Device	$49.99	Communications	
8	362	Persuasive Pencil	$1.99	Communications	
9	363	Multi-Purpose Rubber Band	$1.99	Munitions	
10	364	Universal Repair System	$4.99	Tools	
11	359	Escape Vehicle (Water)	$1,299.99	Travel	
12	367	The Incredible Versatile Paperclip	$1.49	Munitions	

Sheet1 Sheet2 Sheet3

Figure 14-19: *Conditional formatting helps smoke out a product with the same name—"Escape Vehicle (Water)"—that happens to be in two different places in the table.*

Removing Duplicates Automatically

Once you've found duplicate records, it's up to you what to do with them. You can leave them in your table or delete them by hand. However, if you don't want to keep the duplicates, Excel has a quicker solution, thanks to a built-in feature that hunts for duplicates and automatically removes the offending rows. Here's how to use it:

1. **Move to any one of the table cells.**

   Although you can technically use the duplicate removal feature with any range of cells, it works best with tables because you don't need to select the full range of data yourself.

2. **Choose Data→Data Tools→Remove Duplicates.**

   Excel shows the Remove Duplicates dialog box, where you can choose which columns to search for duplicates.

3. **Decide how many columns need to match in order for Excel to consider the record a duplicate. Add a checkmark next to each column you want to inspect.**

   For example, you may decide that any record with the same Product ID is a duplicate. When Excel scans the table, it removes any subsequent record that has the same Product ID as an earlier record, even if the rest of the data is different.

   On the other hand, you may decide that you want to match several columns to prevent the possibility for error. For example, there may be two different products that share the same Product ID due to a minor typo. To avoid this possibility, you could inspect several columns (as shown in Figure 14-20), so that records are removed only if every column value matches. (You can't do this maneuver with conditional formatting and the Highlight Duplicates rule. When you use conditional formatting, you're limited to finding duplicates in *one* column—if you include more than one column, Excel treats them as one big batch of cells.)

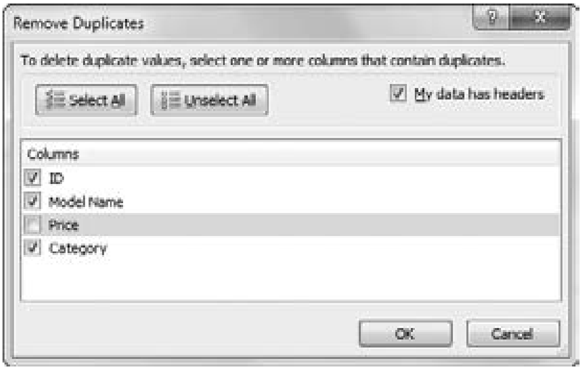

Figure 14-20: *In order to be considered a duplicate in this operation, the record has to have the same Product ID, Model Name, and Category. (The Price can vary.) If you need to select all the columns in a hurry, use the Select All button.*

4. **Click OK to remove the duplicates.**

   Excel scans your table looking for duplicates. If it finds any, it keeps the first copy and deletes those that appear later on in the table. When Excel is finished, it pops up a message box telling you how many rows it removed and how many still remain.

   You don't get a chance to confirm the deletion process, but you can reverse it by using the Undo feature (hit Ctrl+Z) immediately after it finishes.

Remember, when using the Remove Duplicates feature, you have no way of knowing what records Excel's deleted. For that reason, people often use the Highlight Duplicates rule first to check out the duplicates and make sure they don't belong.

Note: The Remove Duplicates feature and the Highlight Duplicates rule don't work in exactly the same way. For instance, Remove Duplicates treats repeated empty cells (blank values) as duplicates, and removes them. The Highlight Duplicates feature ignores empty cells and doesn't highlight them.

Performing Table Calculations

Excel tables have several nifty features that help you out when you're performing calculations on the data inside a table.

One example's the way that automatic expansion works when you're adding new columns that contain calculations. For example, say you've got a table with three columns: Product ID, Model Name, and Price. If you want to add a new column that tracks the discounted price of each product (say, 90 percent of the regular price), then a table's a great timesaver. Once you've created the new discount price column for one row (see Figure 14-21), Excel fills the same calculation into every other row (Figure 14-22). You wind up with a table that shows the discounted price of every product. No copying and pasting required.

Figure 14-21: *This worksheet shows a new column with a formula that's in the process of being entered. See Figure 14-22 for what happens next.*

Figure 14-22: *Once you finish the formula shown in Figure 14-21, and then hit Enter, Excel expands the table to include the new column, and then fills the formula down to include every row. And if you scroll to the bottom of the table and start adding a new row, Excel is intelligent enough to automatically copy this new formula into the new row. In other words, once you add a formula to a column, that formula is sticky.*

Once you've added a new column with a calculation, the only thing left to do is to give it a good name. Click the column header cell (which will say something like Column1), and then type in something more relevant (like Discount Price).

Dynamic Calculations

One of Excel tables' nicest features is the way they handle calculations. You can build formulas that use the data in a table, and Excel adjusts them automatically as the table grows or shrinks.

For example, imagine creating a table that stretches from row 1 to row 5. The first row contains the column headers, followed by four product rows. You create the following formula to add the total price of all the items:

```
=SUM(C2:C5)
```

If you now add a new item (that is, a new row) to the table, Excel automatically updates your formula so it includes this new item:

```
=SUM(C2:C6)
```

The same magic happens if your table shrinks because you've deleted an item. You just need to make sure that your range includes the whole column when you first create the formula. For example, consider the following formula, which omits the first item in the table:

```
=SUM(C3:C5)
```

If you expand the table now, Excel doesn't modify the formula. Fortunately, the program's got your back: It'll remind you if you create a formula that includes only a portion of your table by showing a green triangle in the corner of the cell, which is Excel code for "Is this really what you want to do?" When you move to the cell, an exclamation mark icon appears. Click it, and you get a short explanation and a menu of error-fixing options, as shown in Figure 14-23.

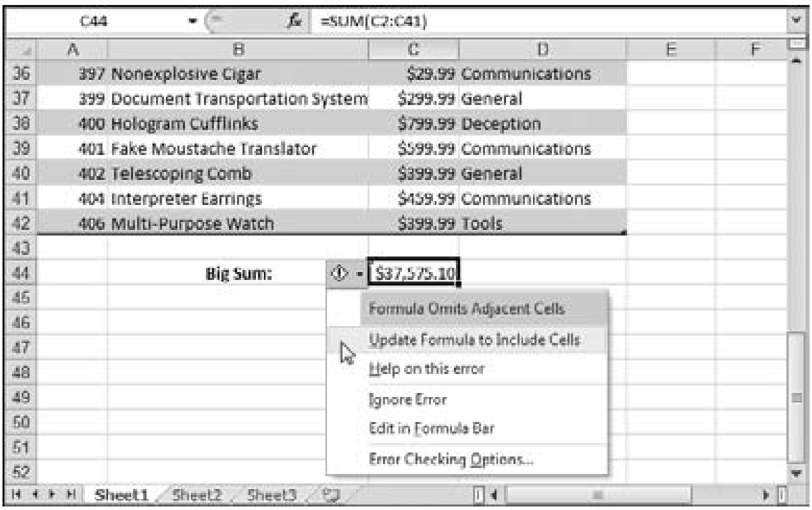

Figure 14-23: *Excel warns you if you create a function that leaves out part of the information in a table. For example, the formula in C44 mistakenly left out cell C42. Excel offers to remedy the problem—just choose Update Formula to Include Cells.*

Filtering settings don't affect your formulas. If you filter the table so that only some products are visible, the SUM() function still produces the same result, which is the total price of all products. If you don't want this behavior, you can use the *total row* (page 422) or Excel's SUBTOTAL() function (page 423).

Column Names

So far, the table calculation examples you've seen have ignored a super-convenient feature called *column names*. You've already learned how Excel lets you create your own named ranges to refer to frequently used parts of your worksheet. (Page 371 has the full story.) But even if you don't create any named ranges of your own, you'll find that Excel equips every table with some built-in names.

First of all, Excel creates column names that you can use when you write formulas inside your table. For example, look back at Figure 14-21, which uses this formula to discount the current price:

```
=C2*0.9
```

When Excel copies this formula down the table, it adjusts the formula automatically using the familiar magic of relative references (page 252), so that each discounted price refers to the product price in the same row. For example, in row 10 (technically, the ninth row of data), the formula becomes:

```
=C10*0.9
```

But there's a shortcut. Instead of referring to the specific cell, you can use the name of the column (which is whatever text you've placed in the column header) inside square brackets. That shortcut means you can write the formula like this:

```
=[Price]*0.9
```

The [Price] name automatically refers to the value of the Price column in the current row, no matter where in the table you stick it.

Column names aren't only major timesavers, they also make your worksheets much easier to understand. After all, who could mistake the following formula?

```
=[Price]*[Tax Rate]+[Shipping Charge]
```

Note: If you use Excel's point-and-click formula creation, Excel uses the column names rather than the cell reference.

Table Names

You can't use column names in formulas that are in other parts of your worksheet. After all, Excel would have no idea what row you're trying to use. However, Excel gives you another shortcut. You can refer to your entire table anywhere on your worksheet *by name*.

This raises one excellent question—namely, how does Excel decide what name your table should have? When you first create a table, Excel picks a rather unexciting name like Table1, Table2, and so on. To change this name, click anywhere inside your table, and then edit the text in the Table Tools | Design→Properties→Table Name box. For example, ProductList makes a good name.

On its own, your table name refers to the entire range of cells that contains the data for your table. That means it includes the entire table, minus the column headers. You may want to use these cells with a lookup function like VLOOKUP(), as described on page 341. For example, if you want to get the price (from the third column) of the product named Persuasive Pencil, you can use this easy-to-read formula:

```
=VLOOKUP("Persuasive Pencil", ProductList, 3, FALSE)
```

Excel includes another treat that lets you dig deeper into your table. You can use the table name *in conjunction* with a column name to get the range of cells that holds the data for just one column. Here's an example:

```
=SUM(ProductList[Price])
```

This gets the ProductList table, pulls out the cells in the Price column, and then passes them to the familiar SUM() function, which generates a total. You can quickly and efficiently get a hold of parts of a table. Figure 14-24 shows you how Excel helps you create formulas like these with its Formula AutoComplete feature.

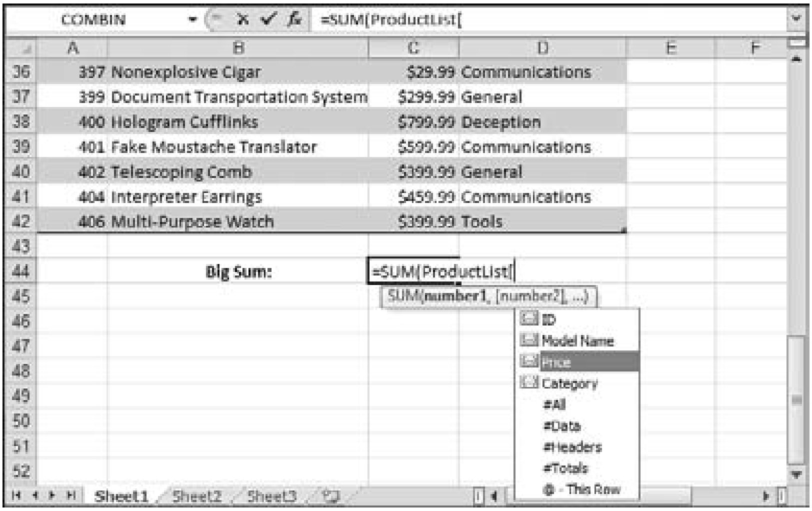

Figure 14-24: *Once you type a table name followed by an open bracket, Excel shows a list of columns in that table. You can select one using the up and down arrow keys, and then insert it into your formula by pressing Tab.*

The Total Row

Excel tables make it easy to calculate totals, averages, standard deviations, and other common formulas by using a dedicated summary row. To show this row, just select Table Tools | Design→Table Style Options→Total Row. Excel adds an extra row at the bottom of the table.

When the total row first appears, it shows only one piece of information: the number of records currently displayed. If you want to show some other type of information, choose a column, and then click the total row cell at the bottom of that column. A drop-down list appears with preset options. Choose one, and the total row displays the calculation in that cell, as shown in Figure 14-25.

Of course, you can create your own formulas to show the same information as the total rows do. However, the total row requires no work. It also uses only the rows that are currently visible, ignoring all filtered rows. You may or may not want this behavior. (It's great if you want to calculate totals for a small subset of filtered data, but it's not so good if you want to create grand totals that include everything.) If you don't want this behavior, you can write your own summary formula using functions like SUM(), COUNT(), and AVERAGE().

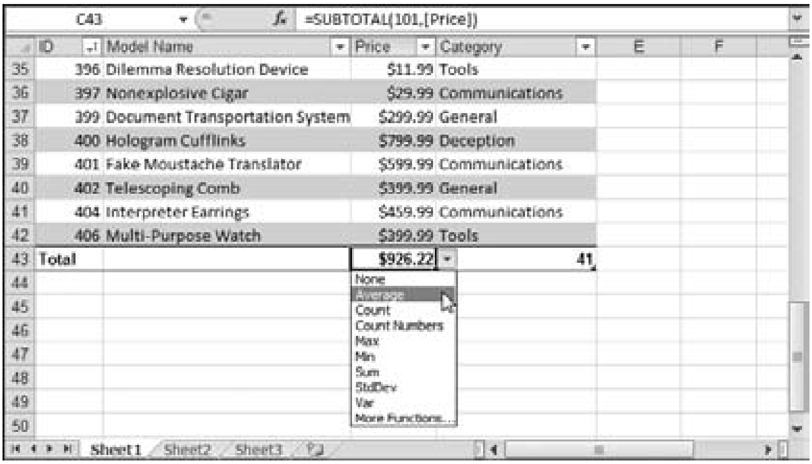

***Figure 14-25:** The Total row lets you perform common calculations with a single mouse click. In this example, the Total row displays the average price of the records that match the filter condition.*

The SUBTOTAL() Function

The total row conveniently works only with the currently visible rows in a table. To see this phenomenon in action, simply click one of the cells in the total row. If you look in the formula bar, you see that these cells use the SUBTOTAL() function to perform their calculations. That's because the SUBTOTAL() function is the only Excel function that takes table filtering into account.

The SUBTOTAL() function is the perfect solution for all the calculations in the total row, including sums, averages, counts, and more. The trick is that the first argument of the SUBTOTAL() function is a numeric code that tells Excel what type of calculation it should perform, while the second argument is the range of cells for the entire table column, from the first row to the last.

```
SUBTOTAL(function_code, column_range)
```

Table 14-1 lists all the function codes you can use with SUBTOTAL(). Note that each calculation type actually has two function codes associated with it. Function codes *above* 100 ignore hidden rows (rows that you've hidden using the Home→Cells→Format→Hide & Unhide→Hide Rows command). These function codes are the ones the total row uses. Function codes *under* 100 don't ignore hidden rows, but they still ignore rows you've filtered out. Figure 14-26 shows the SUBTOTAL() function in action.

Table 14-1. Function codes for SUBTOTAL()

Function used for calculation	Code that ignores hidden rows	Code that includes hidden rows
AVERAGE()	101	1
COUNT()	102	2
COUNTA()	103	3
MAX()	104	4
MIN()	105	5
PRODUCT()	106	6
STDEV()	107	7
STDEVP()	108	8
SUM()	109	9
VAR()	110	10
VARP()	111	11

In Figure 14-26, compare the results of the SUBTOTAL() functions in cells C2 to C6 to the formulas in cells D2 to D6, which just use the regular functions, like AVERAGE(), SUM(), MAX(), MIN(), and so on. In this figure, the formulas show the totals after filtering's limited the table to items in the Communications category.

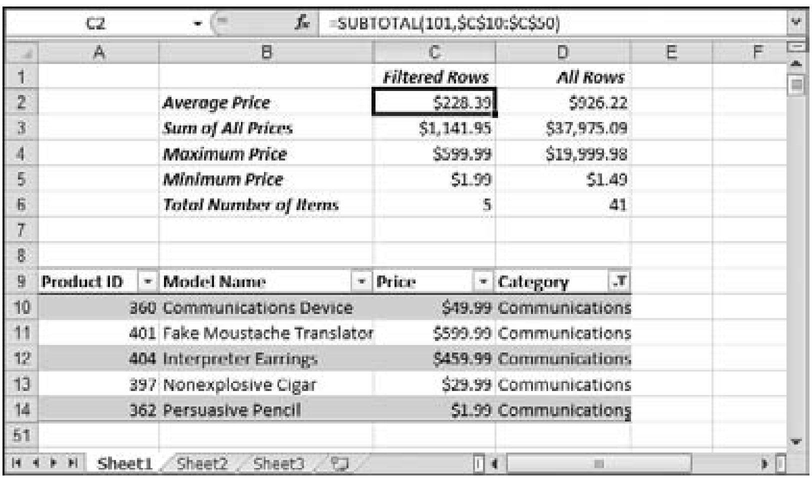
C2 =SUBTOTAL(101,$C$10:$C$50)

	A	B	C	D	E	F
1			Filtered Rows	All Rows		
2		Average Price	$228.39	$926.22		
3		Sum of All Prices	$1,141.95	$37,975.09		
4		Maximum Price	$599.99	$19,999.98		
5		Minimum Price	$1.99	$1.49		
6		Total Number of Items	5	41		
7						
8						
9	Product ID	Model Name	Price	Category		
10	360	Communications Device	$49.99	Communications		
11	401	Fake Moustache Translator	$599.99	Communications		
12	404	Interpreter Earrings	$459.99	Communications		
13	397	Nonexplosive Cigar	$29.99	Communications		
14	362	Persuasive Pencil	$1.99	Communications		
51						

Figure 14-26: *This worksheet puts the SUBTOTAL() function to work calculating averages, sums, and more. The formulas that use the SUBTOTAL() function are in cells C2 to C6, while the formulas that use the closest ordinary Excel function (and don't take filtering into account) are in cells D2 to D6.*

The Database Functions

Excel also includes functions exclusively for use with long tables. These are the *database functions*, a set of 11 functions that let you analyze groups of data.

The database functions are very similar to the basic statistical functions like SUM(), AVERAGE(), and COUNT(). In fact, the database functions have the exact same names, but with an initial letter D—so you find a DSUM(), DAVERAGE(), DCOUNT(), and so on. The database functions differ from their nondatabase counterparts in that they can selectively filter out rows. In other words, when you use DSUM(), you can specify a set of criteria that a record must match in order to be included in the sum. (The filtering that you apply with the AutoFilter feature makes no difference to the database functions. They don't ignore hidden rows.)

Table 14-2 lists the database functions, along with comparable statistical function.

Table 14-2. *Database functions*

Function	Similar to	Description
DAVERAGE()	AVERAGE()	Calculates the average in rows that meet the specified criteria.
DCOUNT()	COUNT()	Counts the number of rows that meet the specified criteria.
DCOUNTA()	COUNTA()	Calculates the number of non-blank values in rows that meet the specified criteria.
DGET()	No equivalent	Returns the value that meets the specified criteria. If more than one value matches, DGET() returns the *#NUM!* error. If no records match, it returns the *#VALUE!* error.
DMAX()	MAX()	Returns the maximum value in rows that meet the specified criteria.
DMIN()	MIN()	Returns the minimum value in rows that meet the specified criteria.
DPRODUCT()	PRODUCT()	Calculates the product produced by multiplying all values in rows that meet the specified criteria.
DSTDEV()	STDEV()	Calculates the standard deviation in rows that meet the specified criteria.
DSUM()	SUM()	Calculates the sum of values in rows that meet the specified criteria.
DVAR()	VAR()	Estimates the variance of a sample population in the rows that meet the specified criteria.
DVARP()	VARP()	Estimates the variance of an entire population in the rows that meet the specified criteria.

DGET() is the only function without a statistical counterpart. DGET() works a little like the VLOOKUP() and HLOOKUP() functions (page 341), and it returns a single value in a row that meets the specified criteria.

Each database function uses the exact same three parameters:

```
DFUNCTION(table_range, field, criteria)
```

- The *table_range* is the range that the function uses. The table range should include the entire table, including the column you want to use for your calculation and the columns to use for your criteria. The table range has to include the column headers, because that's how the database functions identify each column and match it up with the criteria. That means you can't use the automatically generated table names you learned about earlier (page 421).
- The *field* is the name of the column you want to use for the calculation. For example, if you're using DSUM(), the field is the numeric column you want to total. Excel scans the column headers until it finds the column that has the same name.
- The *criteria* is a range of cells that specifies all the conditions that rows must meet to be included in the calculation. This range can be as large as you want, and you're free to define conditions for multiple columns, or multiple conditions that apply to the same column. If the range contains only empty cells, the database function operates on all the items in the table.

To use a database function successfully, you need to create a suitable range of cells that you can use for criteria. Excel expects this range of cells to be in a strict arrangement. Here are some rules you'll need to follow:

- Each condition needs two cells. One cell specifies the name of the field, and the other specifies the filter condition. For example, you can enter "Category" for the field name and "Tools" for the filter condition.
- The cell with the filter condition must be directly under the cell with the field name. If you put them side by side, the database function just gives you an error.
- You can add as many conditions as you want. You can even use multiple criteria that impose different conditions on the same columns. However, you must place them in columns next to each other. You can't stack them on top of one another.

Filter conditions follow the same rules that you used for search criteria in the data form window. Thus, you can use comparison operators like less than (<) and greater than (>) to create conditions like *<500* (all prices under $500) or <>*Travel* (all products not in the Travel category). If you don't specify a comparison operator, Excel acts as though you've chosen the equal sign. In addition, it allows partial matches with text criteria. For example, the category criteria T will match both Travel and Tools.

You have total freedom to decide where on the worksheet to place your filter cells. Figure 14-27 shows an example that puts the filter cells at the top of the worksheet. To find the average price of all the products that match these criteria, you use the following formula:

```
=DAVERAGE(A9:D51, "Price", C1:E2)
```

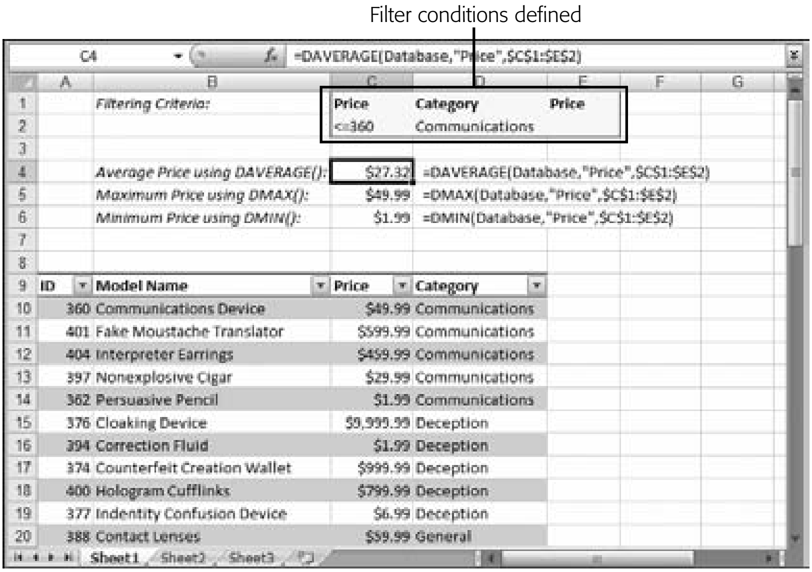

Figure 14-27: *In this example, you could use cell E2 to add another condition specifying a minimum price. However, because this cell is currently blank, Excel doesn't apply a price filter. Instead, it filters by matching a category and setting a maximum price.*

The database functions are nice because they force you to define your criteria on the worksheet. If you want to change the criteria, you simply need to modify the condition cells, which makes database functions a perfect solution for building dynamic reports.

CHAPTER 15

Grouping and Outlining Data

As you saw in the previous chapter, Excel's tables are great tools for managing big collections of data made up of long, uniform columns (like lists of customers or products). But tables don't work so well if you need to show more complex information, especially if it's split into separate groups, each with its own subtotals.

Imagine a company sales report that lists a year's worth of quarterly results for each of its regions around the world. If you try to cram all this data into one long list—including subtotals for each region—you're likely to end up with a spreadsheet that looks like a numerical version of Twister. You'd be better off breaking out each region into a separate group of cells, and then tying everything together with another batch of cells that sums it all up.

Excel's grouping and outlining features are perfect for dealing with multilayered information like this. They help you quickly and easily expand and collapse big chunks of data and, in the process, make calculating summary information much easier. These tricks are remarkably easy to implement, but they rank as one of Excel's best-kept productivity secrets.

Basic Data Grouping

When you want to simplify your worksheets, you first have to learn how to *group* data. Grouping data lets you tie related columns or rows into a single unit. Once you've put columns or rows into a group, you can *collapse* the group, temporarily hiding it and leaving more space for the rest of your data. At its simplest, grouping is just a quick way to help you easily view what you want to see, when you want to see it. It's most handy when you want to show all the key summary information but none of the numbers that you need to crunch to arrive at these results.

Note: Excel's grouping settings also affect how your worksheet prints out—collapsed rows or columns don't appear in your printout.

UP TO SPEED

Comparing Grouping and Tables

Overall, grouping isn't as powerful as Excel's tables (covered in Chapter 14). Whereas tables include features for sorting, filtering, and searching, groups simply try to make it easier to work with data tables that use one or more levels of subtotals.

Of course, there's no reason you can't use *both* grouping and tables. You may group one or more columns in a big table so that you can see just the columns you want (all the columns that track address components, for instance).

This approach works well if you're grouping multiple columns, but it doesn't work as well if you're grouping rows. Excel tables grow dynamically, so you could have a problem. If you add a new row at the bottom of a table (or a new column on either side), the table automatically expands to incorporate the new information. Groups don't have the same behavior. As a result, if your table grows, some of it may slip out beyond an edge of the group.

Creating a Group

To see how grouping can simplify complex worksheets, check out the sales report data shown in Figure 15-1. In this example, the information fits easily into the viewable area of the Excel window, but in a real-world company, you could easily end up needing to extend the worksheet with more columns and rows.

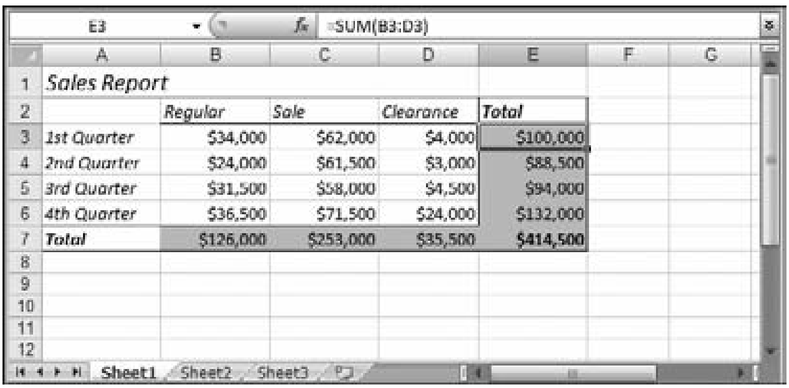

	A	B	C	D	E
1	Sales Report				
2		Regular	Sale	Clearance	Total
3	1st Quarter	$34,000	$62,000	$4,000	$100,000
4	2nd Quarter	$24,000	$61,500	$3,000	$88,500
5	3rd Quarter	$31,500	$58,000	$4,500	$94,000
6	4th Quarter	$36,500	$71,500	$24,000	$132,000
7	Total	$126,000	$253,000	$35,500	$414,500

Figure 15-1: *This worksheet shows the products sold at a retail store, with the numbers broken down in two ways: by quarter, and by the type of merchandise (regular priced, on-sale, and clearance). The data is sub-totaled so you can tell how much the company sold in a particular category over the whole year (cells B7:D7), how much it sold in a quarter in all categories (cells E3:E6), and how much it sold in total (cell E7).*

You could add a variety of columns (representing different types of promotions, different discounts from regular prices, or different departments in the store). Or, you may add more rows to cover sales quarters from more than one year, or to track monthly, weekly, or even daily sales totals. In either case, the data would become fairly unwieldy, forcing you to scroll up and down the worksheet, and from side to side. By grouping columns (or rows) together, you can pop them in and out of view with a single mouse click.

Note: You may remember from Chapter 5 that you can also use *hiding* to temporarily remove rows and columns from sight. But grouping is a much nicer approach—you can more easily pop data out of and back into view without going hunting through the menu. It also makes it more obvious to the person reading the worksheet that there's more data tucked out of sight. People more often use hiding to remove gunk that no one *ever* needs to see. Technically, grouping uses Excel's hiding ability behind the scenes to make grouped cells disappear temporarily.

Here are the steps you need to follow in order to group several columns together using the sales report example:

1. **Select the columns or rows you want to group.**

   In most cases, you'll want to group the *detail columns* that provide fine-grained information (like columns B, C, and D). You shouldn't include in your group any subtotals (like column E) that summarize your information, because you're likely to want to see those subtotals.

   To make your selection, drag over the column headers so that you've selected the entire column. This way, Excel can tell right away that you want to group columns. In the sales report example, you'd choose columns B, C, and D.

2. **Choose Data→Outline→Group.**

   If Excel isn't sure which groups of cells you want to group, it shows a Group dialog box that lets you specify whether you're grouping columns or rows. (If you selected cells B3:D3, Excel would consider it an ambiguous selection because the selection could represent an attempt to group the columns *or* rows.) If the Group dialog box appears, make your selection, and then click OK.

When you group your columns, the worksheet doesn't change, although a new margin area appears at the top of the worksheet, as shown in Figure 15-2. This margin lets you collapse and expand your groups (see Figure 15-3). To collapse a group, click the minus sign (–), which then changes into a plus sign (+). To expand a collapsed group, click the plus sign (+).

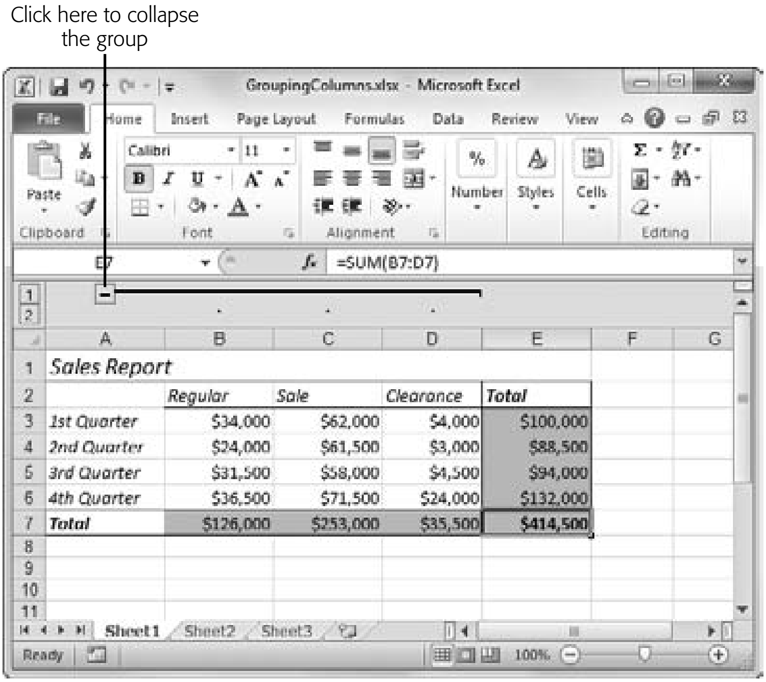

Figure 15-2: *Here, columns B, C, and D have been grouped together.*

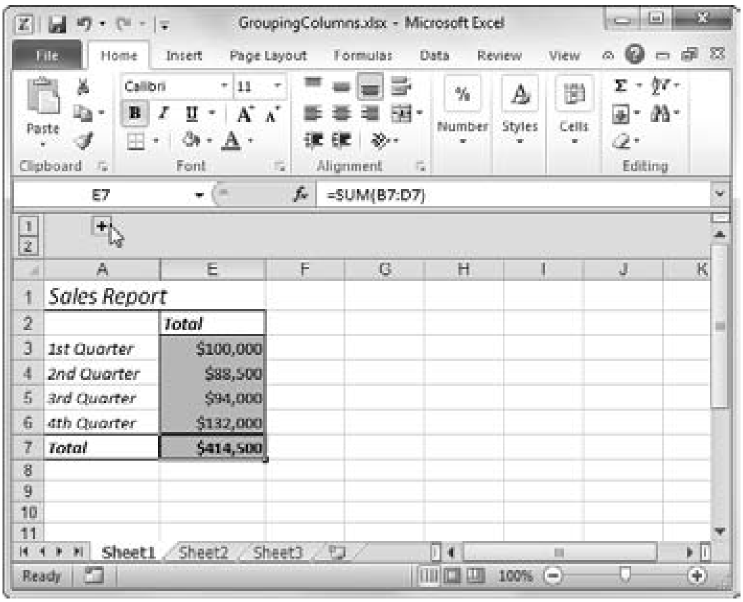

Figure 15-3: *After clicking the minus icon in the worksheet's upper margin (as shown in Figure 15-2), you can quickly collapse this group, hiding columns B, C, and D. When you hide columns, Excel still uses them in calculations (like those in column E).*

To remove a group after you've established it, expand the group, select all the columns, and then choose Data→Outline→Ungroup. If you select only some of the columns, Excel removes the selected columns from the group, but leaves the rest of the group intact—as long as at least two grouped columns remain. (If you remove a column from the middle of a group, you're actually left with two groups, one with the columns on the left, and one with the columns on the right.)

Tip: You can also ungroup a set of columns while the group's collapsed. Just select a range of columns that includes the hidden columns, and then use the Data→Outline→Ungroup command. Even though you've removed the group, the columns remain hidden and trapped out of sight! You can expose these columns after you've removed the group only if you select a range of columns that includes the hidden columns, and then choose Home→Cells→Format→Hide & Unhide→Unhide Columns.

You can use the same technique to group rows. In the sales report, you may want to group all the detail rows that give the quarter-by-quarter sales numbers, so that you can quickly collapse them and just see the totals. To do so, select rows 3 to 6 by clicking the first row number, and then drag down over all the rows you want to select. Then, select Data→Outline→Group. This time, a margin appears just to the left of the worksheet, letting you quickly hide the grouped rows. You can also have groups of rows *and* groups of columns, as shown in Figure 15-4.

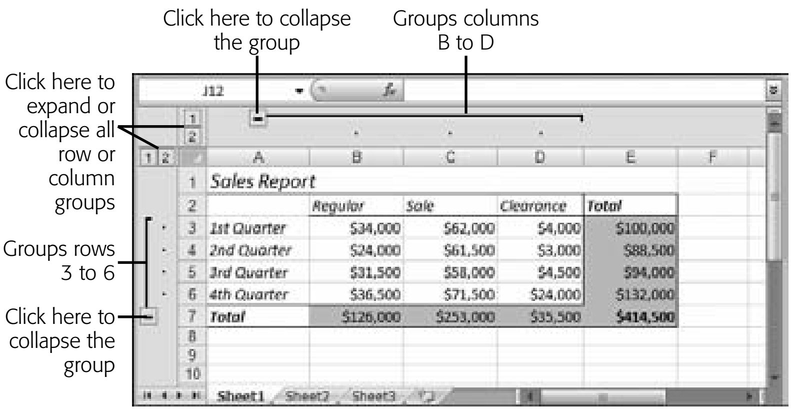

Figure 15-4: *The trick to mastering Excel's grouping abilities is understanding how the grouping bars work. The grouping bars show you the range of cells that are bound together in a single group. The minus sign (–) icon always appears just outside the group (either over the column to the left, or under the last row). That's because this icon has to remain visible after you've collapsed the group so that you can expand it.*

Note: If you insert a new row between two grouped rows using the Home→Cells→Insert→Insert Sheet Rows command, Excel automatically places the new row into the group. The same's true if you insert a new column between two grouped columns using Home→Cells→Insert→Insert Sheet Columns.

You can add as many or as few groups as you like. If you have some data with far too many columns, then you can put all the extra columns into a single group, or into several separate groups. What's the difference? With one group, you can hide all the extra information with a single mouse click. With multiple groups, you'll need to do more clicking—but you have the ability to choose exactly what you want to see.

Keep in mind that if you create adjacent groups of columns, you always need to leave at least one column between each group. If you create two groups next to each other (like columns A to C and columns D to F), Excel automatically merges them into a single group. (This little quirk exists because Excel needs a free column to display the expand/collapse icon.) The same holds true for groups of rows.

Tip: If you decide to use multiple groups, it makes sense to group related rows or columns. For example, you may group all the columns that have address information into one group.

Nesting Groups Within Groups

As you've seen, grouping lets you temporarily shrink the size of a large table by removing specific rows or columns. Grouping becomes even *more* useful when you create a worksheet that has multiple tables of information, like the one shown in Figure 15-5.

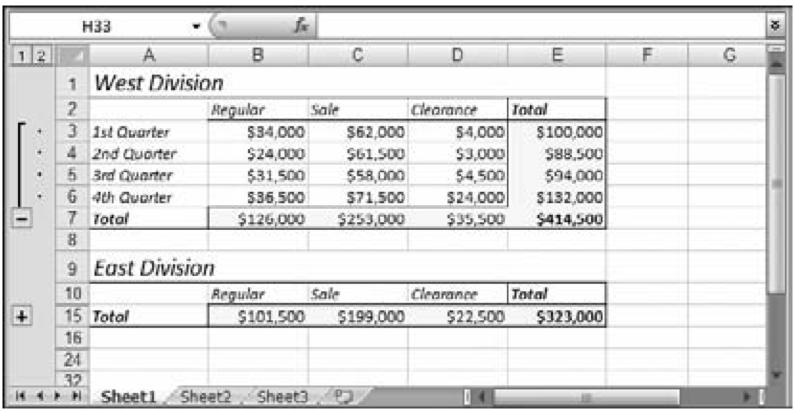

Figure 15-5: *This worksheet includes two separate groups with data. Currently, the lower group is collapsed, so that you see only the totals. Thanks to grouping, you can fill a worksheet with dozens of separate tables with detailed information, but you can still see the summary information for them all when you collapse the groups.*

Note: Don't get confused by thinking about Excel's structured tables, which you saw in Chapter 14. With grouping, you can use any combination of cells. Here, the data falls into a tabular structure with easily identifiable rows and columns, but it isn't an official Excel table.

Not only can you create separate groups for different tables, you can also put one group *inside* another. This approach gives you multiple viewing levels. Consider Figures 15-6 and 15-7, which show two views of the same worksheet. This worksheet includes four tables, each of which has two levels of grouping. By collapsing the first-level group for a table, you can hide everything except the summary information (listed in the Total line). By collapsing the second-level group, you can hide the entire table.

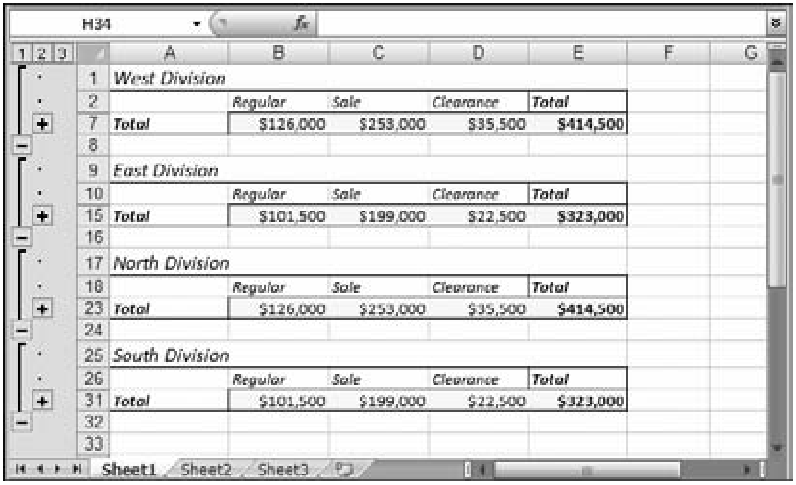

Figure 15-6: *All four tables are partially collapsed, and each has its second-level groups hidden so that they show just the total sales information for each region.*

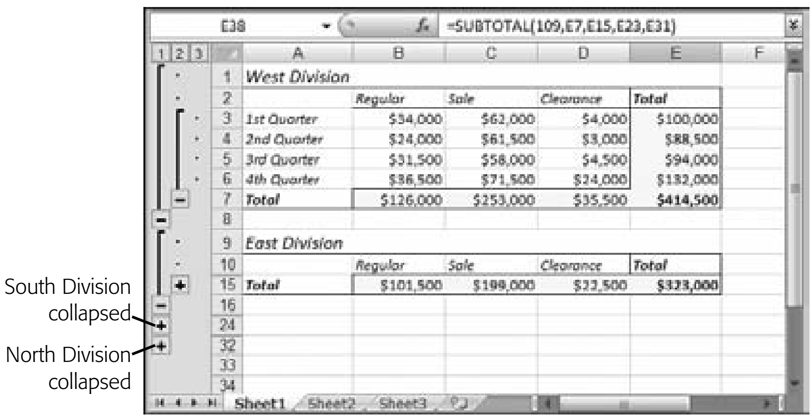

Figure 15-7: *You see only two tables. The West Division is completely visible, because both its levels are expanded. The East Division is partially visible because its second-level group is collapsed. The North and South Division tables aren't visible at all because the first-level group for each table is collapsed.*

Note: The only drawback is that the more levels of grouping you add, the more room Excel needs in the margins at the side of the worksheet in order to display all the grouping lines.

Technically, when you add more than one level of grouping, you're adding collapsible *outline* views to your worksheet, which let you quickly move back and forth between a bird's-eye view of your data and an up-close glimpse of multiple rows.

TIMESAVING TIP

Collapsing and Expanding Multiple Groups at Once

If you create a worksheet with multiple levels of grouping, you can find yourself doing a lot of clicking to expose all the information you're interested in (or to hide all the details you don't care about). Fortunately, Excel includes a helpful shortcut that lets you expand or collapse multiple groups at once: *grouping buttons* (shown in Figure 15-8).

Each *grouping button* is labeled with a number. If you click number 1, Excel collapses all the column or row groups. If you click number 2, Excel collapses all the groups *except* the top level. If you click number 3, Excel collapses all the groups except the first two levels, and so on. In the sales worksheet shown in Figure 15-7, you can click 3 to show the whole worksheet, 2 to show just the subtotals, or 1 to hide all the sales tables.

Excel displays a different number of grouping buttons, depending on your worksheet's levels of grouping. The largest numbered button you'll see is 8, because Excel allows a maximum of seven levels of nested groups.

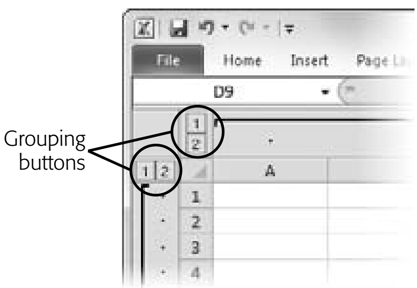

Figure 15-8:
Excel provides two sets of grouping buttons: one for row grouping (just above the row numbers), and one for column grouping (just to the left of the column letters).

Summarizing Your Data

Understanding how to quickly collapse and expand tables is all well and good, but before long, you'll want to *add up* the totals you've collected in these accordion-style tables. In the subtotaled sales report shown in Figure 15-7, the perfect way to complete this worksheet is to add a final table that sums up *all* the Total rows listed in the last row of each division.

To create a grand summary, you need to employ just a few more formulas. These formulas add the separate subtotals (contained in columns B, C, D, and E) to arrive at a final series of grand totals. The following formula would calculate the total sales in all divisions and all quarters for regular merchandise:

```
=B7+B15+B23+B31
```

Note: You could, of course, also calculate any of these subtotals by using the SUM() function, which is covered on page 268.

You'll notice that no matter how you expand or collapse the groups in your worksheet, the result of these formulas is always the same. That's because, whether you're using straight addition (of multiple cells) or the SUM() function, Excel takes into account visible *and* hidden cells. If you're using grouping, however, it may occur to you that it would be handy if you could perform a calculation that deals only with the visible cells. That way, you could choose the portions of a worksheet you want to consider, and the formula could recalculate itself automatically to give you the corresponding summary information.

Good news: Excel gives you the power to build this sort of formula. All you need is the awesome SUBTOTAL() function. As you may recall from Chapter 14, the SUBTOTAL() function is a useful tool for calculating totals in filtered lists. But it's just as helpful when you apply it to outlines because it ignores collapsed rows and columns.

As explained on page 423, the first argument (known as the calculation code) you use with the SUBTOTAL() function tells Excel the type of math you want to do (summing, averaging, counting, and so on). If you want to ignore the hidden cells in collapsed groups, you have to use the calculation codes above 100. If you want to perform a sum operation, then you need to use the code 109. (For a full list of the SUBTOTAL() calculation codes, see page 424.)

Once you've chosen the calculation code, you need to specify the cells or range of cells you want to add. Here's an example that rewrites the earlier formula to sum only visible cells:

```
=SUBTOTAL(109,B7,B15,B23,B31)
```

Figure 15-9 shows the difference between using the SUBTOTAL() function and the SUM() function.

Figure 15-9: *This worksheet calculates summary information using two different approaches. The formulas in row 37 use the SUM() function, while the formulas in row 38 use the SUBTOTAL() function. In the SUBTOTAL() calculation, Excel doesn't use the subtotals for the completely hidden tables, giving a different result than that provided by the SUM() function.*

Combining Data from Multiple Tables

Instead of writing your summary formulas by hand, you can generate a summary table automatically that takes advantage of Excel's ability to *consolidate* data. Consolidation works if you have more than one table with precisely the same layout. You can use consolidation to take the different tables shown in the sales report (West Division, East Division, and so on) and calculate summary information. Excel creates a new table that has the same structure but *combines* the data from the other tables. You can choose how Excel combines the data, including whether it totals, averages or multiplies the numbers, and so on.

Note: Data consolidation works with any sort of tabular data. You don't need to create the structured tables you learned about in Chapter 14.

Data consolidation works with any worksheet (with or without grouping); you can even use it to analyze data in identically structured tables from different worksheets.

Tip: For best results, you should consolidate data with the *exact same layout* only. Although you can coax Excel into combining differently sized ranges of data, it's all too easy to confuse yourself about what is and isn't combined. To make life easier, only consolidate ranges that have numeric data or identical labels (like column or row titles). Leave out the overall table title, because there's no way to consolidate it.

To consolidate the sales report data, follow these steps:

1. **Move to the location where you want to insert the summary table.**

   Excel inserts the summary table, starting at the current cell. Make sure you've scrolled down past all your data, so you don't overwrite important information.

2. **Choose Data→Data Tools→Consolidate.**

   The Consolidate dialog box appears, as shown in Figure 15-10.

3. **From the Function pop-up menu, choose how you want to combine numbers.**

   In the sales report, Sum is the best choice to calculate total sales. However, you may want to create separate tables that pick out the best or worst sales using Max and Min.

4. **Click inside the Reference text box. Now, drag to select the first table you want to consolidate in the worksheet.**

   If Excel's main window isn't already visible, click the icon at the right end of the Reference text box to collapse the Consolidate dialog box. You have to click this icon again to restore the window when you're done selecting the cells you want.

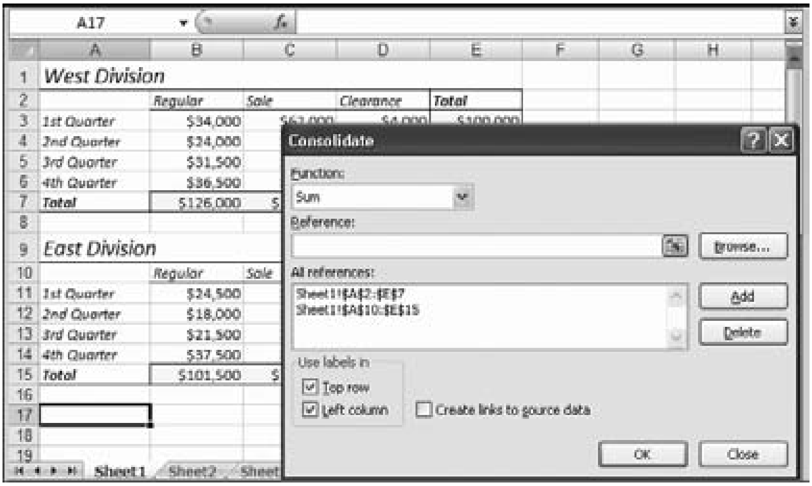

Figure 15-10: *This worksheet contains two detailed sales tables that are about to be consolidated. Both ranges have been added to the Consolidate dialog box (they're listed in the "All references" list), and the "Top row" and "Left column" checkboxes have been turned on so that Excel can find and use the headings on either side of the table.*

5. **Once you've selected the appropriate cells, click the Consolidate dialog box's Add button.**

   The range appears in the "All references" list.

6. **Return to step 4 to select the next table you want to consolidate.**

   Repeat steps 4 and 5 for each table you want to consolidate.

7. **If your selection includes labels (like row or column titles), select the "Top row" or "Left column" checkbox to tell Excel where the labels are.**

   If you don't tell Excel where the labels are, it ignores these cells, and the corresponding cells in the summary table wind up blank. But if you use these checkboxes to tell Excel where the labels are, Excel simply copies the labels directly to the summary table.

Note: If the labels don't match exactly in the ranges you're using, you may need to clear the "Top row" and "Left column" checkboxes to perform your consolidation. Otherwise, depending on the placement of these labels, Excel may refuse to consolidate your data.

8. **Click OK to generate the summary table.**

   When creating the consolidated data, Excel copies headings and calculates numbers, but it doesn't copy any of the source formatting. Figure 15-11 shows the result.

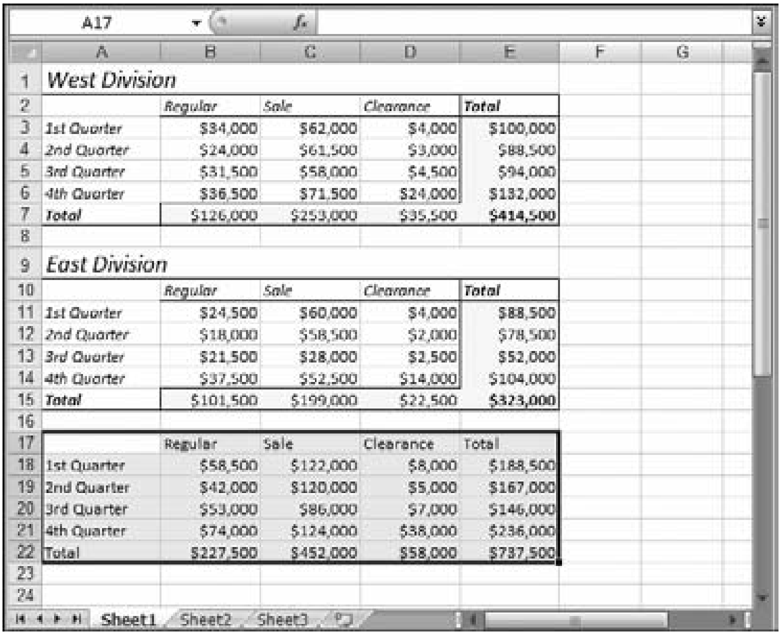

	A	B	C	D	E
1	West Division				
2		Regular	Sale	Clearance	Total
3	1st Quarter	$34,000	$62,000	$4,000	$100,000
4	2nd Quarter	$24,000	$61,500	$3,000	$88,500
5	3rd Quarter	$31,500	$58,000	$4,500	$94,000
6	4th Quarter	$36,500	$71,500	$24,000	$132,000
7	Total	$126,000	$253,000	$35,500	$414,500
8					
9	East Division				
10		Regular	Sale	Clearance	Total
11	1st Quarter	$24,500	$60,000	$4,000	$88,500
12	2nd Quarter	$18,000	$58,500	$2,000	$78,500
13	3rd Quarter	$21,500	$28,000	$2,500	$52,000
14	4th Quarter	$37,500	$52,500	$14,000	$104,000
15	Total	$101,500	$199,000	$22,500	$323,000
16					
17		Regular	Sale	Clearance	Total
18	1st Quarter	$58,500	$122,000	$8,000	$188,500
19	2nd Quarter	$42,000	$120,000	$5,000	$167,000
20	3rd Quarter	$53,000	$86,000	$7,000	$146,000
21	4th Quarter	$74,000	$124,000	$38,000	$236,000
22	Total	$227,500	$452,000	$58,000	$737,500

Figure 15-11: *Here, the newly created consolidated data table is at the bottom of the worksheet. In this case, someone used the Sum option, which means that the consolidated data shows the totals you get by adding the values from the separate tables. Cell B18 has the consolidated data for regularly priced merchandise sold in the first quarter in both the West and East Divisions (cells B3 and B11).*

Data consolidation's big disadvantage is that it generates a table filled with numbers, rather than formulas. Microsoft's engineers probably designed the consolidation feature this way so you can consolidate data from multiple files and not worry about losing the information if the source files move or their structure changes.

But this behavior means that if you modify any of the sales figures, the summary table doesn't update. Instead, you need to generate a completely new summary table by using the Data→Data Tools→Consolidate command. Fortunately, the second time around should be quite a bit faster, because Excel keeps track of all the ranges you selected for consolidation, so you don't need to define them again.

Grouping Timesavers

So far, you've learned how to tame an intimidating worksheet and neatly organize it into groups. Why stop there? This section covers a few other tools that, along with the grouping tools you've just mastered, will make your life easier. These tools include automatic outlining and subtotaling (both of which can create data groups automatically).

Auto Outline

Adding groups to a large worksheet can be tedious. You can leap over that tedium in a single bound with Auto Outline, a feature that gets Excel to examine your worksheet, and then create all the column and row groups you need automatically.

Like most automatic features, Auto Outline is a great solution when it works, and it's absolutely no help the rest of the time. Automatic outlining is an all-or-nothing affair, giving you either a grouped worksheet or the unhelpful error message, "Cannot create an outline". You also don't have any ability to configure how many levels of nesting Excel uses, how many groups it creates, and whether it creates column groups, row groups, or both.

With a bit of planning, however, you can set up your worksheets so they're Auto Outline-friendly. Here are the key points. Excel makes all its guesses about how to implement Auto Outline based on any formulas it finds that give it clues about the structure of your worksheet. If you have a formula that uses SUM() to total the cells in multiple columns in the same row, Excel assumes that it can place these columns in a single group so that you can collapse the details and leave just the totals visible. If Excel finds a formula that uses the data in the rows above it, the program assumes that a row group's in order. In other words, Auto Outline requires formulas to determine the grouping levels it should add. If you don't use any summary formulas at all, or if your worksheet uses a table, the Auto Outline feature doesn't work.

Note: If you don't have any formulas at all on your worksheet, you're guaranteed that Excel can't create an outline.

To use automatic outlining, select the portion of the worksheet that you want to outline, and then choose Data→Outline→Group→Auto Outline. Generally, it's easiest to select the whole worksheet (click the square just outside the top-left corner of the worksheet) to apply grouping to all your data at once.

If you don't like the results, you can't undo the outlining operation, but you can remove all the groups from your worksheet by choosing Data→Outline→Group→Clear Outline.

Automatic Subtotaling

Auto Outline isn't the only outlining trick that Excel has up its sleeve. The Subtotal tool lets you create groups *and* subtotals all in one click—a feature that can save you scads of time.

To use automatic subtotaling, you need to have a long table of repetitive data (product catalogs, sales transactions, calorie-counting meals on the Atkins plan). You can't use more than one table—instead, the goal is to get Excel to break your table into summarized subtables for you. Each table gets its own subtotal (for any or all of the columns that you choose) and all the subtotals get added together for your very own Grand Total.

To subtotal and outline your data simultaneously, follow these steps:

1. **Sort your data so that the rows you want to group are together in the list.**

   Before subtotaling your data, you may need to sort it into the right order. If you have a list of products and you want to subtotal the product information by category, you need to sort the table by category before you begin. This step ensures that all the products in each category are grouped together. For detailed information about sorting, see page 405.

2. **Once you have the list in the correct order, select the range of cells that contains the list.**

   Usually, the easiest option is to select each of the columns in your list by dragging across the column headers.

3. **Choose Data→Outline→Subtotal.**

   The Subtotals dialog box appears.

4. **In the "At each change in" pop-up menu, choose the field under which you want the subtotal to appear. The menu provides a list with all the column titles in your selection.**

   If, say, you wanted to group together (and create subtotals for) products that have the same Category label (as shown in Figure 15-12), you'd choose Category. Excel inserts a subtotal row each time the category label changes. (In the figure, the Communication products subtotal first, the Deception products next, and so on.) In order for this process to work, the list must be sorted—as you did in step 1—so that all records with the same category are already next to each other.

5. **In the "Use function" pop-up menu, choose the function you want to use to calculate the subtotal information.**

   This list includes everything supported by the SUBTOTAL() function, including averages, counts, and subtotals. Unfortunately, you can't total different columns using different functions.

Figure 15-12: *This worksheet has a long list of product data that's about to be subtotaled. Excel creates the subtotals for each group of products that has a different label in the Category column. The subtotal row includes the average price of the products in each group.*

6. **The "Add subtotal to" list includes all the column names in your selection. Put a checkmark next to each column you want to generate a subtotal for.**

   To calculate the subtotals, Excel uses the function you chose in step 5.

7. **If you want to start each group on a new page, turn on the "Page break between groups" checkbox.**

   This option works well if you have large groups, and you want to separate them in a printout.

8. **If you want to display the summary information at the end of group, choose "Summary below data." Otherwise, the totals appear at the beginning of the group.**

   If you're dividing products into category groups, choose "Summary below data" to make sure you'll see the listing of products, followed by the subtotal row.

9. **Click OK to group and subtotal the data.**

   Excel inserts a row in between each group, where it adds the formulas that calculate the subtotals for the group (Figure 15-13). All of the formulas use the SUBTOTAL() function. Excel also calculates grand totals for all the subtotals that it's generated.

Figure 15-13: *The product list, with subtotaling applied. Rows 7 and 13 contain the newly calculated subtotals. Excel has also added a grand total at the bottom of the list (which isn't shown in this figure).*

If you want to remove subtotals, you have two choices. The easiest approach is to choose Data→Outline→Subtotals to open the Subtotals dialog box, and then click the Remove All button. Alternatively, you can replace the existing subtotals with new subtotals. To do this, choose different options in the Subtotal dialog box, and make sure the "Replace current subtotals" checkbox is turned on.

CHAPTER 16

Templates

Spreadsheets are rarely one-of-a-kind creations. After you've built the perfect sales forecast, expense report, or personal budget, you'll probably want to reuse your hard work instead of starting from scratch. One approach is to save an extra copy of your workbook and just change the data for each new spreadsheet you want to create. That works fine, but it's not terribly convenient.

Excel provides a more streamlined option with *templates*, which are spreadsheet blueprints that you can use to create new files. Templates don't necessarily hold any data (although they can if you want them to). Instead, their main role in life is letting you format them to your heart's content—adding things like column titles, fancy shading, and complex formulas—so that every time you want a worksheet that looks like the template, all you have to do is select the template and voilà! A new file appears, containing all the design elements you created in the original template.

For example, you could create a monthly expense report template containing all the appropriate formulas and formatting you need, and use it to create a fresh report each month. If you ever need to change your report's formatting or calculations, then you simply modify the template, and all future expense reports will use the new version.

In this chapter, you'll learn about Excel templates, and how to get the most out of them. You'll begin by exploring how you can use Excel's existing set of templates, which provide a helpful starting point for all kinds of common spreadsheets. Then you'll learn how to create your own equally professional templates.

Understanding Templates

Before you begin using and creating templates, it's important to understand what they really are. Many template novices assume that templates are a special sort of file that's *similar* to Excel spreadsheets. However, the crafty individuals who created Excel actually designed template files so that they're *exactly the same* as spreadsheet files. (Whether this feature is a brilliant masterstroke of simple, elegant design or the product of terminally overworked programmers is up to you to decide.)

If you use a tool like Windows Explorer to look in a folder that contains spreadsheet files and template files, you'll notice that templates have different files extensions. Whereas ordinary Excel files have the extension *.xlsx* (or *.xlsm* if they contain macro code), templates have the extension *.xltx* (or *.xltm* with macros). For example, Invoice.xltx could be a template file for creating invoices, while Invoice-01.xlsx would be an actual invoice. However, you can put the exact same type of data in both files: elaborately formatted worksheets, formulas, numbers, text, and so on.

Note: Astute readers will notice that there's just a single letter of difference between the file format extensions. In Excel spreadsheets, the s stands for spreadsheet (.xlsx), whereas in template files the *t* stands for template (.xltx).

Based on this similarity, you might wonder why you should bother using templates at all. The difference is in *how* Excel handles templates. Here are some of the differences:

- Templates are usually stored in a special folder, so that the latest version of the template's always easily accessible.
- It's almost impossible to accidentally overwrite a template file. That's because Excel automatically creates a new worksheet when you select a template (more on how to actually use templates in a moment).
- Spreadsheet experts often create templates brimming with advanced Excel features (things like data validation tools, which you'll learn about in Chapter 24, and fancy formula design). These features help you prevent errors (like entering incorrect data or entering it in the wrong place) and provide additional guidance to others who use the template.

Finally, it's easy to distinguish template files (the blueprints) from real spreadsheets because of the different file extension they use. The icon also looks slightly different.

Most organizations and businesses maintain a group of templates that define things like standard layout and formulas for common types of spreadsheets like invoices and expense reports. Some organizations host these templates in a central location on a network file server or web server (the two best options), or just distribute them to each employee who needs them.

So how do you put templates to work? The trick is to understand your options when creating them.

Creating a New Workbook from a Template

So far, every example worksheet in this book has started from scratch, with nothing more than Excel's empty grid of cells. This approach is a great way to learn the nuts-and-bolts of how Excel works, but it's not always necessary. In many cases, you can find a template that matches the worksheet you want to create. If so, you'll save yourself a good deal of formatting and formula-writing work.

To use a template, begin by selecting File→New. Excel switches to backstage view, as shown in Figure 16-1. But here's the trick: instead of creating an ordinary, blank workbook (by choosing "Blank workbook", and then clicking Create), you click one of the other options to begin hunting for a suitable template.

Here are your choices:

- "Recent templates" shows a list of templates you've used recently. This gives you a handy way to jump straight to your favorite template. But at first this list is empty, because you haven't used any templates.

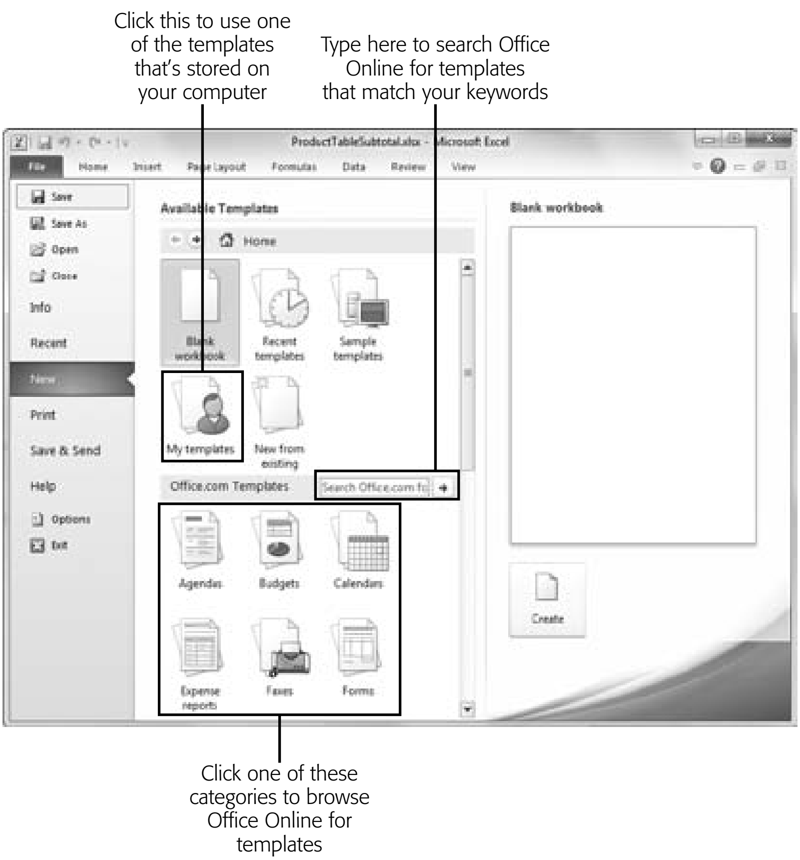

Figure 16-1: *Using Excel's backstage view, you can create a new, empty workbook or create a workbook based on a template.*

- "Sample templates" shows a small set of templates that come with Excel. This category has very few templates to choose from, but you'll find classics like Expense Report and Personal Monthly Budget.
- "My templates" lets you choose from one of the custom templates that you've created and saved on your computer, or one of the templates you've downloaded from Office Online.
- "New from existing" displays a dialog box that lets you choose a spreadsheet file. When you click OK, Excel opens a duplicate version of the selected file, giving it a new name. When you save the spreadsheet, Excel prompts you to supply a new name or to use the one it generated (a variation based on the name of the original workbook). Either way, you end up saving a new copy of the original file.

Note: The "New from existing workbook" option is conceptually similar to using a template, in that it creates a new workbook based on an existing workbook file. However, it's a better idea to create a workbook using a template for several reasons. First of all, because templates are stored in a central place, you don't need to hunt for the file you want to use. Additionally, if the template is properly fine-tuned, you don't have to bother removing old data.

- The categories under the heading "Office.com Templates" let you browse the vast treasure trove of freely downloadable templates that Microsoft provides on its website. You'll find hundreds of handy templates here, organized into logical categories like Agendas, Invoices, and Reports. As you'll see shortly, you can also use the search box to hunt for online templates by keyword (page 451).

These options are all you need to start using templates. Still feeling template intimidated? The next section walks you through all the steps you need to use a template.

Downloading Templates (Method 1: From Backstage View)

Excel's backstage view lets you choose from Microsoft's expansive online catalog of templates without leaving the comfort of Excel. Not only does this ability give you a way to get the latest template innovations, it also lets you dig up specialty templates like a secret Santa gift exchange list, a baseball scorecard, and a baby shower planner. Best of all, the whole process is so seamless you don't even notice you're downloading a template from the Web.

Note: Obviously, if you're using a computer that doesn't have an Internet connection, you're limited to the templates installed on your computer. However, you're not completely cut off from the rest of the word. You can surf to the Office Online website when you're connected, and then use your web browser to search and download Excel template files that seem interesting. You'll learn how in the next section.

Here's how it all goes down:

1. **Choose File→New.**

   Excel switches to backstage view and shows a list of template options.

2. **Choose one of the categories under the "Office.com Templates" heading. If none of the categories seems quite right, choose "More categories" to see some more exotic options.**

   A list of templates appears (Figure 16-2).

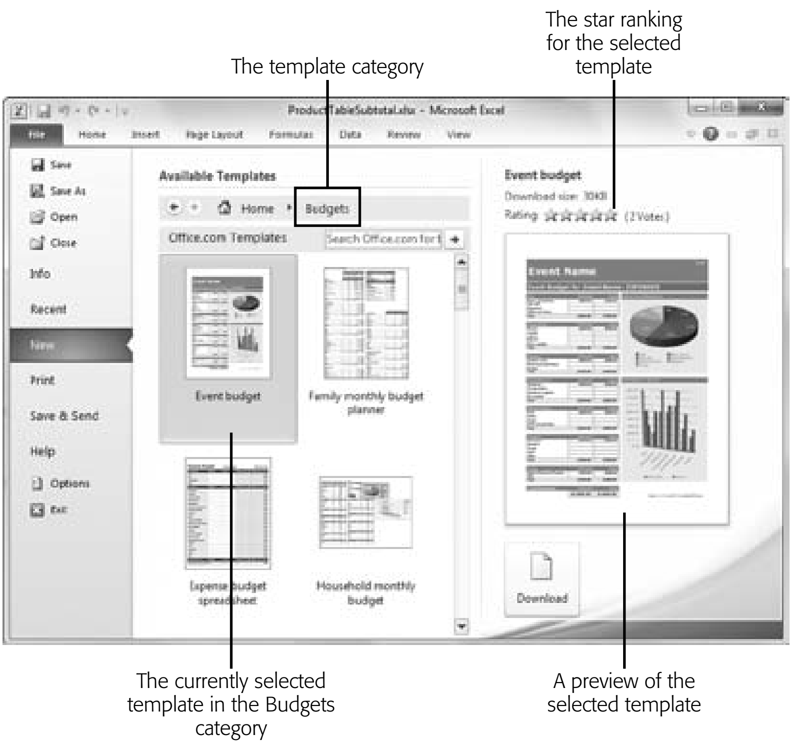

Figure 16-2: *Browsing by category is the fastest way to find a template that interests you. For example, under the Budget category you'll find templates for personal budgets (family budgets, wedding costs, and even a gardening budget) and templates for business budgets (marketing plan, business trips, operating expenses, and so on).*

3. **If you find a template that seems right, click it once to preview it.**

   The preview information appears at the far right of the window. The information includes the template's file size, its star rating (as ranked by other Office fans), and a preview of what it looks like.

Note: The template rating plays the same role as customer reviews on Amazon.com—other folks who have downloaded the template can give it a score of up to five stars (but they can't write a description).

Note that you can't rate a template from inside Excel. Instead, you need to surf to the Office Online site to submit a star rating. See page 450 for details.

---

4. **If you like the template, click Download to create a new workbook based on it.**

   A progress indicator appears while the template downloads to your computer. Because templates are quite small, it rarely takes more than a minute to download one. If you have a high-speed connection, you'll probably have the template in two or three seconds.

Note: When you download a template, Microsoft uses some fancy tricks to inspect your current installation of Windows and Excel to make sure you're not running pirated software. So, if you bought your copy of Excel for $7.99 from a street vendor in Chinatown, you probably can't download new templates.

After Excel downloads the template, it creates a new workbook based on that template (Figure 16-3). You can fill your data in this workbook, and, when you save it, Excel prompts you to choose a new file name. Excel suggests a name based on the name of the template. For example, if you select the loan amortization template, Excel suggests a name like "Loan Amortization Schedule1.xlsx" (you can, of course, change this name to anything you want).

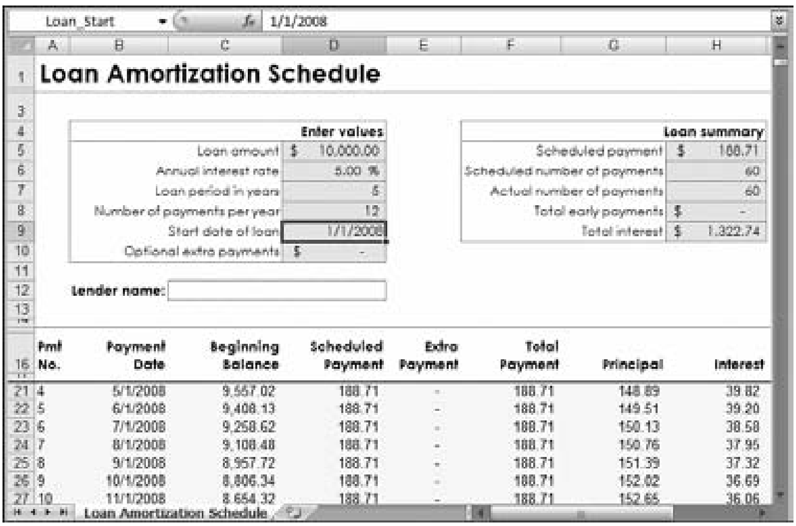

Figure 16-3: *When you choose the loan amortization template (from the Schedules category), Excel automatically creates a complete worksheet with all the calculations you need, showing each payment you'll need to make over the life of the loan. You need to enter only the loan amount, interest rate, duration, payment frequency, and start date (cells D5 to D9); Excel calculates the rest automatically using the financial functions you learned about in Chapter 10.*

Downloading Templates (Method 2: The Office Online Website)

Instead of getting all your templates served to you right in Excel, you can download them the old-fashioned way, from the Office Online website where they live. You may make the trip to Office Online for a few reasons:

- You want to download a template to use on another computer. Perhaps your home computer lacks Internet access, and you're using a friend's computer to get the template goodness you crave.
- You want to give template feedback (so yours can be one of the thousands of votes that make up a typical template ranking). You can also submit feedback with requests for new templates you'd like to see added to the collection.
- You want to search for all Office Online templates without worrying about what program uses them. Maybe you've decided you want a way to print out a calendar for 2011, but you don't know whether Excel or Word has the better template.

TIMESAVING TIP

Searching for a Template

The process described in the last section lets you find a useful template by digging through various categories. Excel also gives you another choice—you can search for a template by *keyword*. This approach is great if you want to quickly locate a template for a specific function or industry (you could search for "invoice" or "real estate"). Keywords aren't as useful if you just want to browse what's available. Of course, it's up to you to decide which way to go. In fact, many Excel gurus use both, depending on the task at hand.

To search by keyword, type your search word(s) into the "Search Office.com for templates" text box (Figure 16-4), and then press Enter. You can enter as many words as you want. Excel finds templates that contain *any* of the words you specify (it gives preference to templates that have *all* the search words). Try using a word that you think may appear in the template name (like "calendar") or a phrase in one of Microsoft's online categories (like "personal finance").

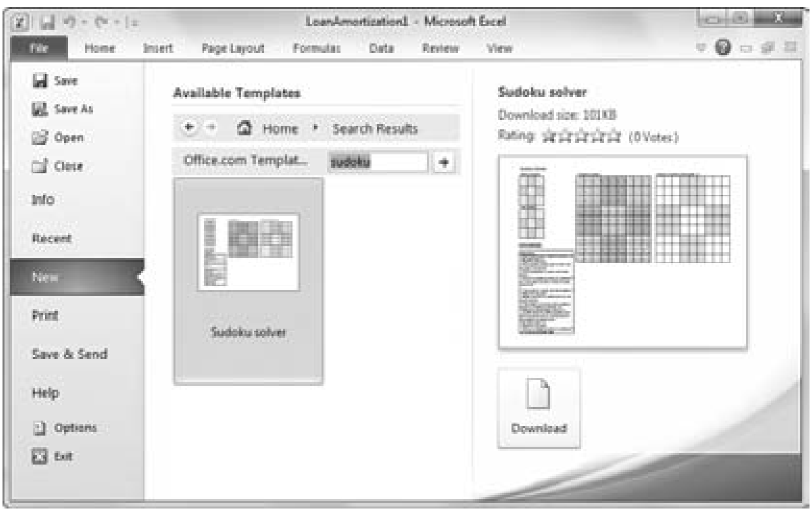

Figure 16-4: *In this example, a search for the word "sudoku" turns up an Excel template for solving Sudoku number puzzles.*

- You want to browse the other news and features on the Office Online website. If it's September, you may find quick links to some popular back-to-school lesson-planning templates.
- You're just happier in a web browser. After all, you've been staring at the Excel worksheet grid for 451 pages already.

Here's how to download a template from the Office Online website:

1. **Open a web browser, and then surf to** ***http://office.microsoft.com/templates*****.**

   It's best to use Internet Explorer, because it supports the ActiveX standard, which Office Online uses for its download-a-template feature. If you use a different browser that doesn't support ActiveX (like Firefox), you'll get a warning message explaining the issue when you try to download a template. You'll then be forced to download a compressed template file, which you must unzip on your own before you can use it. (The Office Online website provides more information about what to do, but you can save the headache altogether by using Internet Explorer.)

2. **Scroll down the web page to the Browse Templates section. Click a category link that interests you.**

   The top portion of the page is reserved for new and featured templates, some of which Microsoft created in response to requests (send in your request by clicking the "Suggest a template" link on the left side of the web page); other templates vary by season. For example, in the beginning of May, you'll find quick links to Mother's Day card and gift label templates. If you don't want to browse by category, you can perform a keyword search using the search box at the top of the page (and skip ahead to step 4).

3. **Depending on the category you've chosen, you may see another list of subcategories. If you do, just click the appropriate subcategory link.**

   You'll end up at a list that shows all the matching templates. The list shows the name, creator, and rating of each template.

Note: The Templates web page provides templates for all Office applications, including Word, Excel, Access, and InfoPath (a program that lets companies create souped-up data entry forms). Of course, you can't use a Word template in Excel, and vice versa. To tell which program a given template requires, look at the icon immediately to the left of the name. If you see the familiar Excel application icon (a green square with an "X" symbol), you can use this template to create new workbooks. Some templates (like greeting cards) are likely to be for Word, while others (like financial forms and worksheets) are usually for Excel.

4. **Choose a template file from the list by clicking its name.**

   A page appears with detailed template information (Figure 16-5). This information includes the size of the template, the required software version (most templates work equally well on all Excel versions since Excel 2000), and a preview graphic that shows what the template looks like.

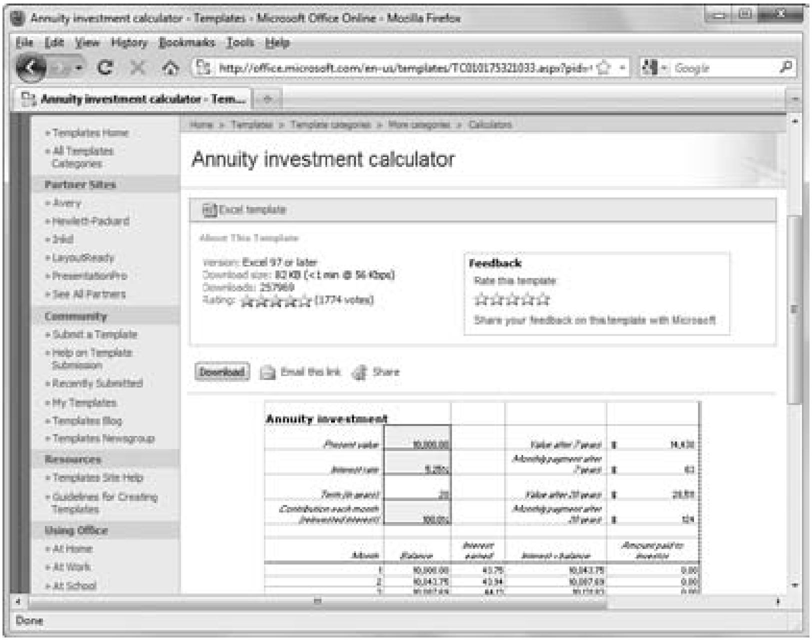

Figure 16-5: *Once you've chosen a template, you'll see a detailed page with information about the size of the file and the time it takes to download it—if you connect to the Internet using a telephone line and a meager 56K modem. (Template downloading is almost always very fast, regardless of your connection.) You'll also see an image that shows a surprisingly detailed preview of the template, complete with sample data.*

5. **If you're still happy with the template, click Download. Or, if you haven't found the right template, then you can click your browser's back button to navigate back to the template list.**

   You may be presented with a Microsoft license agreement. (This essentially states that you won't sue Microsoft if the template makes your computer explode. Don't worry; it won't.) To continue, click Accept.

   While the template is downloading, a pop-up browser window appears with a progress indicator. Once the template downloads (usually, just a fraction of a second later), a Save As dialog box will appear. Excel points the dialog box to your personal template folder (page 455), a location on the current computer where you can store all the templates you download and the templates you create.

6. **Click Save to save the template in your templates folder.**

   Optionally, you can give the template a new name or move to a different folder before you save it. For example, on page 459 you'll learn how to store templates in the workgroup templates folder, a location that lets multiple people use the same templates.

   What happens once the template downloads depends on whether you have the right program installed. If you're downloading an Excel template and you have Excel installed, then Internet Explorer launches Excel, and Excel creates a new workbook based on the template you chose, just as if you'd picked it from the backstage view. If you don't have Excel installed, you'll need to take the template file and email it or copy it to an Excel-enabled computer before you can use it.

Creating Templates

Ready-to-use templates are a fantastic innovation because they provide fine-tuned worksheets without forcing you to write a single formula. Of course, these templates also have a major drawback: No matter how clever Microsoft programmers are, they can't predict your every need.

For example, the Travel Services Invoice provides a generic worksheet that a travel agency may use to bill its customers. But what if you need to group different types of expenses separately, apply different discount rates to various groups, include a late fee, and tack on a few miscellaneous charges to pad your company's bottom line? If that's the case, you'd need to add your own formulas, restrictions, and formatting to the template. In that case, starting over from scratch and creating a template with the exact features you need is probably easier.

Fortunately, creating custom templates isn't difficult at all.

Tip: You can use one of Excel's templates as a starting point for your own templates. Just follow the process described earlier, edit the template (by adding information, changing formatting, and so on), and then save it as a new template. However, you'll generally find it easier (at least at first) to create your own templates from scratch.

Understanding Custom Templates

Earlier in this chapter, you learned how to create a new workbook by choosing a template in backstage view. In order to understand how to create your own templates, you need to know how this feature works—namely, how does Excel *find* the templates that are installed on your computer?

The answer is a little tricky, because Excel actually has three types of templates. They include:

- **Installed templates**. These templates are part of the Office package and are included automatically when you install Excel. You can find them (depending on where you installed Office) in a folder like *C:\Program Files\Microsoft Office\Templates\1033*. (1033 is the language ID for U.S. English.)

Note: This list includes only the templates on ***your*** computer. It doesn't include the templates from Office Online because those are stored on Microsoft's web servers.

UNDER THE HOOD

Tracking Down Templates

It's a good idea to take note of where your Templates folder is. This information helps when you want to share your templates with other people (or take their templates and copy them to your computer). Even if you aren't planning on sharing templates, it still makes sense to pay attention to their location, so you can back them up for safekeeping.

The Templates folder location varies depending on the current user name (the account under which you've logged into Windows). For example, if you've logged yourself in under the user account billjones on a Windows Vista or Windows 7 computer, you'll probably find the templates in *C:\Users\billjones\AppData\Roaming\Microsoft\Templates.* In Windows XP, the path is subtly different—probably *C:\Documents and Settings\billjones\Application Data\Microsoft\Templates*. Incidentally, the templates folder includes user-defined templates for *all* installed Office programs, including Word.

If you're having trouble tracking down your templates folder, there's a handy trick that can help you locate it. In Excel, choose File→Save As, and then select the "Excel Template (*.xltx)" file type. Excel will send you straight to the Templates folder.

- **Local templates**. These templates are the custom ones you create. (They're also known as "My templates.") Excel stores them in a special folder on your computer (more on that later). When you first install Excel, this category is empty, because you haven't made any custom templates yet.
- **Workgroup templates**. These templates are also custom templates. The only difference is that they're stored in a shared location where more than one person can access them. This way, other people using the same computer (or connected via a network) can make use of your hard work. When you first install Excel, the workgroup template folder doesn't exist yet. If you want to use this feature, you need to set it up yourself; see page 459 for details.

You can't change the list of installed templates. Even if you save your own custom files to the folder where the installed templates are stored, Excel steadfastly ignores them. For that reason, you should never try to add or remove an installed template.

On the other hand, you can freely add custom templates to the local and workgroup folders. Best of all, Excel automatically checks these folders. If Excel finds two custom templates in the local template folder, and three more in the workgroup template folder, then it shows all five in the "My templates" category. You can use these templates to create new spreadsheets in exactly the same way you use a built-in template.

So, all you need to know to integrate your custom templates with Excel's is where to save your files. As it turns out, Excel can take you there automatically, as you'll see in the next section.

WORD TO THE WISE

Creating Bulletproof Templates

By now, you probably realize that templates aren't just a way to eliminate repetitive work when you need to create similar spreadsheets. They're also a way to let ordinary people—those, like you, who aren't familiar with Excel's dark arts—to record information, fill out forms, and analyze data. These people need a little guidance, and templates are there to help them.

Unfortunately, Excel isn't always particularly forgiving. Even if you craft the perfect template, an Excel novice can accidentally delete or overwrite a formula just by pressing the wrong key. And it's almost as easy to put the wrong information in a cell (for example, by entering a date incorrectly so that it's interpreted as text). Furthermore, a template is no help at all when the person using it doesn't know where to start typing or what the different headings really mean. All these problems can happen (and regularly do happen), even if you think your template is a small miracle of straightforward design and organization.

When you want to create a truly bulletproof template, you can use two more Excel features:

- *Data validation* prevents people from entering the wrong type of data in a cell (or warns them when they do). It also lets you set up a handy drop-down list of values that the person editing the worksheet can choose from.
- *Worksheet protection* prevents people from changing certain cells, and (optionally) stops them from viewing the formulas inside.

Microsoft designed both these features for ordinary workbooks, but they make good sense in templates too. If you'd like to learn more about how they work, skip straight to Chapter 24 to get the full story.

Building a Custom Template

To create a custom template, you simply need to build a new workbook, add any headings, formatting, and formulas you desire, and then save it as a template. You can start this process from scratch by opening a new, blank workbook, or you can refine an existing built-in template. Either way, you should follow the same process of perfecting your workbook until it's ready for template status. Here are some tips:

- **Clear out the junk**. Your template should be a blank form politely waiting for input. Clear away all the data on your template, unless it's generic content. For example, you can leave your company name or the worksheet title, but it probably doesn't make sense to have sample numbers.
- **Assume formulas won't change**. The ideal template is one anyone can use, even Excel novices who are too timid to edit a formula. If you have a formula that contains some data that might change (for example, the sales commission, interest rate, late fee, and so on), *don't* type it directly into your formulas. Instead, put it in a separate cell, and use a cell reference within the formula. That way, the person using the template can easily modify the calculation just by editing the cell.

- **Don't be afraid to use lists and outlining**. These features are too complicated for many mere mortals (those who haven't read this book, for example), but they make spreadsheets easier to use and more powerful. By putting these advanced frills into the template, you ensure that people can use them in their spreadsheets without having to learn how to apply them on their own. Charts and pictures, which you'll learn more about in Part Four, are also good template additions.
- **Turn off worksheet gridlines**. Many templates don't use Excel's gridlines. That way, it's easier to see custom borders and shading, which you can use to draw attention to the important cells. To turn off gridlines, select View→Show→ Gridlines.
- **Add the finishing touches**. Once you have the basics—titles, captions, formulas, and so on—it's time to create a distinct look. You can add borders, change fonts, and inject color. (Just remember not to go overboard with cell shading, or the output may be impossible to read on a black-and-white printer.) You may also want to tweak the paper size and orientation to ensure a good printout.
- **Delete extra worksheets and assign good names to the remaining worksheets**. Every workbook starts with three worksheets, named Sheet1, Sheet2, and Sheet3. The typical template has only one worksheet, and it's named appropriately (such as Expense Form). For information about deleting and renaming worksheets, flip back to Chapter 4.
- **Consider adding custom macros to make a really slick spreadsheet**. For a real treat, you can build a toolbar with custom macros (and even attach them to custom buttons) that perform related tasks. For more information about how to build macros, see Chapter 28.

Once you've perfected your template, you're ready to save it. Follow these steps:

1. **Select File→Save As.**

   The Save As dialog box appears.

2. **In the "Save as type" drop-down list box at the bottom of the window, choose "Template (*.xltx)".**

   Excel automatically browses to the Templates folder, which is where the templates you create are stored on your computer. Typically, this folder is one like *C:\Users\UserName\AppData\Roaming\Microsoft\Templates* (where *UserName* is the name of the Windows account you used to log in).

3. **Type the template name, and then click Save.**

   The saved template file automatically appears in the "My templates" category. You can use it to build new spreadsheets. Just select File→New, and then choose "My templates". A dialog box appears with all the custom templates you've created or downloaded from Office Online, as shown in Figure 16-6.

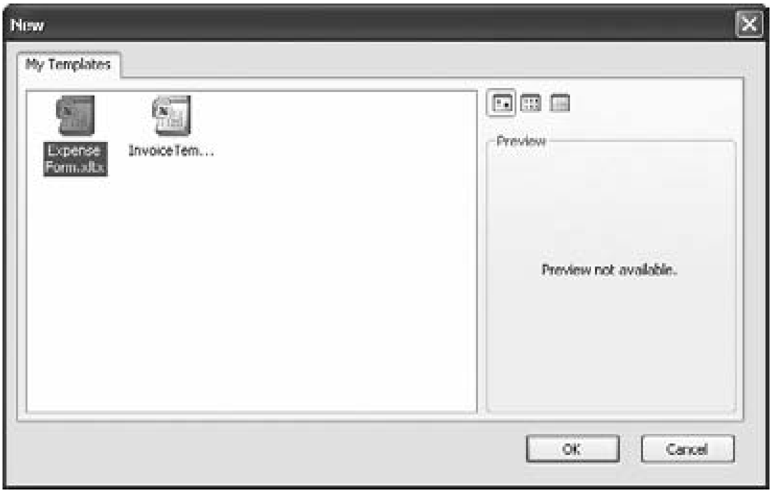

Figure 16-6: *This dialog box displays all your home-made templates. If you want to get rid of a template, right-click it in this window, and then choose Delete.*

Note: Of course, your templates don't *need* to appear in the "My templates" list. You can save a template file anywhere on your computer, and open it later by browsing to the right folder, and then choosing the template file. However, this approach has a couple of drawbacks. When you open a template file directly, Excel doesn't automatically create a new workbook. It's easy to overwrite the template by accident. Also, if your templates aren't in the "My templates" list, you may have trouble remembering where you put them.

Sharing Templates with Others

As you've already seen, once you place a template in the Templates folder, it appears in Excel's backstage view, where you can use it to create new workbooks. However, you may be disappointed to find out that you're the only one who can benefit from all your hard work. Coworkers using other computers won't be able to access your templates. In fact, your template won't even be available to other people who use the same computer if they log in with a different user name and password. That's because the Templates folder is a user-specific setting. You may use the folder *C:\Users\billjones\AppData\Roaming\Microsoft\Templates*, but if Sarah Cheng logs onto the same computer, then Excel bothers to check only *C:\Users\sarahcheng\AppData\Roaming\Microsoft\Templates*.

You can use several strategies to get around this problem.

- You can give a copy of your template to everyone who wants to use it. Choose any way you like to transfer the template (by email, on a floppy disk or CD, over a network drive, and so on), but make sure you tell folks to put it in their own local Templates folder, as explained earlier.

- You can create a workgroup templates folder. In this case, multiple people will use the same workgroup template folder, and they'll all be able to use the template files you've put there.

You can use the second option to share files between more than one user account on the same computer or, for even better results, on a network drive that a whole team of people can use. Unfortunately, Excel nearly throws a wrench into the whole process because it doesn't let you specify which folder to use for workgroup templates. However, Microsoft Word gives you a back door to the answer. Because Excel and Word both use the same template system, if you configure the workgroup templates folder in Word, it also takes effect in Excel. (Don't bother asking why Word provides this service and Excel doesn't—it's just one of those enduring Microsoft mysteries.)

Tip: If you're unlucky enough to not have Word installed, you still have one more painful option. You can write a custom Excel macro that sets the location of the workgroup templates folder. The line of code you need is `Application.NetworkTemplatesPath` = "[*path goes here*]", but you'll have to learn a fair bit about the intricacies of VBA (Visual Basic for Applications) programming before you can actually use it. For more information, see Chapter 29.

Once you accept the fact that you need to rely on Word, the process actually becomes quite easy. Here's what you need to do:

1. **Decide which folder you want to use as a shared folder.**

   You may want to create a new folder now. If you do, the easiest approach is to use a tool like Windows Explorer.

2. **Start Word, if it's not already running. In Word, choose File→Options.**

   The Word Options dialog box appears. The Word Options dialog box looks fairly similar to the Excel Options dialog box.

3. **From the list on the left, choose the Advanced section.**

4. **Scroll down to the General heading, and then click File Locations.**

   This action opens the File Locations dialog box (Figure 16-7), which is the place where you tell Office applications where they should look for certain types of files. You can actually make three interesting changes here:

   - If you modify the Documents folder, Excel and Word start off in that folder the first time you choose to save or load a file after launching the application.

   - If you change the "User templates" folder, you can specify where Excel and Office look for user-created templates stored on your computer.

   - Finally, if you set the "Workgroup templates" folder, you can tell Word and Excel where to look for shared templates. This option is the one that interests you.

Figure 16-7:
If you haven't configured the workgroup templates folder before, you'll notice that it starts off blank. That's because Office has no idea where to put shared templates until you give it the right instructions.

Note: Remember, when Excel looks for templates, it searches both the user templates *and* the workgroup templates folders. Then it shows all the templates it found in the "My templates" list.

---

5. **Select the "Workgroup templates" entry in the list, and then click Modify.**

   A Modify Location dialog box appears. This dialog box looks more or less the same as the standard dialog boxes you use to open and save files.

6. **Browse to the folder you want to use, and then click OK.**

   Remember, you aren't limited to your local computer. Feel free to jump to a network drive, or even browse your network places to find a specific server.

7. **Click OK again to close the Word Options dialog box.**

   Now, the change has been made. Remember, in order for template sharing to work, *everyone* needs to perform this same set of steps to configure their copies of Excel to look in the same shared folder.

Template sharing is a simple idea that can become incredibly useful in a company environment. Instead of sending template files whizzing back and forth in emails, trying to keep a group of overworked employees in sync each time the template changes, you simply need to modify the templates in the shared location. That way, everybody always has the latest versions available, and there are no distribution headaches.

CHAPTER
17

Creating Basic Charts

As you become more skilled with Excel, you'll realize that entering numbers, organizing your layout, and formatting cells aren't the most important parts of spreadsheet creation. Instead, the real work lies in *analyzing* your data—in other words, figuring out a way to tell the story that lies *behind* your numbers. Excel's charting tools may be just what you need.

Charts depict data visually, so you can quickly spot overall trends. They're a fabulous way to help you find the meaning hidden in large amounts of data. You can create many different types of charts in Excel, including pie charts that present polling results, line charts that plot rising or declining assets over time, and three-dimensional area charts that show relationships between environmental conditions in a scientific experiment.

Excel's charting tools are enormously flexible: You can generate a simple chart with standard options in a couple of mouse clicks, or you can painstakingly customize every aspect of your chart's appearance (including colors, scale, titles, and even 3-D perspective). This chapter takes the first approach and explains how to generate straightforward charts, which you'll examine in detail. You'll also learn which chart types are out there. In the next chapter, you'll learn how to fine-tune your charts for maximum effect.

Note: All charts are *not* created equal. Depending on the chart type you use, the scale you choose, and the data you include, your chart may suggest different conclusions. The true chart artist knows how to craft a chart to draw out the most important information. As you become more skilled with charts, you'll acquire these instincts, too.

Charting 101

Excel provides a dizzying number of chart types, but they all share a few things. In this section, you'll learn about basic Excel charting concepts that apply to almost all types of charts; you'll also create a few basic charts. At the end of this chapter, you'll take a chart-by-chart tour of each and every one of Excel's many chart types.

To create a chart, Excel needs to translate your numbers into a graphical representation. The process of drawing numbers on a graph is called *plotting*. Before you plot your information on a chart, you should make sure your data's laid out properly. Here are some tips:

- Structure your data in a simple grid of rows and columns.
- Don't include blank cells between rows or columns.
- Include titles, if you'd like them to appear in your chart. You can use category titles for each column of data (placed in the first row, atop each column) and an overall chart title (placed just above the category-title row).

Tip: You can also label each row by placing titles in the far-left column, if it makes sense. If you're comparing the sales numbers for different products, list the name of each product in the first column on the left, with the sales figures in the following columns.

If you follow these guidelines, you can expect to create the sort of chart shown in Figure 17-1.

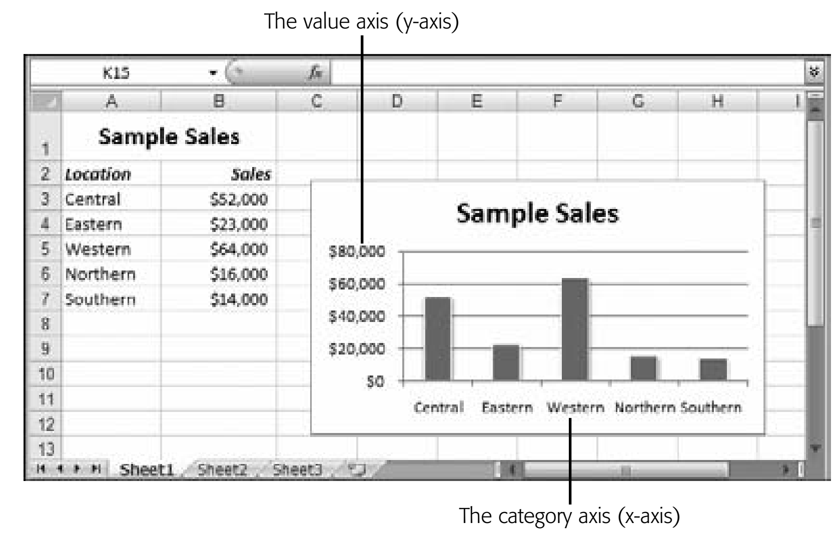

Figure 17-1: *This worksheet shows a table of data and a simple column chart based on Excel's standard chart settings. Nothing fancy, but it gets the job done.*

To create the chart shown in Figure 17-1, Excel performs a few straightforward steps (you'll learn the specifics of how to actually create this chart in the next section). First, it extracts the text for the chart title from cell A1. Next, it examines the range of data (from $14,000 to $64,000) and uses it to set the value—or Y-axis—scale. You'll notice that the scale starts at $0, and stretches up to $80,000 in order to give your data a little room to breathe. (You could configure these numbers manually, but Excel automatically makes commonsense guesses like these by looking at the data you're asking it to chart.) After setting the vertical scale, Excel adds the labels along the bottom axis (also known as the X-axis or category axis), and draws the columns of appropriate height.

Embedded and Standalone Charts

The chart in Figure 17-1 is an *embedded* chart. Embedded charts appear in a worksheet, in a floating box alongside your data. You can move the chart by dragging the box around your worksheet, although depending on where you put it, you may obscure some of your data.

Your other option is to create a *standalone* chart, which looks the same but occupies an entire worksheet. That means that your chart data and your chart are placed on separate worksheets. (See Chapter 4 if you need a refresher about how to create new worksheets, switch between them, and change their order.)

Usually, you'll use an embedded chart if you want to create printouts that combine both your worksheet data and one or more charts. On the other hand, if you want to print the charts separately, it's more convenient to use standalone charts. That way, you can print an entire workbook at once and have the charts and the data on separate pages.

Note: If you use embedded charts, you still have the option of printing just the chart, sized so that it fills a full sheet of paper. Simply select the chart, and then choose File→Print. If you create a standalone chart, you don't have a choice–Excel always prints your chart on a separate page.

Creating a Chart with the Ribbon

So how do you create a chart like the one shown in Figure 17-1? Easy—all it takes is a couple of clicks in the ribbon. Here's how it works:

1. **Select the range of cells that includes the data you want to chart, including the column and row headings and any chart title.**

   If you were using the data shown in Figure 17-1, you'd select cells A1 to B7.

   For speedier chart building, just position your cursor somewhere inside the data you want to chart. Excel then automatically selects the range of cells that it thinks you want. Of course, it never hurts to remove the possibility for error by explicitly selecting what you want to use before you get started.

Tip: And for even *easier* charting, start by creating an Excel table (Chapter 14) to hold the data you want to chart. Then, if you position yourself somewhere inside the table and create a new chart, Excel automatically selects all the data. It also automatically updates the chart if you add new rows or remove existing data.

2. **Head to the ribbon's Insert→Charts section. You'll see a separate button for each type of chart (including column charts, line charts, pie charts, and so on). Click the type you want.**

   When you choose a chart type, you get a drop-down list of subtypes (Figure 17-2).

   The different chart types are explained in more detail later in this chapter. For now, it's best to stick to some of the more easily understood choices, like Bar, Column, or Pie. Remember, the chart choices are just the starting point, as you'll still be able to configure a wide range of details that control things like the titles, colors, and overall organization of your chart.

3. **Click the subtype you want.**

   Excel inserts a new embedded chart alongside your data, using the standard options (which you can fine-tune later).

Note: If you don't want to make *any* choices, you can actually build a chart with one key press. Just highlight your data and press F11. This step creates a column chart on a new worksheet. Although you can't undo this operation, you can always delete the new chart worksheet and start over.

Figure 17-2: *Under each chart choice are yet more subtypes, which add to the fun. If you select the Column type (shown here), you'll get subtypes for two- and three-dimensional column charts, and variants that use cone and pyramid shapes. If you hover over one of these subtypes, a box appears with a brief description of the chart.*

POWER USERS CLINIC

Browsing Excel's Chart Gallery

Excel pros sometimes find that the ribbon approach is a bit awkward when you're trying to find a less commonly used chart type. In this situation, you may prefer to look at the full list of chart types and subtypes. To do so, head to the ribbon's Insert→Charts section, and then click the dialog launcher (the square-with-an-arrow icon in the bottom-right corner). You see the Insert Chart dialog box (Figure 17-3).

The Insert Chart dialog box doesn't just let you create charts. You can also designate the default chart type (the one that's used if you select some cells, and then press F11 to create a chart in a single bound). To designate a default chart, select it, and then click "Set as Default Chart". Lastly, the Insert Chart dialog box lets you use a custom chart template that you've previously prepared, as described on page 518.

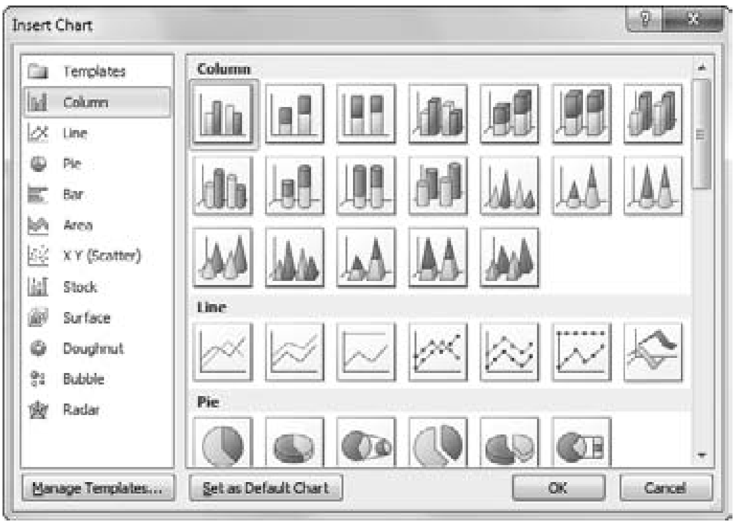

Figure 17-3: *The gallery on the Insert Chart dialog box's right side has a thumbnail of every chart subtype, grouped by type. You can scroll through them all, or you can choose a type from the list on the left to jump straight to a specific section. When you find what you want, click OK to create it.*

The Chart Tools Ribbon Tabs

When you select a chart, Excel adds three new tabs to the ribbon under the Chart Tools heading. These tabs let you control the details of your charts, and they are:

- **Design**. This tab lets you change the chart type and the linked data that the chart uses. It also lets you choose a chart style and layout, two ways to improve the appearance of a chart that you'll consider in Chapter 18.
- **Layout**. This tab lets you configure individual parts of the chart. You can add shapes, pictures, and text labels, and you can configure the chart's gridlines, axes, and background.

- **Format**. This tab lets you format individual chart elements, so you can transform ordinary items into eye candy. You can adjust the font, fill, and borders uses for chart titles and shapes, among other things.

In this chapter, you'll spend most of your time using the Chart Tools | Design tab. In the next chapter, you'll begin fine-tuning your charts, and you'll branch out to the other two tabs.

Basic Tasks with Charts

Unlike the orderly rows of numbers and labels that fill most worksheets, charts float *above* your data, locked inside special box-like containers. To take advantage of these chart boxes, you need to understand a little more about how they work.

Moving and Resizing a Chart

When you insert a chart into an existing worksheet, it becomes a floating object, hovering above your worksheet. Depending on where Excel puts it, it may temporarily obscure your data. The chart box doesn't damage your data in any way, but it can end up hiding your worksheet's numbers and text (both onscreen and in your printouts).

You have to learn to grab hold of these floating boxes and place them where you really want them. The process is pretty straightforward:

1. **Click the chart once to select it.**

   You'll notice that when you select a chart, Excel highlights the worksheet data the chart uses, as shown in Figure 17-4. At the same time, three new tabs appear in the ribbon, under the Chart Tools heading.

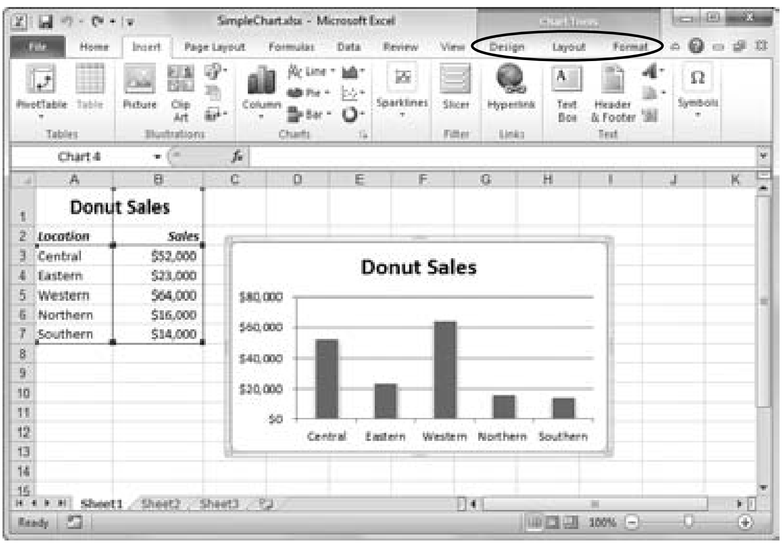

***Figure 17-4:** You'll know you've selected the chart when three new charting tabs appear in the ribbon under the Chart Tools heading.*

2. **Hover over the chart border until the mouse pointer changes to a four-way arrow.**

   Figure 17-5 shows what you're looking for.

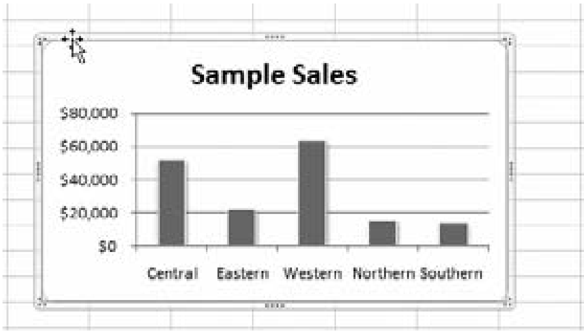

Figure 17-5:
The four-way arrow is a signal that you can click here to move the chart. If you move to one of the corners, you'll see an angled two-way arrow, which lets you resize the chart.

3. **Click and drag with your mouse to move or resize the chart.**

   Using the four-way arrow, you can drag the chart anywhere on your worksheet, releasing the mouse button when it's in the right spot.

   Using the two-way arrow, you can drag the border to make the chart larger or smaller. Once you make a chart box larger, you may also want to resize the individual components inside the chart to better use the available space. Page 507 tells you how to select chart elements and resize them.

Tip: To remove a chart in one fell swoop, just select it with the mouse, and then press Delete.

4. **When you're finished, click a cell anywhere in the worksheet to go back to your data.**

   At this point, life returns to normal, and the Chart Tools tabs disappear.

Tip: You can resize a chart in another, slightly more circuitous way. You can set the Height and Width boxes in the Chart Tools | Format→Size section of the ribbon. Although this isn't as quick as dragging the chart edge, it lets you set the size exactly, which is indispensable if you have several charts on the same worksheet and you need to make sure they're all the same size.

UNDER THE HOOD

How Excel Anchors Charts

Although charts appear to float above the worksheet, they're actually anchored to the cells underneath. Each corner of the chart is anchored to one cell (these anchor points change, of course, if you move the chart around). This fact becomes important if you decide to insert or delete rows or columns anywhere in your worksheet.

For example, consider the chart shown in Figure 17-1. Its top edge is bound to row 2, and its bottom edge is bound to row 12. Similarly, its left edge is bound to column C, and its right edge to column I. That means if you insert a new row above row 2, the whole chart shifts down one row. If you insert a column to the left of column C, the whole chart shifts one column to the right.

Even more interesting is what happens if you insert rows or columns in the area that the chart overlaps. For example, if you insert a new row between the current row 10 and row 11, the chart stretches, becoming one row taller. Similarly, if you delete column D, the chart compresses, becoming one column thinner.

If it bugs you, you can change this sizing behavior. First, select the chart, and then head to the ribbon's Chart Tools | Format→Size section. Click the dialog launcher (the square-with-an-arrow icon in the bottom-right corner). When the Size and Properties dialog box appears, choose Properties in the list on the left. You'll see three "Object positioning" options. The standard behavior is "Move and size with cells", but you can also create a chart that moves around the worksheet but never resizes itself ("Move but don't size with cells") and a chart that's completely fixed in size and position ("Don't move or size with cells").

Creating a Standalone Chart

Even without your input, Excel usually makes commonsense choices, so you can often build a chart without needing to tweak any of these options.

You have two options for placing charts in a workbook. You can create an embedded chart, which appears in an existing worksheet (usually next to the appropriate data), or you can create a standalone chart, which appears in a new worksheet of its own (Figure 17-6).

Ordinarily, when you pick a chart type from the ribbon, Excel creates an embedded chart. However, you can easily switch your chart over to a new worksheet if you're running out of room—just follow these steps:

1. **Right-click the chart, and then choose Move Chart (or, select the chart, and then choose Chart Tools | Design→Location→Move Chart).**

   The Move Chart dialog box appears (Figure 17-7).

Figure 17-6: *A standalone chart lives in a separate worksheet that doesn't have any other data and doesn't include the familiar grid of cells.*

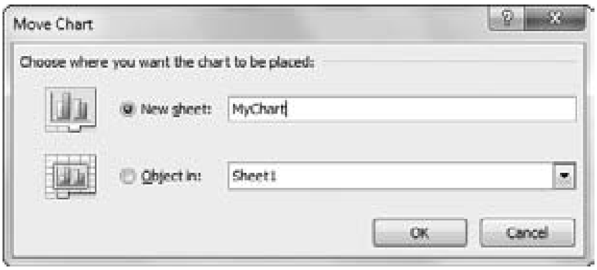

Figure 17-7: *Using the Move Chart dialog box, you can transfer the chart to a standalone worksheet (as shown here) or shuffle it over to another worksheet and keep it as an embedded worksheet. (If you want the latter option, it's just as easy to select the chart, and then use a cut-and-paste operation to move it to a new worksheet.)*

2. **Choose "New sheet", and then enter the name for the new worksheet.**

3. **Click OK.**

   Excel creates the new worksheet and places the chart on it. The new worksheet goes in front of the worksheet that contains the chart data. (You can always move the worksheet to a new position in your workbook by dragging the worksheet tab.)

Note: You can move or resize only embedded charts—the ones that appear in floating boxes inside other worksheets. If you've created a standalone chart, you can't move or resize it. Instead, it automatically shrinks or enlarges itself to match the Excel window's display area.

Editing and Adding to Chart Data

Every chart remains linked to the source data you used to create it. When you alter the data in your worksheet, Excel refreshes the chart with the new information automatically. As long as Excel is set to automatic recalculations (and it almost always is), there's no such thing as an out-of-date chart. (Page 386 has more about changing Excel's calculation mode.)

Note: Excel has no restriction on linking multiple charts to the same data. So you can create two types of charts (like a pie and a column chart) that show the same data. You can even create one chart that plots all the data and another chart that just uses a portion of the same information.

However, there's one tricky point. Any range you define for use in a chart is *static*, which means it doesn't grow as your data grows. That means that if you add a new row at the bottom of the range, it doesn't appear on the chart because it's outside of the chart range.

If you do want to insert additional data to a range of data used in a chart, you have several options:

- You can use the Home→Cells→Insert→Insert Sheet Rows command. If you do, Excel notices the change, and automatically expands the range to include the newly inserted row. However, this command works only if you add a row into the middle of your data. If you try to tack a row onto the end, Excel still ignores it, and you'll need to use the solution described in the next bullet point.
- After you insert new rows, you can modify the chart range to include the new data. This approach is the most common, and it's quite painless. First, select your chart by clicking the edge of the chart box, or a blank space in the chart. Excel highlights the linked worksheet data with a colored border. Click this colored border, and drag it until it includes all the new data you want. When you release the mouse button, Excel refreshes the chart with the new information.

Excel is smart enough to adjust your chart range in some situations. If you drag your chart data to a new place on your worksheet, Excel updates the chart to match automatically.

Changing the Chart Type

When you create a chart, you choose a specific chart type (page 464). However, in many situations you may want to try several chart types with the same data to see which visualization tells your story better. Excel makes this sort of experimentation easy. All you need to do is click your chart to select it, and then make a different choice from the ribbon's Insert→Charts section. You can use this technique to transform a column chart into a pie chart.

POWER USERS' CLINIC

Charting a Table

You can use the Excel table feature (discussed in Chapter 14) with charts. Tables and charts make a perfect match. Tables grow and shrink dynamically in size as you add or delete records. If a chart's bound to a table, the chart updates itself as you add new information or remove old data.

You've already learned how to build a new chart using an existing table. (Just move inside the table, and then make a selection from the ribbon's Insert→Charts section.) But even if you've already created the chart with an ordinary range of cells, you can still use a table—all you need to do is convert the linked range to a table.

In the sales report example shown in Figure 17-1, here's what you'd need to do:

1. Select the range of cells that contain all the data, not including the chart's title (cells A2 to B7).
2. Select Insert→Tables→Table.

Now, as you add new items to the table, Excel adds them to the chart immediately.

When you chart a table, you also gain the ability to use other features, like easier sorting and filtering. You can use sorting to determine the order that items appear within a chart (which is occasionally useful), and you can use filtering to hide rows and to chart only a portion of the data (which is often indispensable). If you apply a filter condition that shows only the three best performing regions, the chart updates itself so that it shows only this data. You'll find this technique particularly handy when you're creating charts that use multiple series, as described later in this chapter.

For more information about filtering and the ever-impressive table feature, refer to Chapter 14.

You can also choose Chart Tools | Design→Type→Change Chart Type to make a choice from the Change Chart Type dialog box, which looks just like the Insert Chart dialog box shown in Figure 17-3.

Printing Charts

How you print a chart depends on the type of chart you've created. You can print embedded charts either with worksheet data or on their own. Standalone charts, which occupy separate worksheets, always print on separate pages.

Embedded charts

You can print embedded charts in two ways. The first approach is to print your worksheet exactly as it appears in the Excel window, with its mix of data and floating charts. In this case, you need to take special care to make sure your charts aren't split over a page break or positioned over some data you want to read in the printout. You can check for both issues using Page Layout view (choose View→Workbook Views→Page Layout View) or the smaller print preview that's shown in backstage view when you're about to print your worksheet (make sure the chart isn't selected, and then choose File→Print).

You can also print an embedded chart on a separate page, which is surprisingly easy. Just click the chart to select it, and then choose File→Print. When you do so, Excel's standard choice is to print your chart using landscape orientation, so that the long edge of the page is along the bottom, and the chart's wider than it is tall. Landscape is usually the best way to align a chart, especially if it holds a large amount of data, so Excel automatically uses landscape orientation no matter what page orientation you've configured for your worksheet. Of course, you can change this as you would with any other printout; just choose Portrait Orientation in the list of print settings before you click the big Print button.

Excel also includes two print options that are specific to charts. To see these options, click the Page Setup link at the bottom of the list of print settings. When the Page Setup dialog box appears, choose the Chart tab. You'll see an option to print a chart using lower print quality ("Draft quality"), and in black and white instead of color ("Print in black and white").

Standalone charts

If you're using a standalone chart, your chart always prints out on a separate page, sized to fit the whole page. To print out just the chart page alone (rather than the whole workbook), switch to the chart's worksheet, and then choose File→Print. Excel automatically sets all chart worksheets to Landscape orientation, which orients the page so that the long edge runs horizontally across the bottom. If this layout isn't what you want, change the orientation setting to Portrait Orientation before you print.

If you want to print the entire workbook, choose File→Print from any worksheet. Then, change the first print setting from Print Active Sheets to Print Entire Workbook.

Practical Charting

Figure 17-1 showed you how to chart a list that contains two columns you want to graph—one with text labels and one with numeric data. But, in real life, you'll probably need to deal with many types of data that occupy many different configurations on your worksheet.

Consider all the possible variations on the simple sales chart shown in Figure 17-1. You may need to compare the sales figures but, rather than showing region-to-region comparisons, you want to show how well (or poorly) each of your firm's products sold. Or perhaps you want to chart the quarterly performance of different stores over a five-year period, or determine the relationship between sales and profitability. All these charts require a slightly different arrangement of data. In the next section, you'll get a quick introduction to all these possibilities, using just the simple column chart and line chart.

Charts with Multiple Series of Numbers

A *series* is the sequence of numbers that you plot on a graph. In the simple chart example (Figure 17-1), there's one series of numbers, which represents the sales figures for a company's different regions. Of course, a real chart usually adds extra layers of detail. You may want to compare the sales figures from several different years. In this case, you'd add a separate column to your worksheet data for each year. Then you'd add each column to your chart as a separate series.

It doesn't take any extra expertise to create a chart that uses multiple series—you just select the right range of cells, and then pick a chart option from the ribbon, just as you would for a chart that has a single series. Different types of charts handle multiple series in different ways. The clustered column chart creates a separate bar for each value in a row, as shown in Figure 17-8. A line chart, on the other hand, shows a separate line for each series (as shown in the next section). For more possibilities, take a look at the "Chart Types" section that starts on page 481.

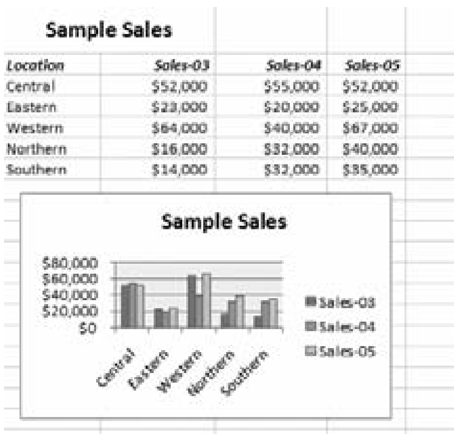

Figure 17-8:
This chart has three series of sales figures (one for each year) and five sets of columns (one for each region). Each of the five sets has three bars, one for each data series. The regions are labeled on the category axis, but you'll need to consult the legend to determine which year each column represents.

Tip: You can add multiple series to an existing chart without starting over from scratch. First, select the chart to highlight the linked data. Then, click the rightmost edge, and drag it to the right to expand the range so that it includes the new columns (which, of course, you've already added to your worksheet).

Controlling the Data Excel Plots on the X-Axis

Excel's charting tool has a dirty little secret. You may not realize it right away, but sooner or later, whether it's your first chart or your 40th, you'll stumble onto the fact that Excel makes a fairly important decision for you about what data shows up in your chart's X-axis. Unfortunately, this decision may not be what you want. Fortunately, you can change it.

UP TO SPEED

Data in Different Scales

Remember when your mother told you not to compare apples and oranges? The same rule applies with charts. When you add multiple series, each series should use the same *scale*. In other words, the points for each series should be plotted (placed on the chart) using the same measurement system.

The worksheet in Figure 17-8 works perfectly well because the different series of sales figures all use the same unit—dollars. But if one series recorded sales totals in dollars and another recorded them in Euros (or even worse, recorded totally different data like the number of units sold), the chart would be inconsistent.

Excel doesn't complain if your series use different scales—in fact, it has no way of noticing that anything's amiss. And if you don't notice either, you'll create a misleading chart. Your chart may imply a comparison that isn't accurate or, if the scale is radically different, the chart can get so stretched that it starts to lose detail. If you have sales figures from $50,000 to $100,000 and units sold from 1 to 100, the scale stretches from 1 to 100,000, and the differences in sales totals or units sold are too small to show up at all.

What's the solution? Don't mix different scales. Ideally, convert values to the same scale (in this case, use the currency exchange rate to turn Euros into U.S. dollars before you create the chart). Or just create two charts, one for each data series. But if you really want to compare the changes in different types of data across the same categories, there's a way. Page 534 shows you how to build combination charts that fuse together two incompatible sets of data in a logical way.

But what causes the situation in the first place? Excel creates your charts according to the way the data's organized in your worksheet. A simple example shows you the effect.

The worksheet in Figure 17-9 looks at sales based on two factors: the year when the sales were recorded, and the region where the sales were made. In technical charting terms, the regions form the *category axis*, while the sales figures form the *value axis*. In other words, Excel creates a separate series for each year. But it makes just as much sense to organize the table in a different way, by making the year the category axis and creating a separate series for each region! Figure 17-9 contrasts these two different ways of looking at the same data, and shows how they affect the way Excel groups your data in a column chart.

The column chart example is fairly innocent. Although you may prefer one way of looking at the data over the other, they're relatively similar. However, most Excel charts aren't as forgiving. The line chart's a classic example.

In a line chart, each line represents a different series. If you list the sales years on the category axis (as shown on the left side of Figure 17-10), you end up with a separate line for each region that shows how the region has performed over time. But if you invert the table (shown on the right side), you end up with a chart that doesn't make much sense at all: a series of lines that connect different regions in each year. Figure 17-10 shows the problem.

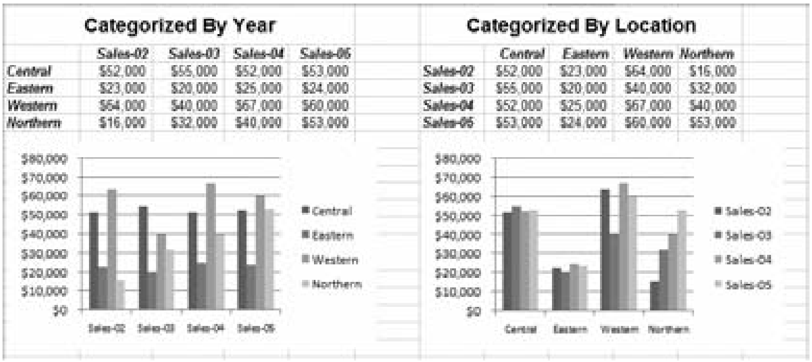

Figure 17-9: *This worksheet shows the same data charted in two different ways. In the first table (left), the category axis lists the sales years, which are used to group the regions. In the second table (right), the category axis lists the regions, which are used to group the years.*

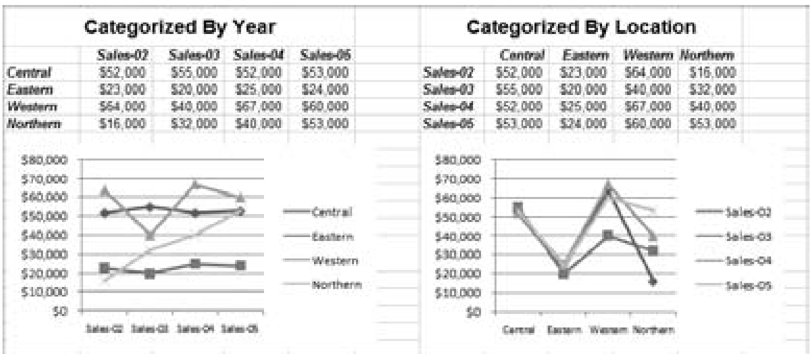

Figure 17-10: *The chart on the left is pretty straightforward. The chart on the right shows a line for each year, which makes sense if you concentrate on what's being depicted, but mostly shows how people can use computers to complicate things.*

Clearly, when you create a line chart, you need to make sure the chart ends up using the data in a way that makes the most sense. So, how does Excel decide how to plot the data? Essentially, Excel makes a best guess about your data. If you have more rows than columns, Excel assumes that the first column represents the category axis. If you have more columns than rows (or if you have the same number of rows and columns), Excel assumes that the first *row* represents the category axis, as in Figure 17-10.

Fortunately, you have the power to override Excel's choice if you need to. Just select your chart, and then choose Chart Tools | Design→Data→Switch Row/Column. If you try this action on the charts in Figure 17-10, you reverse the results. Thus, the chart on the left would group the data into yearly series, and the chart on the right would group the data into regional series. To return them to normal, you can select each chart, and then click Switch Row/Column again.

UP TO SPEED

The Difference Between a Column and a Line

With simple column charts, life is easy. It doesn't matter too much what data you choose to use for your category axis because your choice simply changes the way data's grouped. Other chart types that follow the same principle include pie charts (which only allow one series), bar charts (like column charts, but oriented horizontally instead of vertically), and donut charts (where each series is a separate ring).

The same isn't true for line charts and most other types of Excel charts. The category axis you use for a line chart is important because the values in each series are connected (in this case, with a line). This line suggests some sort of "movement" or transition as values move from one category to another. That means it makes sense to use a line to connect different dates in a region (showing how sales have changed over time), but it probably doesn't make sense to use a line to connect different regions for each date. Technically, this latter scenario (shown on the right side of Figure 17-10) should show how yearly sales vary as you move from region to region, but it's just too counterintuitive for anyone to interpret it properly.

As a general rule of thumb, use time or date values for the category axis. You should do this *especially* for chart types like line and area, which usually show how things change over time.

Data That Uses a Date or Time Scale

As the previous example shows, using time or date values for the category axis makes a lot of sense for charting progress or spotting long-term trends. However, the example does cheat a little. Even though any sentient human knows that the labels Sales-03, Sales-04, and Sales-05 represent consecutive years, Excel is oblivious to what these labels actually mean. You could chart a bunch of years that are far from sequential (like Sales02, Sales04, and Sales08) and Excel would obediently (and misleadingly) place each value on the category axis, spaced out evenly.

This snafu doesn't present a problem in the previous example, but it's an issue if you need to chart years that aren't spread out evenly. Fortunately, Excel offers an easy solution. Instead of entering text labels, you can enter actual dates or times. Because Excel stores dates and times as numbers, it can scale the chart accordingly (this process is sometimes called *category axis scaling*). Best of all, Excel automatically notices when you're using real dates, and kicks into action, making the appropriate adjustments, as shown in Figure 17-11.

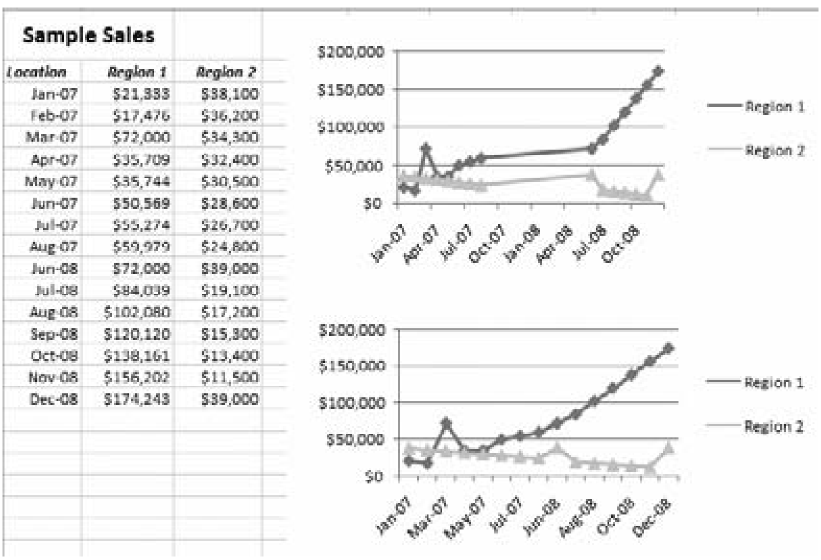

Figure 17-11: *The top chart uses category axis scaling to properly space out dates, even when values are missing. The bottom chart doesn't.*

What's happening in Figure 17-11 is worth examining in a bit of detail. The pictured worksheet shows two charts that show the exact same data: a series of monthly sales figures from two regions (covering the time period between January 2010 and December 2011). The diamonds and triangles on the line charts indicate the data points for which sales data is available. The twist is that a big chunk of data (the months between August 2010 and June 2011) is missing. To make sure Excel handles this omission correctly, you have to enter real date values (rather than text labels) for the category axis. If you take that step, the chart Excel creates automatically uses a continuous timescale, as shown in the top chart. (As you can see by looking at the data points, no values fall in the middle of the series.)

On the other hand, if you enter the labels as text (as was done when creating the bottom chart), you'll see an incorrect result: The data from August 2010 and June 2011 are placed close together—even though they record months that are almost a year apart.

Optionally, you have the ability to tell Excel to disregard any values you've used in your column or row labels, thereby spacing the dates out evenly, as though they're ordinary text labels. That's how the incorrect chart in Figure 17-11 was created. (Why you'd want to do it is another question, but someone, somewhere, is probably in desperate need of this feature.) To change how Excel scales the category axis, select the chart, and then choose Chart Tools | Layout→Axes→Axes→Primary Horizontal Axis→More Primary Horizontal Axis Options to show the Format Axis dialog box. Next, change the Axis Type setting, as shown in Figure 17-12.

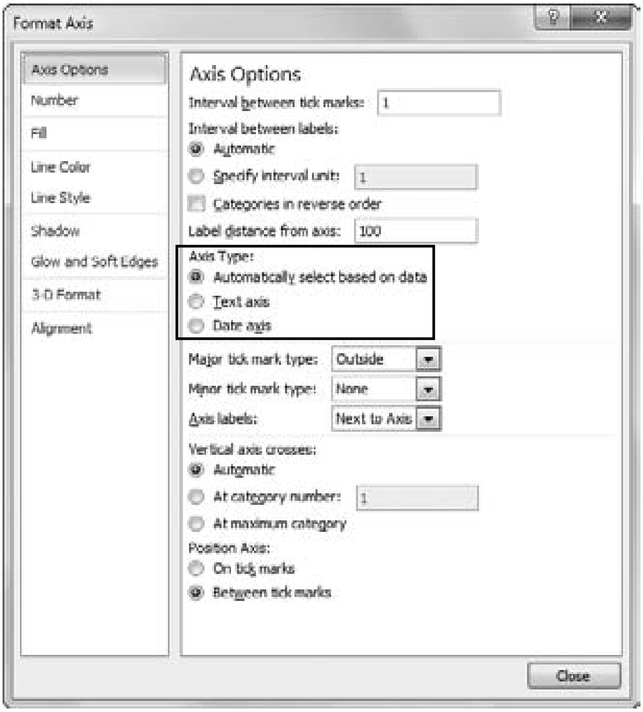

Figure 17-12:
Under the Axis Type heading, pick one of the following: "Text axis" (treat the category values as labels), "Date axis" (treat the category values as date vales), or "Automatically select based on the data" (let Excel decide based on what it thinks is best). Chapter 18 has more information about the other options in this dialog box.

Category axis scaling works with more than just dates. You can scale any category axis values, as long as they're numeric, which is particularly useful if you're trying to determine the relationship between two different values. If you wanted to determine the relationship between students' IQs and their test scores, you could use the numeric IQ for the category axis, and use the test scores for the value axis. If you want to create a chart like this that compares two sets of numbers, you must use a *scatter chart* instead of a line chart. Scatter charts look similar to line charts. Page 488 covers scatter charts in more detail.

Noncontiguous Chart Ranges

So far, all the chart examples have assumed the data you want to chart is placed in a single, tightly packed table. But what if your information is actually scattered across your worksheet? This scenario may seem unlikely, but it actually happens quite often when you need to chart only *part* of the data in a table. Say you want to create a chart using two or three columns, and these columns aren't next to each other. In this case, you need to take a few extra steps when you create your chart.

Imagine you have a table that records the monthly sales of 10 different regional offices. However, you want to create a chart that compares only two of these offices. Your chart will use the category information in column A (which contains the month in which the sales were recorded), along with the values in column C and column D (which contain the total amount of sales for the two regions in which you're interested).

The easiest way to create this chart is to start by selecting the noncontiguous range that contains your data. Chapter 3 describes this technique in detail (page 99). Here's what you need to do:

1. **First, use the mouse to select the data in column A.**

   Excel surrounds the data with a marquee. Don't click anywhere else yet.

2. **Then, hold down the Ctrl key while you drag to select the data in columns C and D.**

   Because you're holding down the Ctrl key, column A remains selected (see Figure 17-13).

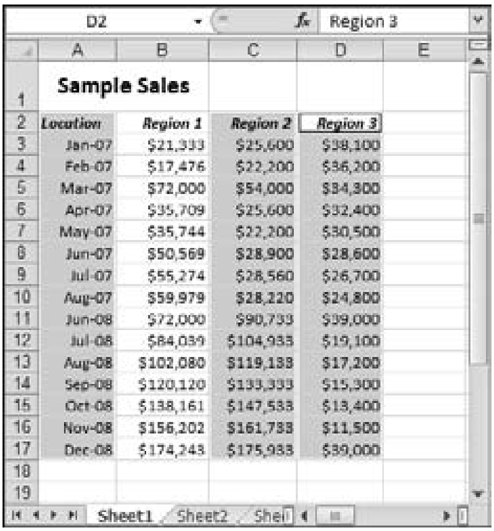

	A	B	C	D	E
1	Sample Sales				
2	Location	Region 1	Region 2	Region 3	
3	Jan-07	$21,333	$25,600	$38,100	
4	Feb-07	$17,476	$22,200	$36,200	
5	Mar-07	$72,000	$54,000	$34,300	
6	Apr-07	$35,709	$25,600	$32,400	
7	May-07	$35,744	$22,200	$30,500	
8	Jun-07	$50,569	$28,900	$28,600	
9	Jul-07	$55,274	$28,560	$26,700	
10	Aug-07	$59,979	$28,220	$24,800	
11	Jun-08	$72,000	$90,733	$39,000	
12	Jul-08	$84,039	$104,933	$19,100	
13	Aug-08	$102,080	$119,133	$17,200	
14	Sep-08	$120,120	$133,333	$15,300	
15	Oct-08	$138,161	$147,533	$13,400	
16	Nov-08	$156,202	$161,733	$11,500	
17	Dec-08	$174,243	$175,933	$39,000	
18					
19					

Figure 17-13:
This worksheet shows a noncontiguous selection that ignores the numbers from region 1. When you create the chart, Excel includes only two series in the chart: one for region 2, and one for region 3.

3. **Now choose Insert→Charts, and then pick the appropriate chart type.**

   Excel creates the chart as usual, but uses only the data you selected in steps 1 and 2, leaving out all other columns.

This approach works most of the time. However, if you have trouble, or if the columns you want to select are spaced *really* far apart, then you can explicitly configure the range of cells for any chart. To do so, follow these steps:

1. **Create a chart normally, by selecting part of the data, and then, from the Insert→Chart section of the ribbon, choosing a chart type.**

2. **Select the chart, and then choose Chart Tools | Design→Data→Select Data.**

   The Select Data Source dialog box appears (Figure 17-14).

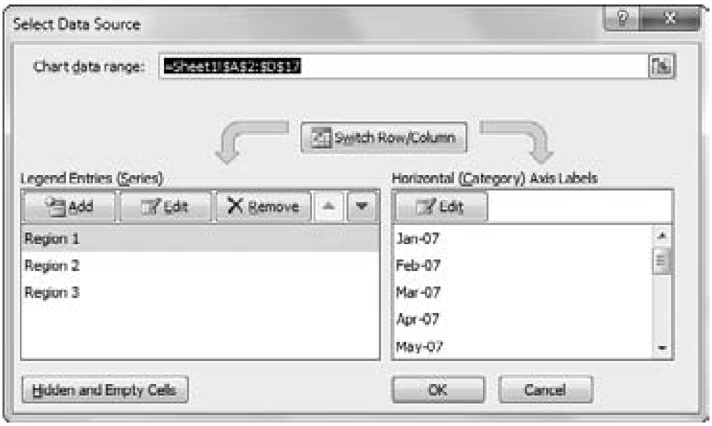

Figure 17-14: *This dialog box shows a handy secret about Excel charting. Excel not only records the whole range of cells that contain the chart data (as shown in the "Chart data range" text box), it also lets you see how it breaks that data up into a category axis and one or more series (as shown in the Legend Entries (Series) list).*

3. **Remove any data series you don't want and add any new data series you do want.**

   To remove a series, select it in the Legend Entries (Series) list, and then click Remove.

To add a new series, click Add, and then specify the appropriate cell references for the series name and the series values.

You can also click Switch Row/Column to change the data Excel uses as the category axis (page 474) and you can adjust some more advanced settings, like the way Excel deals with blank values, and the order in which it plots series (as explained in the following sections).

Changing the Order of Your Data Series

If your table has more than one series, Excel charts it in the order it appears on your worksheet (from left to right if your series are arranged in columns, or from top to bottom if they're arranged in rows). In a basic line chart, it doesn't matter which series Excel charts first—the end result is still the same. But in some charts, it *does* make a difference. One example is a stacked chart (Figure 17-16), in which Excel plots each new series on top of the previous one. Another example is a 3-D chart, where Excel plots each data series behind the previous one.

You can easily change your data series' order. Select your chart, and then choose Chart Tools | Design→Data→Select Data. Now select one of the series in the Legend Entries (Series) list, and then click the up or down arrow buttons to move it. Excel plots the series from top to bottom.

Changing the Way Excel Plots Blank Values

When Excel creates a chart, its standard operating procedure is to *ignore* all empty cells. The value of 0 doesn't count as an empty cell and neither does text (Excel plots any cells that contains text as a 0).

So what's the difference between an ignored cell and a cell that contains the number 0? In some types of charts, there's no difference. In a bar or pie chart, the result is the same—you don't see a bar or a pie slice for that data. However, in some charts, there *is* a difference. In a line chart a 0 value is plotted on the chart, but an empty cell causes a break in the line. In other words, the line stops just before the missing data, and then starts again at the next point. This broken line indicates missing information.

If you don't like this behavior (perhaps because your empty cells really do represent 0 values), you can change it. Select your chart, and then choose Chart Tools | Design→Data→Select Data to get to the Select Data Source dialog box. Then, click the Hidden and Empty Cells button, which pops open a dialog box with three choices:

- **Gaps**. Excel leaves a gap where the information should be. In a line chart, this breaks the line (making it segmented). This option is the standard choice.
- **Zero**. Excel treats all blank cells as though they contain the number 0.
- **Span with line**. Excel treats all blank cells as missing information and tries to guess what should go in between. If a line chart goes from 10 to 20 with a blank cell in between, Excel interpolates the data point 15 and plots it.

You can also switch on or off the "Show data in hidden rows and columns" setting to tell Excel whether it should include hidden cells when creating a chart. This setting determines how Excel deals with data when you use filtering in a table, or when you explicitly hide rows or columns using the Home→Cells→Format→Hide & Unhide menu. Ordinarily, Excel treats these missing values just like blank values, and ignores them.

Chart Types

Although there's a lot to be said for simple column charts—they can illuminate trends in almost any spreadsheet—there's nothing quite as impressive as successfully pulling off the exotic bubble chart. This section covers the wide range of charts that Excel offers. If you can use these specialized chart types when they make sense, you can convey more information and make your point more effectively.

Note: The following sections explain all of the Excel chart types. To experiment on your own, try out the downloadable examples, which you can find on the Missing CD page at *www.missingmanuals.com/cds*. The examples include worksheets that show most chart types. Remember, to change a chart from one type to another, just select it, and then make a new choice from the ribbon's Insert→Charts section, or use the Chart Tools | Design→Type→Change Chart Type command.

Column

By now, column charts probably seem like old hat. But column charts actually come in several different variations (technically known as *subtypes*). The main difference between the basic column chart and these subtypes is how they deal with data tables that have multiple series. The quickest way to understand the difference is to look at Figure 17-15, which shows a sample table of data, and Figure 17-16, which charts it using several different types of column charts.

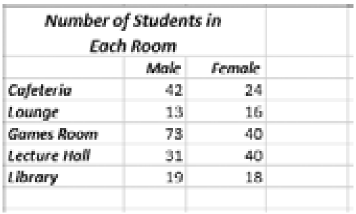

Number of Students in Each Room		
	Male	Female
Cafeteria	42	24
Lounge	13	16
Games Room	73	40
Lecture Hall	31	40
Library	19	18

Figure 17-15:
This simple table of data records the number of female and male students in several rooms at a university. The category is the room name, and there are two data series: the numbers of male students, and the numbers of female students. This data is perfect for a column chart, but different subtypes emphasize different aspects of the data, as you can see in Figure 17-16.

Note: In order to learn about a chart subtype, you need to know its name. The name appears when you hover over the subtype thumbnail, either in the Insert→Charts list (Figure 17-2) or the Insert Chart dialog box (Figure 17-3).

Here's a quick summary of your column chart choices:

- **Clustered Column**. In a clustered column, each value's shown in its own separate column. To form a cluster, the columns are grouped together according to category.
- **Stacked Column**. In a stacked column, each category has only one column. To create this column, Excel adds together the values from every series for each category. However, the column is subdivided (and color-coded), so you can see the contribution each series makes.
- **100% Stacked Column**. The 100% stacked column is like a stacked column in that it uses a single bar for each category, and subdivides that bar to show the proportion from each series. The difference is that a stacked column always stretches to fill the full height of the chart. That means stacked columns are designed to focus exclusively on the percentage distribution of results, not the total numbers.

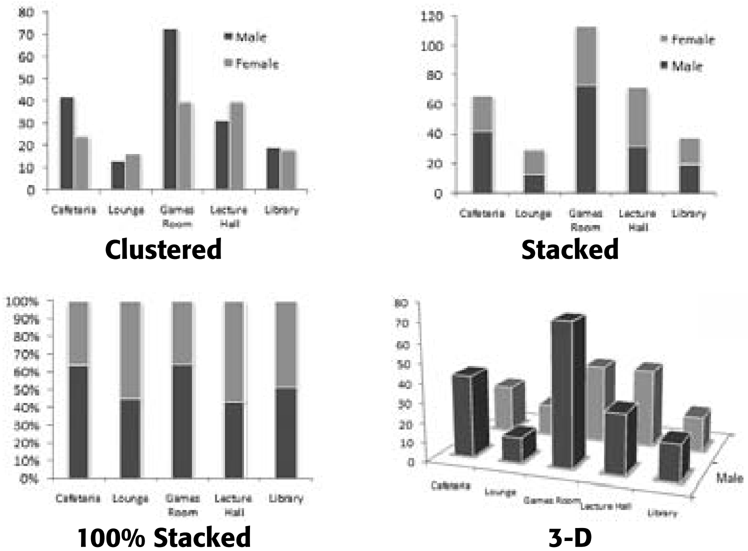

Figure 17-16: *The Clustered Column makes it easy to compare the gender of students in each room, but makes it somewhat more difficult to compare different rooms. The Stacked Column is an elegant way to compress the data, and it lets you compare the total number of students in each room without losing the gender information. The 100% Stacked Column makes each column the same height, so it's useless for comparing total student numbers but perfect for comparing how the gender breakup varies depending on the room. (Notice the scale also changes to reflect that you're comparing percentage values.) Finally, the 3-D chart shows you all the data at once by placing the male student counts in front of the female student counts.*

- **3-D Clustered Column, Stacked Column in 3-D, and 100% Stacked Column in 3-D**. Excel's got a 3-D version for each of the three basic types of column charts, including clustered, stacked, and 100% stacked. The only difference between the 3-D versions and the plain-vanilla column charts is that the 3-D charts are drawn with a three-dimensional special effect, that's either cool or distracting, depending on your perspective.
- **3-D Column**. While all the other 3-D column charts simply use a 3-D effect for added pizzazz, this *true* 3-D column chart actually uses the third dimension by placing each new series *behind* the previous series. That means, if you have three series, you end up with three layers in your chart. Assuming the chart is tilted just right, you can see all these layers at once, although it's possible that some bars may become obscured, particularly if you have several series.

Along with the familiar column and three-dimensional column charts, Excel also provides a few more exotic versions that use cylinders, cones, and pyramids instead of ordinary rectangles. Other than their different shapes, these chart types work just like regular column charts. As with column and bar charts, you can specify how cylinder, cone, and pyramid charts should deal with multiple series. Your options include clustering, stacking, 100% stacking, and layering (true 3-D). See Figure 17-17 for an example.

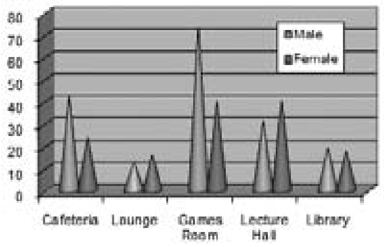

Figure 17-17:
Though a cone chart looks a little different, it's really just a column chart in disguise.

Bar

The venerable bar chart is the oldest form of data presentation. Invented sometime in the 1700s, it predates the column and pie chart. Bar charts look and behave almost exactly the same as column charts—the only difference being that their bars stretch horizontally from left to right, unlike columns, which rise from bottom to top.

Excel provides almost the same set of subtypes for bar charts as it does for column charts. The only difference is that there's no true three-dimensional (or layered) bar chart, although there are clustered, stacked, and 100% stacked bar charts with a three-dimensional effect. Some bar charts also use cylinder, cone, and pyramid shapes.

Tip: Many people use bar charts because they leave more room for category labels. If you have too many columns in a column chart, Excel has a hard time fitting all the column labels into the available space.

Line

People almost always use line charts to show changes over time. Line charts emphasize trends by connecting each point in a series with a line. The category axis represents a time scale or a set of regularly spaced labels.

Tip: If you need to draw smooth trendlines, you don't want to use a line chart. That's because a line chart connects every point exactly, leading to jagged zigzagging lines. Instead, use a scatter chart (page 488) without a line, and add one or more trendlines afterward, as explained in the next chapter (page 523).

Excel provides several subtypes for line charts:

- **Line**. The classic line chart, which draws a line connecting all the points in the series. The individual points aren't highlighted.

- **Stacked Line**. In a stacked line chart, Excel displays the first series just as it would in the normal line chart, but the second line consists of the values of the first and second series added together. If you have a third series, it displays the total values of the first three series, and so on. People sometimes use stacked line charts to track things like a company's cumulative sales (across several different departments or product lines), as Figure 17-18, bottom, shows. (Stacked area charts are another alternative, as shown in Figure 17-20.) Stacked line charts aren't as common as stacked bar and column charts.

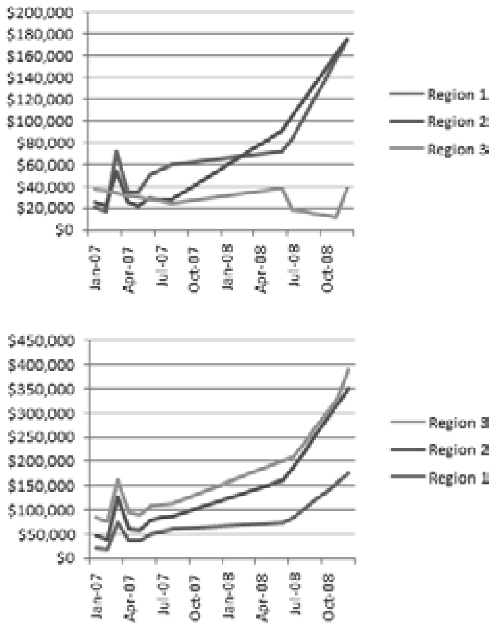

Figure 17-18:
Here are two different line chart variations—both of which show the same information, although you'd never be able to tell that from looking at them quickly.

Top: This chart is a regular line chart that compares the sales for three different regions over time.

Bottom: This chart is a stacked line chart, which plots each subsequent line by adding the numbers from the earlier lines. That makes the stacked line chart a great vehicle for showing cumulative totals. For example, sales in Region 3 for April of 2010 appear to top $150,000. That's because the Region 3 line is stacked. It shows a total made up from three components—$72,000 (Region 1), $54,000 (Region 2), and $34,300 (Region 3). In this example, the stacked line chart clearly shows that sales spiked early on, but have risen over-all, which isn't clear in the top chart. However, the stacked line chart also obscures the differences between the regions. You'd never guess that Region 3 is the underperforming re-gion because this chart reflects the total of all three regions.

Note: Lines can never cross in a stacked line chart, because Excel adds each series to the one (or ones) before it. You can change which line is stacked at the top by changing the order of the series. To do this, either rearrange your table of data on the worksheet (Excel places the rightmost column on top) or refer to page 480, which describes how you can change the order of your series manually.

- **100% Stacked Line**. A 100% stacked line chart works the same as a stacked line chart in that it adds the value of each series to the values of all the preceding series. The difference is that the last series always becomes a straight line across the top, and the other lines are scaled accordingly so that they show percent-ages. The 100% stacked line chart is rarely useful, but if you do use it, you'll probably want to put totals in the last series.

- **Line with Markers, Stacked Line with Markers, and 100% Stacked Line with Markers**. These subtypes are the same as the three previous line chart subtypes, except they add markers (squares, triangles, and so on) to highlight each data point in the series.
- **3-D Line**. This option draws ordinary lines without markers but adds a little thickness to each line with a 3-D effect.

Pie

Pie charts show the breakdown of a series proportionally, using "slices" of a circle. Pie charts are one of the simplest types of charts, and one of the most recognizable.

Here are the pie chart subtypes you can choose from:

- **Pie**. The basic pie chart everyone knows and loves, which shows the breakup of a single series of data.
- **Exploded Pie**. The name sounds like a Vaudeville gag, but the exploded pie chart simply separates each piece of a pie with a small amount of white space. Usually, Excel charting mavens prefer to explode just a single slice of a pie for emphasis. This technique uses the ordinary pie subtype, as explained in the next chapter.
- **Pie of Pie**. With this subtype, you can break out one slice of a pie into its own, smaller pie (which is itself broken down into slices). This chart is great for emphasizing specific data; it's shown in the next chapter.
- **Bar of Pie**. The bar of pie subtype is almost the same as the pie of pie subtype. The only difference is that the breakup of the combined slice is shown in a separate stacked bar, instead of a separate pie.
- **Pie in 3-D** and **Exploded Pie in 3-D**. This option is the pie and exploded pie types in three dimensions, tilted slightly away from the viewer for a more dramatic appearance. The differences are purely cosmetic.

Note: Pie charts can show only one series of data. If you create a pie chart for a table that has multiple data series, you'll see just the information from the first series. The only solution is to create separate pie charts for each series (or try a more advanced chart type, like a donut, described on page 491).

Area

An area chart is very similar to a line chart. The difference is that the space between the line and the bottom (category) axis is completely filled in. Because of this difference, the area chart tends to emphasize the sheer magnitude of values rather than their change over time (see Figure 17-19).

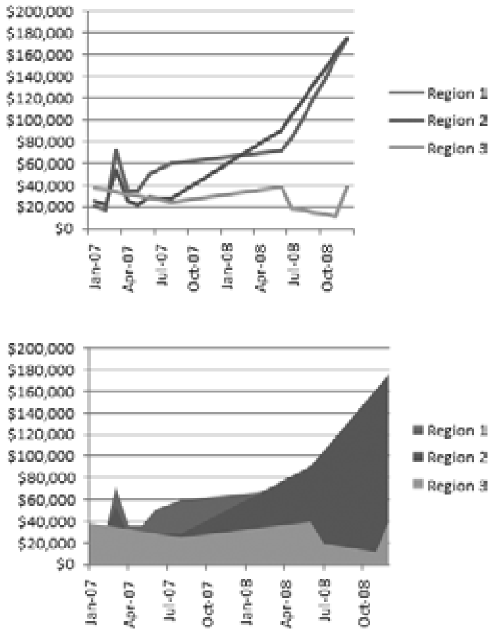

Figure 17-19:
This example compares a traditional line chart (top) against the area chart (bottom). As you can see, the area chart makes a more dramatic point about the rising sales in Region 2. However, it also obscures the results in Region 1.

Area charts exist in all the same flavors as line charts, including stacked and 100% stacked. You can also use subtypes that have a 3-D effect, or you can create a true 3-D chart that layers the series behind one another.

Stacked area charts make a lot of sense. In fact, they're easier to interpret than stacked line charts because you can easily get a feeling for how much contribution each series makes to the total by judging the thickness of the area. If you're not convinced, compare the stacked charts in Figure 17-18 (bottom) and Figure 17-20. In the area chart, it's much clearer that Region 3 is making a fairly trivial contribution to the overall total.

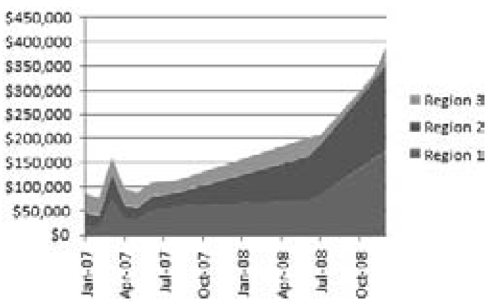

Figure 17-20:
You can create an area chart that doesn't obscure any data, but it needs to be a stacked (as shown here) or 3-D area chart. The stacked area chart shows the combined total of all regions, but it still lets you pick out the most important series. For example, it's clear that Region 3 (the narrow sliver on top of the stack) contributes relatively little to the total, while Region 1 and Region 2 are more important.

XY (Scatter)

XY Scatter charts show the relationship between two different sets of numbers. Scatter charts are common in scientific, medical, and statistical spreadsheets. They're particularly useful when you don't want to connect every dot with a straight line. Instead, scatter charts let you use a smooth "best fit" trendline, or omit the line altogether. If you plot multiple series, the chart uses a different symbol (like squares, triangles, and circles) for each series, ensuring that you can tell the difference between the points.

Why would you want to plot data points without drawing a line? For one thing, you may need to draw conclusions from an inexact or incomplete set of scientific or statistical data. Scientific types may use a scatter chart to determine the relationship between a person's age and his reflex reaction time. However, no matter how disciplined the experimenters, they can't test every different age. In addition, their data will include natural variations from the overall trend. (In other words, if the trend is for older people to have gradually slowing reactions, you're still likely to run across a few exceptionally speedy older folks.) In this case, the best approach is to include no line, or use a smooth "best fit" line that indicates the overall trend, as shown in Figure 17-21.

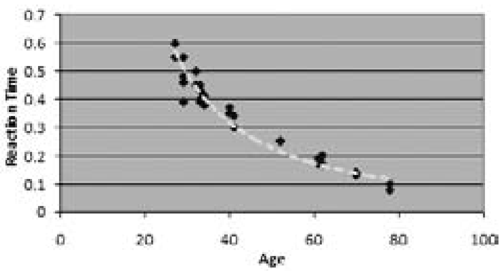

Figure 17-21:
This XY Scatter chart shows the relationship between a person's age and his reflex reaction time.

Excel offers several scatter chart subtypes, including:

- **Scatter with Only Markers**. This scatter chart uses data markers to show where each value falls. It adds no lines.
- **Scatter with Smooth Lines and Markers**. This scatter chart adds a smooth line that connects all the data points. However, the points are connected in the order they appear in the chart, which isn't necessarily the correct order. You're better off adding a trendline to your chart, as explained in the next chapter.

- **Scatter with Straight Lines and Markers**. This subtype is similar to the scatter chart with smoothed lines, except it draws lines straight from one point to the next. A line chart works like this, and this subtype makes sense only if you have your values in a set order (from lowest to largest or from the earliest date to the latest).
- **Scatter with Smooth Lines and Scatter with Straight Lines**. These subtypes are identical to the scatter with smooth lines and markers and the scatter with straight lines and markers. The only difference is they don't show data markers for each point. Instead, all you see is the line.

Stock

A stock chart displays specialized charts for stocks. Usually, these charts show how a stock value changes over a series of days. The twist is that the chart can display information about the daytime high and the daytime low of the stock, along with its opening and closing value. Excel uses all this information to draw a vertical bar from the stock's low point to its high point on a given day. If you're really ambitious, you can even add volume information (which records the number of shares traded on a given day).

Stock charts are more rigid than most other chart types. In order to use a stock chart, you need to create a column of numbers for each required value. The type of columns you need and their order depends on the stock chart subtype that you select. Here are your choices:

- High-Low-Close
- Open-High-Low-Close
- Volume-High-Low-Close
- Volume-Open-High-Low-Close

In each case, the order of terms indicates the order of columns you should use in your chart. If you select Volume-High-Low-Close, then the leftmost column should contain the volume information, followed by another column with the stock's daytime high, and so on. (Technically, you can use the Chart Tools | Design→Data→Select Data command to specify each series, even if it's not in the place Excel expects it to be. However, this maneuver is tricky to get right, so it's easiest to just follow the order indicated by the chart type name.) No matter which subtype you use, a stock chart shows only values for a single stock.

Tip: The simplest stock chart (High-Low-Close) is also occasionally useful for charting variances in scientific experiments or statistical studies. You could also use a stock chart to show high and low temperature readings. Of course, you still need to follow the rigid stock chart format when ordering your columns.

Figure 17-22 shows an example of a Volume-High-Low-Close.

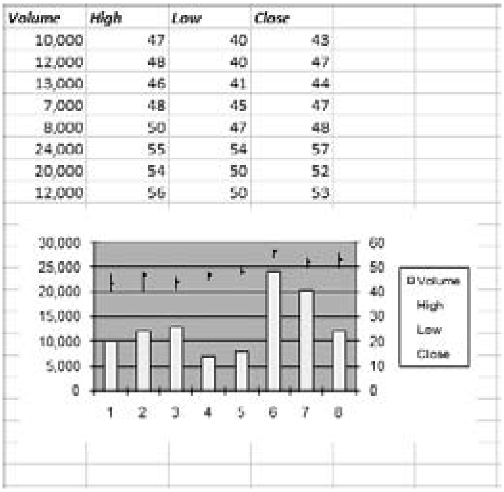

Volume	High	Low	Close
10,000	47	40	43
12,000	48	40	47
13,000	46	41	44
7,000	48	45	47
8,000	50	47	48
24,000	55	54	57
20,000	54	50	52
12,000	56	50	53

Figure 17-22:
The Volume-High-Low-Close chart shows a combination of related information. The columns at the bottom show the number of shares traded (using the value scale on the left). The lines above these columns show the stock price (using the value scale on the right). For each day, a bar that spans from the day's low to the day's high represents the stock price. The closing price is marked with a tick in the middle of the bar. If you like, you can add a trendline to show the movement of the stock price using the techniques explained in the next chapter.

Surface

A surface chart shows a 3-D surface that looks a little like a topographic map, complete with hills and valleys. Surface charts are different from most other charts in that they show the relationship of three values. Two category axes (X and Y) determine a data point's position. The value determines the height of the data point (technically known as the Z-axis). All of the points link to create a surface.

Surface charts are neat to look at, but ordinary people almost never create them, as they're definitely overkill for tracking your weekly workout sessions. For one thing, to make a good surface chart, you need a lot of data. (The more points you have, the smoother the surface becomes.) Your data points also need to have a clear relationship with both the X and Y axes (or the surface you create is just a meaningless jumble). Usually, rocket-scientist types use surface charts for highly abstract mathematical and statistical applications. Figure 17-23 shows an example.

Heat Index Calculator

Temp	Humidity 10	20	30	40	50	60	70	80	90	100
70	71.372408	74.2	76.1	77	76.9	75.9	73.9	70.9	66.9	61.97
75	74.612787	76	76.9	77.5	77.6	77.4	76.7	75.6	74.1	72.17
80	78.115694	78.6	79.2	79.9	80.8	81.8	83	84.2	85.6	87.19
85	81.88113	82	82.9	84.3	86.5	89.3	92.7	96.8	102	107
90	85.909094	86.3	87.9	90.7	94.6	99.7	106	113	122	131.6
95	90.199587	91.5	94.4	99	105	113	123	134	147	161.1
100	94.752609	97.5	102	109	118	129	143	158	176	195.3
105	99.568159	104	112	121	134	149	166	187	209	234.4

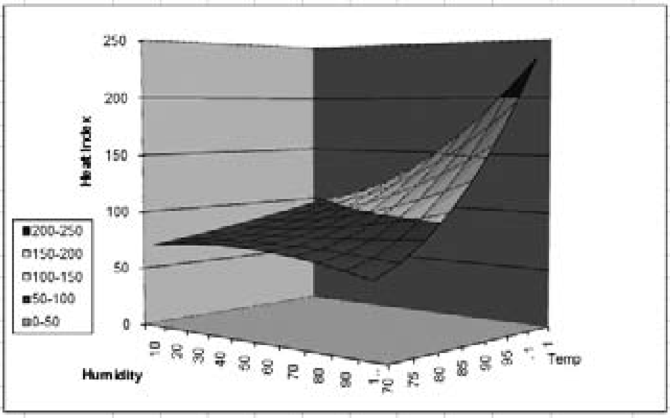

Figure 17-23: *A surface chart usually represents scientific data or mathematical models. Here, the surface chart represents a heat index table. Heat index tables show the perceived temperature under different conditions of humidity. If it's 70 degrees with 90 percent humidity, the dampness makes it feel like 66.9. On the other hand, if it's 85 with 90 percent humidity, it feels to a human observer like a balmy 102.*

Donut

The donut chart (Figure 17-24) is actually an advanced variation on that other classic food-themed chart, the pie chart. But while a pie chart can accommodate only one series of data, the donut can hold as many series as you want. Each series is contained in a separate ring. The rings are one inside the other, so they all fit into a single compact circle.

The donut chart's ideal for comparing the breakdown of two different sets of data. However, the data on the outside ring tends to become emphasized, so make sure this series is the most important. To change the order of the series, see page 480.

Donuts have two subtypes: standard and exploded. (An exploded donut doesn't suggest a guilty snack that's met an untimely demise. Instead, it's a donut chart where the pieces in the topmost ring are slightly separated.)

Although the donut chart can hold as many series as you want, if you add more than two or three, the chart may appear overly complicated. No matter what you do, the center of the donut is never filled in (unless you decide to add some text there using Excel's drawing tools, which are covered in Chapter 19).

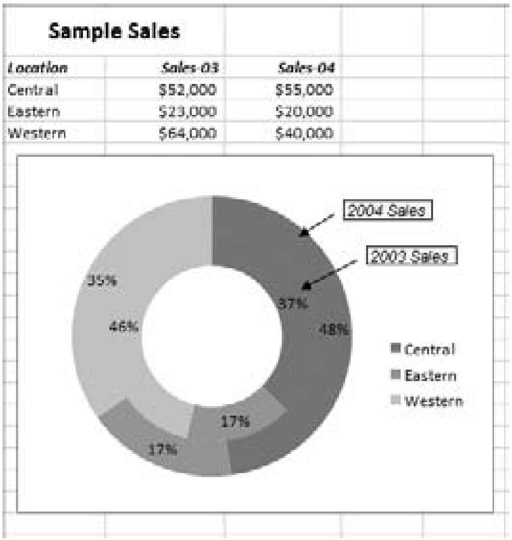

Figure 17-24:
Here, a donut chart compares the sales in two different years. The inner ring shows the sales for 2003, broken down by region. The outer ring shows the sales for 2004. Donut charts sometimes need a bit of customizing before they look right. Here, the chart has labels that detail the percentage of each slice, and the year represented by each ring. (You can learn how to add these labels using Excel's drawing tools in Chapter 19.)

Note: Think twice before you use a donut chart in a presentation. Most Excel gurus avoid this chart, because it's notoriously difficult to explain.

Bubble

The bubble chart is an innovative variation on the scatter chart. It plots only a single series, and it never draws a line. Each point is marked with a circle—either an ordinary circle or a 3-D sphere, depending on the subtype you choose. The extra frill is that the bubble sizes change based on a second set of related values. The larger the value, the larger the data point bubble. In any bubble chart, the largest bubble is always the same size. Excel scales the other bubbles down accordingly. Figure 17-25 shows an example.

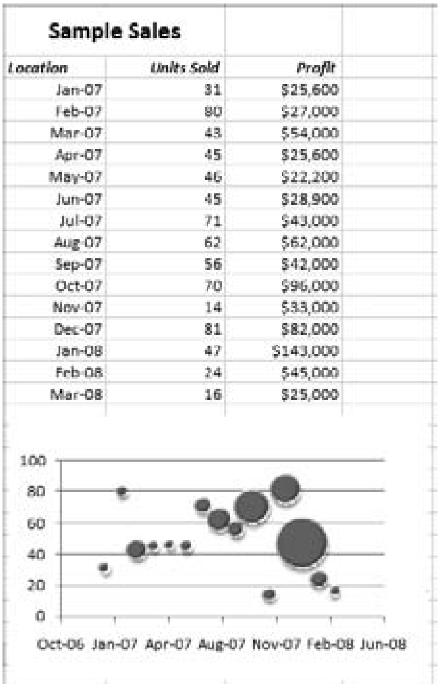

Sample Sales

Location	Units Sold	Profit
Jan-07	31	$25,600
Feb-07	80	$27,000
Mar-07	43	$54,000
Apr-07	45	$25,600
May-07	46	$22,200
Jun-07	45	$28,900
Jul-07	71	$43,000
Aug-07	62	$62,000
Sep-07	56	$42,000
Oct-07	70	$96,000
Nov-07	14	$33,000
Dec-07	81	$82,000
Jan-08	47	$143,000
Feb-08	24	$45,000
Mar-08	16	$25,000

Figure 17-25:
Each bubble's position represents two values: the month (the category axis) and the number of units sold (the value axis). Each bubble's size reflects the profit generated by the units sold.

Radar

The radar chart (Figure 17-26) is a true oddity, and it's typically used only in specialized statistical applications. In a radar chart, each category becomes a spoke, and every spoke radiates out from a center point. Each series has one point on each spoke, and a line connects all the points in the series, forming a closed shape. All these spokes and lines make the chart look something like the radar on an old-time submarine.

You have three radar subtypes: the standard radar chart; a radar chart with data markers indicating each point; and a filled radar, where each series appears as a filled shape, somewhat like an area chart. No matter what subtype you use, choosing data for a radar chart isn't easy.

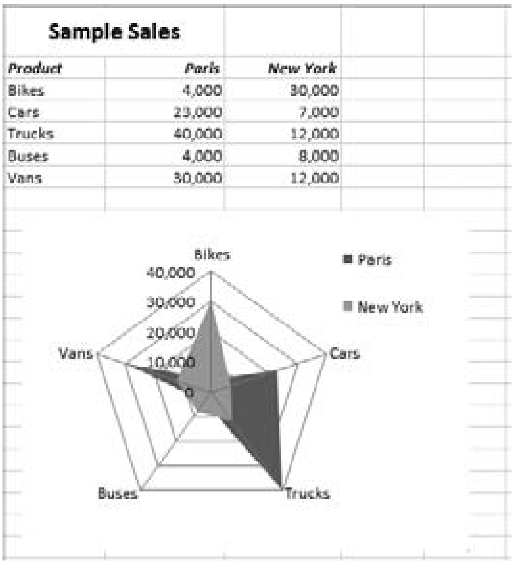

Figure 17-26:
This filled radar chart compares the products sold in sales offices in two cities. Because all the categories (in this case, the various products) are joined into a closed shape, the radar chart acts somewhat like an area chart, so you can judge the significance of values by looking at the size and shape of the area. It's easy to see that bicycles are selling well with customers in the New York store, while trucks lead the way in Paris. If the two series had similar results, you couldn't effectively use a filled radar chart, because some of the data would be obscured.

CHAPTER
18

Formatting and Perfecting Charts

In the previous chapter, you learned how Excel charts work and how you can transform ordinary tables of information into graphical representations, complete with columns, bars, lines, and even bubbles. But creating the right chart is only half the battle. The next step is *refining* your charts so they convey their point more effectively. Often this step means tweaking the chart's formatting, inserting labels, and fine-tuning the scale. But if you're really ambitious, you'll want to tackle more advanced professional charting techniques, like trendlines, overlays, and combination charts. These techniques let you turn plain-vanilla charts into polished graphics—like the ones you see in magazines, annual reports, and brochures.

In this chapter, you'll start by looking at how you can use basic formatting techniques to change the color and font of different chart components. Then you'll learn how to set a chart's scale, unleash 3-D views, and make your data stand out no matter what type of chart you use.

Chart Styles and Layouts

In Chapter 6, you learned how *cell* styles let you take ready-made formatting and apply it to your data to glitz up the dullest worksheet. For the same reason, Excel provides a set of *chart* styles that you can use to give the plainest chart a makeover.

Like cell styles, chart styles draw from the colors, fonts, and shape effects that are part of each theme (page 174). If you use the Trek theme, your chart styles draw upon a palette of earthy tones, while the Verve theme gives you a much more vivid set of colors. If you use cell styles, table styles, and chart styles, the fonts and colors are consistent everywhere. You can also swap in a new palette for all these elements just by choosing a new theme.

Tip: Before you choose a chart style, it helps to pick the theme you want to use so you can see the real, final result. To change the theme, make a selection from the Page Layout→Themes→Themes list. Page 175 has the full story.

Chart Styles

Chart styles give you a way to apply shake-and-bake formatting to ordinary charts. Excel includes a wide range of chart styles that vary from simple, flat charts with minor accents (like colored borders) to showier styles that include bevel effects and shadows. You can quickly create plain or opulent charts, depending on your needs.

Before you use a chart style, it's important to understand what it changes (and what it doesn't). Every chart style includes settings that determine:

- The background color.
- The colors for each series (which may be displayed as bars, lines, points, or something else). Some colorful chart styles use all the colors from the current theme. Other chart styles take a single color from the current theme and use multiple shades of that color for a more refined look.
- Shape effects, like softly curved or beveled edges, shadows, and glow.
- Marker styles (for line and XY scatter charts) that distinguish the points in one series from those in another.

Note: Some chart styles use a heavy black background with bold colors. This sort of style isn't designed for worksheets because it can tie the best color printer in knots. But these high-contrast styles look good on computer monitors and projection screens, so use them if you want to cut and paste your chart into a PowerPoint presentation. (Chapter 27 has more about transferring chart objects and other data between Excel and other programs.)

On the other hand, chart styles don't change the font Excel uses for the chart title and labels; instead Excel bases these elements on the current theme. Chart styles also don't change the layout of the chart or the chart settings Excel uses for the legend, scale, axis titles, error bars, and so on. (You haven't yet seen how to tweak all these details, but you'll learn about them later in this chapter.)

To choose a style, select the chart, and then head to the ribbon's Chart Tools | Design→Chart Styles section, which gives you a gallery of choices (Figure 18-1).

Note: If you've resized your Excel window to be very small, the chart styles don't fit in the ribbon and you see a Quick Styles button instead. Click this button to see the full list of styles.

A chart can use only one style at a time, so if you choose a new style, Excel wipes out the existing formatting.

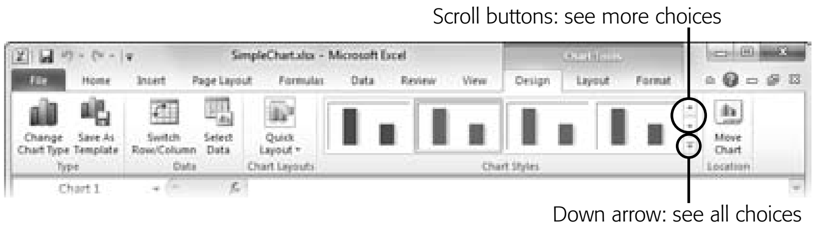

***Figure 18-1:** Click any of the Chart Style options to dress your chart up in different colors. The scroll bars (circled) let you move through the style options not currently shown, or click the down arrow to see a drop-down window with all your choices.*

Note: One key difference between chart styles and other types of styles (like cell styles and table styles) is that you can't create your own chart styles. However, you can choose a chart style to use as a starting point, further customize it, and then save it as a template so you can reuse it again and again. You'll learn how this feature works on page 518.

Chart Layouts

Chart styles make it easy for you to change the colors and visual styling in a chart. Chart layouts are complementary—they let you control the presence and placement of various chart elements, like the chart and axis titles, and the legend.

As you'll learn in the next section, Excel lets you tweak each of these ingredients separately. However, you can choose a prebuilt layout to do it all in one shot. To try it out, head to the ribbon, and then make a choice from the Chart Tools | Design→Chart Layouts section. As with styles, the list of available charts depends on the chart type. Figure 18-2 shows an example.

Note: To make chart layouts as practical as possible, the creators of Excel reviewed thousands of professional charts and identified the most common arrangements. Most Excel pros still want to customize the various parts of their chart by hand. However, a chart layout can provide a great starting point.

Adding Chart Elements

You build every chart out of small components, like titles, gridlines, axes, a legend, and the bars, points, or exotic shapes that actually represent the data. And Excel lets you manipulate each of these details separately. That means you can independently change the format of a label, the outline of bar, the number of gridlines, and the font and color of just about everything.

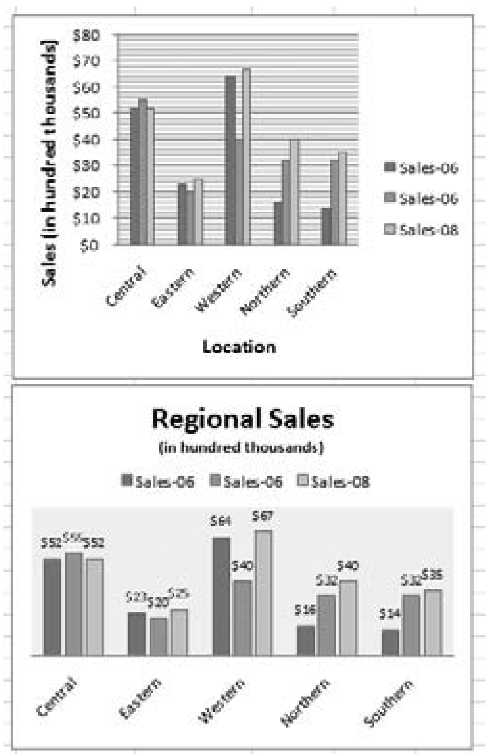

Figure 18-2:
This worksheet shows two copies of the same chart, each with a different layout. The chart at the top includes heavy gridlines, axis titles, and a legend on the right. The chart below has a chart title and places the legend at the top. It also has no gridlines and instead displays the series value above each column.

Figure 18-3 shows the different elements that make up a chart. They include:

- **Title**. The title labels the whole chart. In addition, you can add titles to the axes. If you do, then you can select these titles separately.
- **Legend**. The legend identifies each data series on the chart with a different color. A legend's useful only when the chart contains more than one series.
- **Horizontal and Vertical Axes**. An axis runs along each edge of the chart and determines the scale. In a typical two-dimensional chart, you have two axes: the category axis (typically on the bottom of the chart, running horizontally), and the value axis (typically on the left, running vertically).
- **Plot Area**. The plot area is the chart's background, where the gridlines are drawn. In a standard chart, the plot area is plain white, which you can customize.
- **Chart Area**. The chart area is the white space around the chart. It includes the space that's above, below, and to either side of the plot area.
- **Gridlines**. The gridlines are the lines that run across the plot area. Visually, you line the data up with the gridlines to get an idea of the value of each data point.

Every chart starts out with horizontal gridlines, but you can remove them or add vertical gridlines. You can tell Excel how many gridlines to draw, and even how to format them.

- **Data Series**. The data series is a single set of data plotted on the chart across the category axis. In a line chart, for example, the data series is a single line. If a chart has multiple series, you'll often find it useful to format them separately to make them easier to differentiate or to emphasize the most important one.
- **Data Point**. A data point is a single value in a data series. In a line chart, a data point's a single dot, and in a column chart, a data point is a single column. If you want to call attention to an exceptionally important value, you can format a data point so it looks different from the rest of the series.

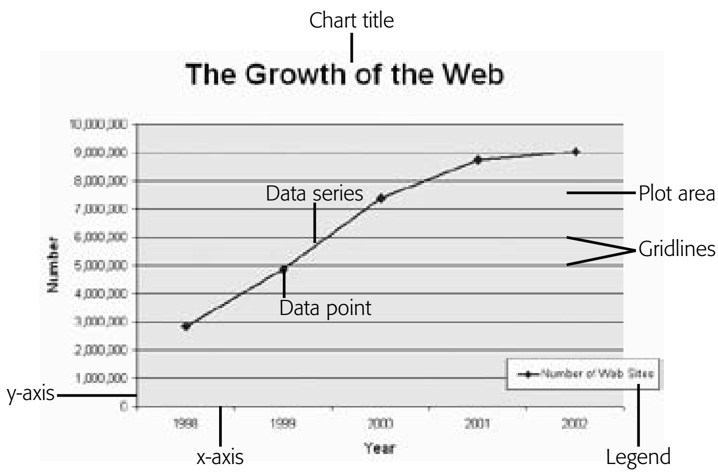

Figure 18-3: *Before you can begin tweaking your chart's formatting, you need to know the names of the different elements you'll find on a chart, all of which are labeled here.*

Not all charts include all these elements. As you learned on page 497, the layout you pick determines whether you begin with a chart title, a legend, gridlines in the background, and so on. However, in many cases you'll want to pick and choose exactly the elements you want. Excel lets you do this with the buttons on the ribbon's Chart Tools | Layout tab. The following sections show you how.

Adding Titles

It doesn't matter how spectacular your chart looks if it's hard to figure out what the data represents. To clearly explain what's going on, you need to make sure you have the right titles and labels.

An ordinary chart can include a main title (like "Increase in Rabbit Population vs. Decrease in Carrot Supplies") and titles on each axis (like "Number of Rabbits" and "Pounds of Carrots"). To show or hide the main title, make a selection from the Chart Tools | Layout→Labels→Chart Title list. Your options include:

- **Above Chart** puts a title box at the very top and resizes the chart smaller to make room.
- **Centered Overlay Title** keeps the chart as is but superimposes the title over the top. Assuming you can find a spot with no data, you get a more compact display.
- **None** hides the title altogether.

Once you select one of those options, you see the title box; click inside it, and then type in new text, as shown in Figure 18-4.

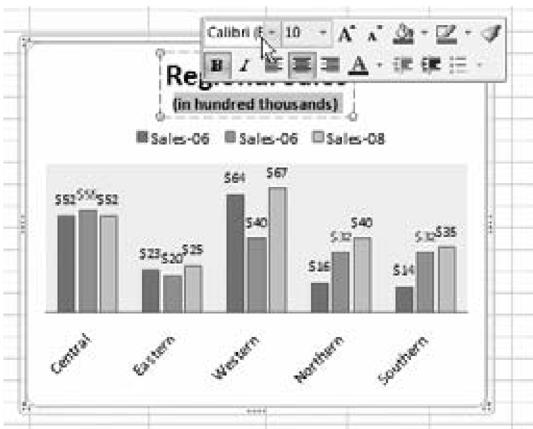

Figure 18-4: *You can type in whatever text you'd like for a chart title. If you select part of the text, a mini bar appears (sadly, of the alcohol-free variety), with formatting options for changing the font, size, color, and alignment. These commands are the same as in the Home→Font section of the ribbon, but it's way more convenient to reach them here.*

You can just as easily add a title to each axis using the Chart Tools | Layout→Labels→Axis Titles→Primary Horizontal Axis Title and Chart Tools | Layout→Labels→Axis Titles→Primary Vertical Axis Title lists. You'll find options for showing your title, hiding it, and (in the case of a vertical axis), showing a title that's rotated to run neatly along the side of your chart.

Note: As with almost all chart elements, you can also format titles by adding a border, a shadow effect, and a fancy background fill. To get these options, right-click the title, and then choose Format Chart Title. You'll learn more about these options throughout this chapter.

Adding a Legend

Titles help explain a chart's overall purpose. Usually, titles indicate what a chart is comparing or analyzing. You may add a chart title like "Patio Furniture Sales" and the axis labels "Gross Revenue" and "Month of Sale" to a chart that shows how patio furniture sales pick up in the summertime. However, the category labels don't help you single out important data. They also don't let you point out multiple series (like the sales results in two different stores). You can fix this problem by adding additional labels or a *legend*. A legend is a separate box off to the side of the chart that contains one entry for each data series in a chart. The legend indicates the series name, and it adds a little sample of the line style or fill style that you've used to draw that series on the chart.

Excel automatically adds a legend to most charts. If you don't already have a legend, you can choose a layout that includes one, or you can make a selection from the Chart Tools | Layout→Labels→Legend list. Different selections let you position the legend in different corners of the chart, although true Excel pros just drag the legend box to get it exactly where they want it.

Legends aren't always an asset when you need to build slick, streamlined charts. They introduce two main problems:

- **Legends can be distracting**. In order to identify a series, the person looking at the chart needs to glance away from the chart to the legend, and turn back to the chart again.
- **Legends can be confusing**. Even if you have only a few data series, the average reader may find it hard to figure out which series corresponds with each entry in the legend. This problem becomes more serious if you print your chart out on a printer that doesn't have the same range of colors as your computer monitor, in which case different colored lines or blocks may begin to resemble each other.

If you don't want to use a legend for these reasons, you can use data labels instead, as described in the next section.

Adding Data Labels to a Series

Data labels are labels that you attach to every data point in a series. This text floats just above the point, column, or pie slice that it describes, clearly identifying each piece of information. Data labels have unrivalled explaining power—they can identify *everything*. The only possible drawback is that adding data labels to a chart that's already dense with data may lead to an overcrowded jumble of information.

To apply data labels, choose a position from the Chart Tools | Layout→Labels→Data Labels list. If you choose Chart Tools | Layout→Labels→Data Labels→Center on a column chart, each bar's value appears as a number that's centered vertically inside the bar. On the other hand, if you choose Chart Tools | Layout→Labels→Data Labels→Outside End, the numbers appear just above the top of each column, which is usually more readable (Figure 18-5).

Tip: No matter how you choose to label or distinguish a series, you're best off if you don't add too many of these elements to the same chart. Adding too many labels makes for a confusing overall effect, and it blunts the effect of any comparison.

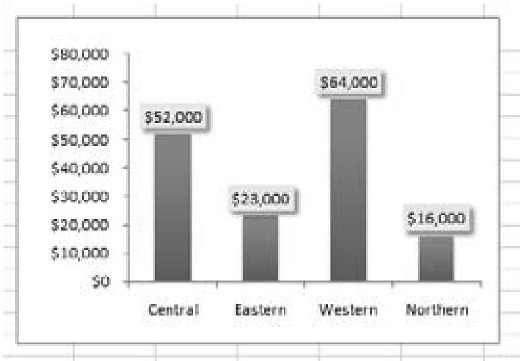

Figure 18-5:
Here, you can see how a value label adds information to a column chart. Even without the labels, you could still get this information by eyeing where the bar measures up to on the value axis on the left, but the labels make it a whole lot easier to get the information with a single glance. The labels have been customized slightly via the Format Data Labels dialog box to shrink their font size and add a simple box border with a shadow effect.

If you're in an adventurous mood, you can create even more advanced labels by choosing Chart Tools | Layout→Labels→Data Labels→More Data Label Options. The Format Data Labels dialog box appears, with a number of options for customizing data labels (Figure 18-6).

Using the Format Data Labels dialog box, you can choose the data label's position (just like you could from the Chart Tools | Layout→Labels→Data Labels list). But the options under the Label Contains heading are more interesting, as they let you chose the information that appears in the label. Ordinarily, the information is simply the value of the data point. However, you can also apply a *combination of* values. Your exact options depend on the type of chart you've created, but here are all the possible choices:

- **Series name**. The series name identifies the series each data point comes from. Because most series have multiple data points, using this option means the same text repeats again and again. In a line chart that compares sales between two stores, using this option would put the label "Store 1" above each data point for the first store, which is probably overkill.
- **Category name**. The category name adds the information from the category axis. If you're using a line chart to compare how sales fluctuated month by month, then this option adds the month above every data point. Assuming you have more than one line in your line chart, this option creates duplicate labels, which crowds out the important information in your form. For that reason, category labels don't work very well with most charts, although you can use them to replace the legend in a pie or donut chart.
- **Value**. Value labels insert the data that corresponds with a data point. This data is the actual information in the corresponding cell in your worksheet. If you're plotting changing sales, this data is the dollar amount of sales for a given month. Value labels are probably the most frequently used type of label.

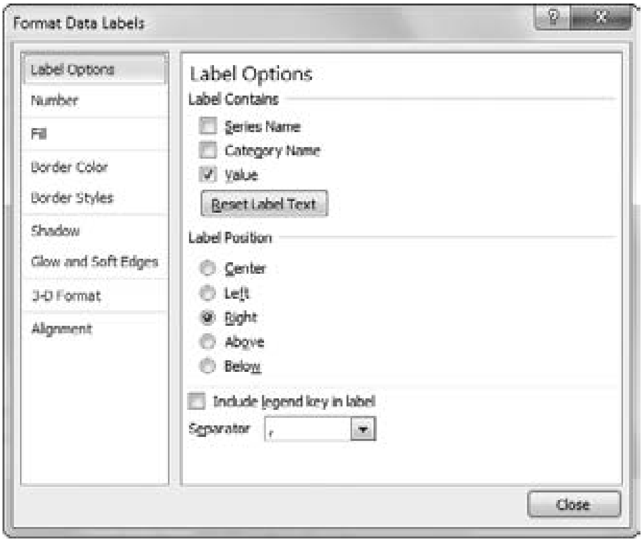

Figure 18-6: *The Format Data Labels dialog box is divided into several sections of settings. (You can see one setting at a time by picking from the list on the left.) For now, you're interested in the Label Options section. You'll learn how to use the other formatting settings, which apply to all chart elements, later in this chapter.*

- **Percentage**. Percentage labels apply only to pie charts and donut charts. They're similar to value labels, except they divide the value against the total of all values to find a percentage.
- **Bubble size**. Bubble size labels apply only to bubble charts. They add the value from the cell that Excel used to calculate the bubble size next to each bubble. Bubble labels are quite useful in bubble charts because bubble sizes don't correspond to any axis, so you can't understand exactly what numeric value a bubble represents just by looking at the chart. Instead, you can judge relative values only by comparing the size of one bubble to another.

Note: In some charts (including XY scatter charts and bubble charts), the checkboxes "Category name" and "Value" are renamed as "X Value" and "Y Value", although they have the same effect as "Category name" and "Value."

When you use multiple items, you can also choose a character from the Separator list box to specify how to separate each piece of text in the full label (with a comma, space, semicolon, new line, or a character you specify). And if you want to display a mini square with the legend color next to the label, then choose "Include legend key in label" (although most people don't bother with this feature).

Figure 18-7 shows more advanced data labels at work.

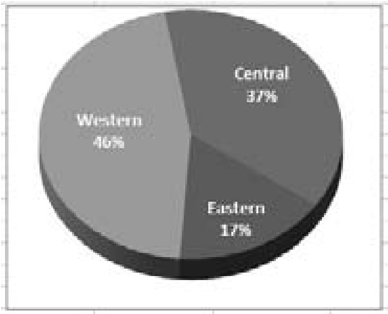

Figure 18-7:
Here's how you can combine percentage and category information to make a pie chart more readable and eliminate the legend altogether.

Note: Wondering what your chart will look like? As you make changes, Excel updates the chart on the worksheet using its handy live preview feature. Just move the Format Data Labels dialog box out of the way to get a sneak peak before you confirm your choices.

Adding Individual Data Labels

In simple charts, data series labels work well. But in more complex charts, data series labels can be more trouble than they're worth because they lead to chart overcrowding, particularly with line charts or any chart that has multiple series. The solution is to add labels to only a few data points in a series—those that are most important. Figure 18-8 shows the difference.

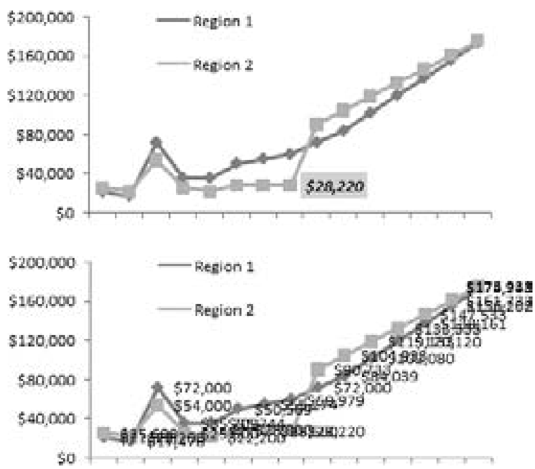

Figure 18-8:
Data point labels work particularly well with line charts and scatter charts because both chart types are dense with information. The two examples here underscore that fact.

Top: Here, a single data point label indicates the point where the sales changed dramatically for the Region 1 office.

Bottom: Here's the mess that results if you add data labels to the whole Region 1 and Region 2 series. No amount of formatting can clear up this confusion.

To add an individual data label, follow these steps:

1. **Click the precise data point that you want to identify. This point could be a slice in a pie chart, a column in a column chart, or a point in a line chart.**

   Selecting a data point is a little tricky. You need to click twice—the first click selects the whole series, and the second click selects just the data point you want. You'll see the handles appear around the specific column or point to indicate you've selected it, as shown in Figure 18-9.

2. **When you have the right data point selected, choose an option from the Chart Tools | Layout→Labels→Data Labels list.**

   These options work the same way they do when you format the whole series, except now they apply only to the currently selected value.

   To remove a data label, click to select it, and then press Delete. If you want to add several data labels, you're best off adding all the data labels (as described in the previous section), and then deleting the ones you don't want.

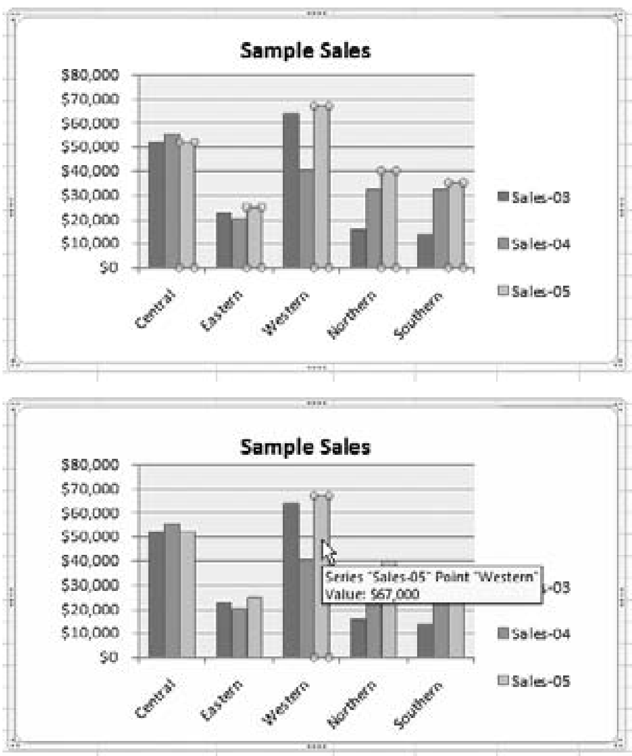

Figure 18-9:
Top: To select a single data point, click it twice. The first click selects the whole Sales-09 series.

Bottom: The second click gets just the data point you want, which is the Sales-05 data in the Western sales office. In some cases, you may not be able to tell which data point's the one you want, especially if you're creating a dense scatter chart. When that happens, just hover over the data point briefly to see a tooltip with the category, series, and value information.

Tip: If a data label doesn't have exactly what you want, you can click inside it and edit the text, just as you do with a chart title.

As with data series labels, right-clicking a data label gives you a choice of formatting options (choose Format Data Labels). It's almost always a good idea to format your labels so they stand out and don't crowd other information on the chart. Page 509 describes more about formatting chart elements.

Tip: Instead of using data labels, you can add arrows and text boxes anywhere on a chart to call out important information. To do so, you need Excel's drawing features, explained in the next chapter.

Adding a Data Table

Trying to pack as much information as possible into a chart—without cluttering it up—is a real art form. Some charting aficionados use labels, titles, and formatting to highlight key chart details, and then use the data on the worksheet itself to offer a more detailed analysis. However, Excel also provides a meeting point between chart and worksheet that works with column charts, line charts, and area charts. It's called the *data table*.

Excel's data table feature places your worksheet data *under* your chart, but lined up by category. You can best understand how this feature works by looking at a simple example, like the one in Figure 18-10.

To add a data table, select your chart, and then choose Chart Tools | Layout→Labels→Data Table→Show Data Table. Or, if you want each series in the data table to have a small square next to it with the same color as the matching data series, then choose Chart Tools | Layout→Labels→Data Table→Show Data Table with Legend Keys. This way, you might not need a legend at all.

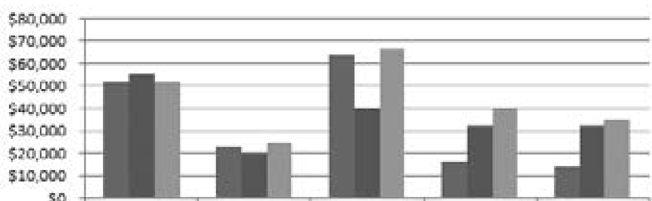

	Central	Eastern	Western	Northern	Southern
Sales-03	$52,000	$23,000	$64,000	$16,000	$14,000
Sales-04	$55,000	$20,000	$40,000	$32,000	$32,000
Sales-05	$52,000	$25,000	$67,000	$40,000	$35,000

Figure 18-10: *This data table removes the need for a legend or data point labels. However, keep in mind that data tables don't work well with large amounts of data.*

Selecting Chart Elements

Some chart elements, like titles and axes, are obvious. Others are a little more subtle. But it's important to realize that everything you add to a chart is its own distinct ingredient, and you can tweak each ingredient independently—even the data labels that identify important values (page 504). That means you can select them, move them, delete them, and so on. Understanding this principle is the secret to creating charts that stand out from the crowd.

So far, you've covered how to add the most popular elements to your chart, using either a readymade layout or the lists in the Chart Tools | Layout tab. But you haven't learned about everything you can do with a chart element once you've created it. In this section (and those that follow), you'll learn how to grab hold of the element you want and further fine-tune it.

If you want to manipulate your chart components, you have to first learn how to select them. The easiest way to do so is to click directly on the element you want to change, as shown in Figure 18-11.

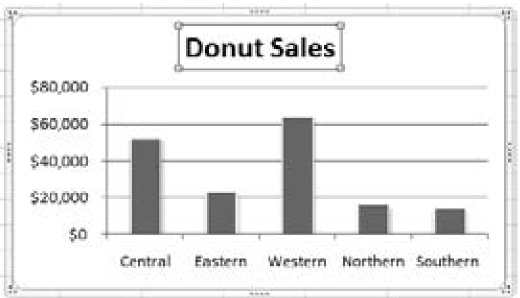

Figure 18-11:
When you select a part of a chart, a rectangle appears around the item, with a circle (known as a handle) in each corner. In this example, the chart title is selected.

Some elements—like titles and legends—are easy to select this way. Others, like gridlines and axes, are a little trickier because you need to click in exactly the right place. If you end up selecting the wrong element, it's sometimes easiest to just use the arrow keys to select another one. Each time you press an arrow key, Excel selects the next chart element. So, if you keep pressing an arrow key, you'll eventually cycle through all the elements that you can select in the current chart.

You can also select chart elements with the ribbon's Chart Tools | Format→Current Selection section, as shown in Figure 18-12. Using this list, you can select any of the chart elements described earlier (page 498), except individual data points.

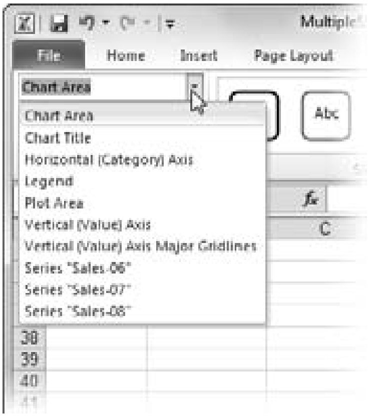

Figure 18-12:
Click the drop-down arrow in the box at the top of this section (Chart Tools | Format→Current Selection) to show a list of all the elements in the current chart. When you choose an item from the list, Excel automatically highlights the corresponding part of the chart.

Tip: To check which element you've selected, just right-click it. The last item in the shortcut menu makes everything clear. If the context menu starts with Format Axis, you know that you've selected a chart axis. If it says Format Gridlines, you've got yourself a gridline. Or, you can check the Chart Tools | Format→Current Selection section of the ribbon.

So what do you do with a chart element once you've selected it? You can perform three basic tasks, although not every chart element supports every task. The tasks include:

- **Deleting an element**. To remove a selected element from your chart, just press Delete. You can delete any chart element, including titles, legends, data series, gridlines, the background, and even an axis. It's possible to restore these elements after you've deleted them by choosing a new layout (page 497) or using the Undo command (Ctrl+Z) immediately after.
- **Moving an element**. You can move a chart element by dragging it, just like the way you move the whole embedded chart. You can easily move some elements, like labels, legends, and the whole chart. Other chart elements—including axes, gridlines, and data points—can't go anywhere.
- **Resizing an element**. You can resize an element by dragging one of the *resizing handles* (the circles that appear at each corner of the element box when you select it). Resizing is mainly useful with legends and the overall chart—there aren't many other chart elements you can resize.

Formatting Chart Elements

Often, you don't select a chart element to delete or move it, but rather to format it with a different border, font, or color. In this case, simply right-click the element, and then choose the format option from the pop-up menu (or, select it, and then choose Chart Tools | Format→Current Selection→Format Selection from the ribbon). The following sections get you started.

Coloring the Background

Now you're ready to start creating spiffy-looking, customized charts. The background color is a good starting point. Initially, this color is a plain white, but it's easy enough to change if you want to add a personal touch. Just follow these steps:

1. **Select the plot area.**

   To select the plot area, click the empty space between gridlines, or choose Chart Tools | Format→Current Selection→Plot Area.

2. **Right-click the plot area, and then choose Format Plot Area, or choose Chart Tools | Format→Current Selection→Format Selection.**

   A Format Plot Area dialog box appears (Figure 18-13). Here you can set the background color (Fill), add a border around it (Border Color and Border Styles), and use advanced Shadow and three-dimensional effects (Shadow and 3-D Format). For now, focus on the Fill tab.

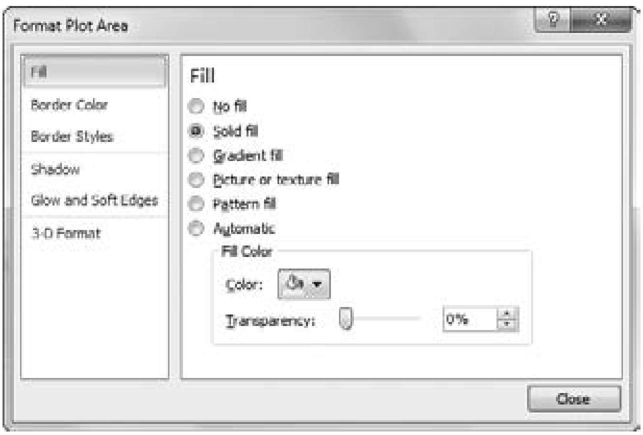

Figure 18-13: *It's worth getting familiar with all the formatting options in the Format Plot Area dialog box, because you'll see many of the same options again when you format other chart elements, like columns, gridlines, and just about everything else.*

3. **Choose "Solid fill", and then click the paint can button to choose from a list of colors.**

   For the most flexible formatting, choose one of the theme colors (which appear under the heading Theme Colors). That way, if you choose a new theme, then the chart switches itself to one of the new theme colors. Don't worry about the fancy fills for now—you'll take them on in the next section.

4. **When you've finished making your changes, click OK.**

   Excel applies the results of your artistry.

Note: Remember, if you don't have a color printer, you need to think about how colors convert when you print them in black and white. In some cases, the contrast may end up being unacceptably poor, leading to charts that are difficult to read. As a general rule, the less powerful your printer, the less you should use graphically rich details like tiles, background images, and gradients–unless, of course, you're planning to view your worksheet only onscreen.

The neat thing about this sequence of steps is that you can use exactly the same process to format *any* chart element. That means you now know enough to give a solid fill to a chart title, the gridlines, the columns in a column chart, and so on. And once you learn your way around the rest of the formatting options, you'll be able to really spiff up your chart.

Tip: If you run rampant changing a chart element and you just want to return it to the way it used to be, select it, and then choose Chart Tools | Format→Current Selection→Reset to Match Style. Excel removes your custom formatting and leaves you with the standard formatting that's based on the chart style.

Fancy Fills

Coloring the background of a chart is nice, if a little quaint. In the 21st century, charting mavens are more likely to add richer details like textured backgrounds or gradient fills. Excel gives you these options and more. And although textured fills don't always make sense, they can often add pizzazz when used in the background of a simple chart. You can apply fancy fills to the chart background (the plot area) or individual chart items, like the columns in a column chart. Figure 18-14 shows some of your fill choices.

To apply a fancy fill, start by selecting the chart element, and then choosing Chart Tools | Format→Current Selection→Format Selection to get to the appropriate formatting window (like Format Plot Area). Then, choose the Fill section. What you do next depends on the type of fill you want.

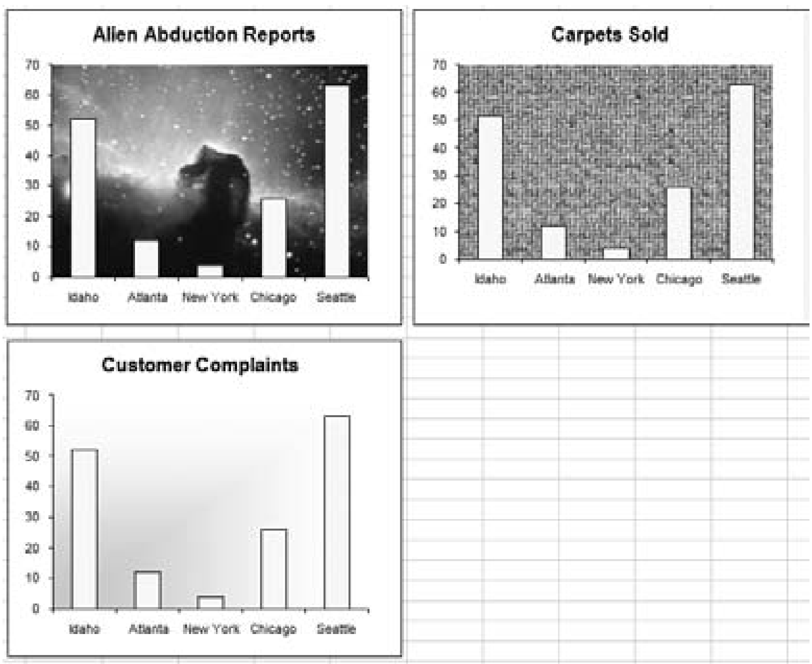

Figure 18-14: *Depending on the type of fill you use, you can subtly change the message your chart conveys. These figures show a picture fill (top left), texture fill (top right), and gradient fill (bottom), each with the same numbers and chart type. Fancy fills are particularly useful if you want to use your charts in a presentation program like PowerPoint. (Chapter 27 covers how to integrate your Excel worksheets with other programs.)*

Gradient fills

A gradient is a blend between two colors. You may use a black-and-white gradient that gradually fades from black in the top-left corner to white in the bottom-right corner. More complex gradients fade from one color to another to another, giving a 1960s-era tie-dye effect. To set a gradient, in the formatting window's Fill section, choose the "Gradient fill" choice.

You can use two basic strategies to choose your gradient: the colors and the shading pattern. For the utmost simplicity, you can use a prebuilt set of color and shading options. You choose this option by clicking the "Preset colors" button, and then clicking one of the thumbnail previews that appears in the drop-down list. Each one has a picturesque name like Late Sunset or Ocean. (As you make a choice here, Excel updates your chart to show the results you'll get if you apply the new fill.) The box below explains your custom gradient options.

Texture fills

A texture is a detailed pattern that's tiled over the whole chart element. The difference between a texture and an ordinary pattern is that patterns are typically simple combinations of lines and shading, while a texture uses an image that may have greater, more photographic detail.

POWER USERS' CLINIC

Custom Gradients

If you're an unrelenting micromanager, you may want to choose your own gradient colors. But be forewarned—Excel lets you design gradient fills in truly ridiculous detail. Here's the recipe you need to follow to cook up your own custom gradient.

If you're just starting out, begin with one of the choices from the "Preset colors" list. You can then modify it.

1. **From the Type list, choose a gradient type**.

   The type determines how the fill's shaded. A normal linear gradient shades colors from one side (like the top) to the other (the bottom). The more exotic radial shades colors in concentric rings, starting from a single point.

2. **From the Direction list, choose an option**.

   The directions you can use depend on the type you chose in the previous step. If you're using a linear fill, you can choose whether the gradient starts at the top, bottom, left, or right. (For each option, a tiny thumbnail preview shows you what it looks like.) If you're using a linear fill, you can also set an angle in the Angle text box to tilt the gradient just a bit, so that it's not completely horizontal or vertical.

3. **In the Gradient Stops section, click the Add and Remove buttons until you have the right number of stops**.

   A gradient stop is a color in a gradient. In a simple gradient (like a blue-to-white shading), you have two colors, which means you have two stops, one for each color. If you have a more complex gradient (like a tacky yellow-green-red number), you need three gradient stops, one to put in each color. Of course, your gradient stops don't need to be different colors; you could use different tints or shades of the same color.

4. **Modify each gradient stop to suit**.

   Select a gradient stop from the list. When you pick one, you can adjust the color and the *stop position*. The stop position is a percentage that determines when the gradient switches from one color another. Say you create three gradient stops, with the colors yellow, green, and red. You may choose gradient stops of 0%, 50%, and 100% to space them out evenly. That way, the fill is yellow at the starting point (0%), green in the mid-point (50%), and red at the end (100%).

5. **Click Close**.

To choose a texture for a fill, click "Picture or texture fill". You can then choose one of the ready-made textures from the drop-down texture list (Figure 18-15).

Further down the window, you see options that let you control exactly how Excel tiles your texture. You can size your texture to be larger or smaller by adjusting the scale percentages, and you can play with the offset settings to alter how the tiles of texture overlap. Finally, you can use the Mirror list to flip the texture around. But in truth, all these options are excessive frills, and you'll rarely need to touch any of them once you have a texture you like.

If the ready-made list of textures doesn't have what you want, you can choose your own texture from a picture you have in a separate file, which is the next section's topic.

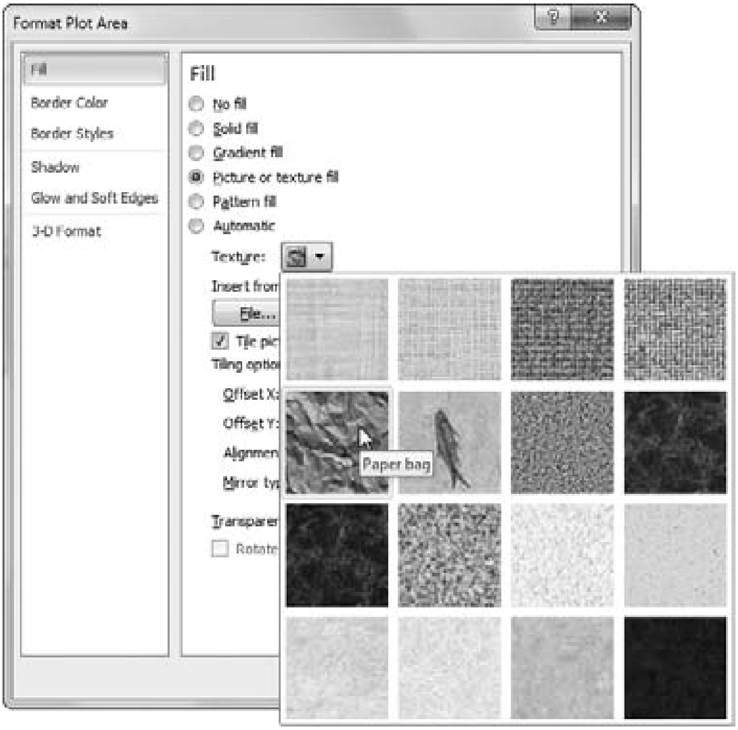

Figure 18-15: *When you choose a texture, you have choices like Woven Mat, Canvas, and Blue Tissue Paper. When you apply the texture, it actually looks like its name. Most textures are great for presentations, although potentially problematic for all but the most capable printers.*

Note: You'll notice that the Fill section has a slider bar that you can use to set the degree of transparency you want. You can make a fill partially transparent so that other elements show through. You can see a chart background through a partially transparent chart column. However, transparency is difficult to get right, and it often makes ordinary charts harder to read. But if your boss is out of the office and you need to fill the next hour, go ahead and experiment!

Picture fills

A picture is a graphical image that goes behind your chart and stretches itself to fit. Excel doesn't provide any ready-made pictures. Instead, you'll need to browse to a graphics file on your computer (a .bmp, .jpg, or .gif file). This option works well if you need a themed chart—like a beach scene behind a chart about holiday travel choices. If you just want to add a company logo somewhere on your chart, you're better off using the drawing tools described in the next chapter to place the logo exactly where you want it.

To use a picture fill, choose the "Picture or texture fill" option, and then click the File button to browse for the picture you want to use. Once you've picked the right picture, you can adjust the other options in the Fill section. Start by making sure the "Tile picture as texture" checkbox isn't selected—if it is, Excel tiles your picture just like the textures you saw in the previous section. (On the other hand, if that's the effect you're looking for, click away.)

Ordinarily, Excel stretches a picture over the surface of the chart. However, if you want your picture to fill just a part of the chart, you can adjust the different offset percentages (Top, Left, Right, and Bottom).

Note: If you just want to add an unstretched image or two somewhere on your chart, you shouldn't use a picture fill. Instead, add a picture object, as described in the next chapter.

Fancy Borders and Lines

Now that you've tweaked the background fill to be slick and sophisticated (or wild and crazy), you're ready to modify other details. Along with the fill, the border is the next most commonly modified detail. You can add a border around any chart element, and your border can sport a variety of colors, line thicknesses, and line styles (like dashed, dotted, double, and so on).

To set a line from a format window (like Format Plot Area), follow these steps:

1. **Choose the Border Color or Line Color section.**

   Both these sections are basically the same. The difference is that the name changes depending on the chart element you're tweaking. Some chart elements (like the plot area) can have a border, so they have a Border Color section. Other elements (like the gridlines) contain lines, so they have a Line Color section.

2. **Choose the type of line you want. Usually, it's "Solid line".**

   Obviously, "No line" removes the line altogether. Your only other option is "Gradient line", which lets you create a line that's painted with a gradient fill of colors (page 511). The reason this odd feature exists may never be entirely clear.

3. **Choose the Border Styles or Line Style section.**

   Here's where you really get the chance to customize your border or line. Once again, some chart elements get a Border Styles section, while others have an almost identical Line Style section.

4. **Set the line settings, and then click Close when you're finished.**

   Figure 18-16 shows your options.

Figure 18-16: *Using the Border Styles section, you can set a width (the thickness of the line), choose a fancy double or triple-edged border (from the "Compound type" list), choose between solid and different types of dotted or dashed lines (from the "Dash type" list), and even set what the line corners look like (the "Join type") and what the line ends look like (the "Cap type"). The "Cap type" doesn't have an effect on borders unless you're using dashed lines, because that's the only situation where the line is broken, and so has visible ends.*

Note: Most of the time, you probably won't bother putting borders around chart elements. However, you can use the options in the Line Style section to configure the gridlines that appear behind your chart data. Just select the gridlines (you can use the Chart Tools | Format→Current Selection list if you're having trouble clicking the right spot), and then choose Chart Tools | Format→Current Selection→Format Selection. You'll see the familiar Border Color and Border Styles sections that let you change the line color, thickness, and dash style.

POWER USERS' CLINIC

Using Shadows

The various formatting windows have another option. You can choose the Shadow section to apply a shadow effect to a title, data series, or some other part of a chart. (A shadow, of course, is a faint shaded region behind a graphical shape or piece of text. Usually, you use a shadow to make something look more three-dimensional, so it seems to float above the rest of the background.) And although the idea sounds a little strange, a soft shadow can make an ordinary chart seem more professional (Figure 18-17).

To apply a shadow, choose the Shadow section, and then, in the Presets list, choose one of the ready-made options. This list includes thumbnails for a wide range of different shadow effects with different degrees of fuzziness and shadow placement. If you're a hard-core artist, you can choose a different shadow color, and you can tweak all the other aspects of your shadow using the sliders that control transparency, shadow size, blurriness, distance from the shadowed chart element, and so on.

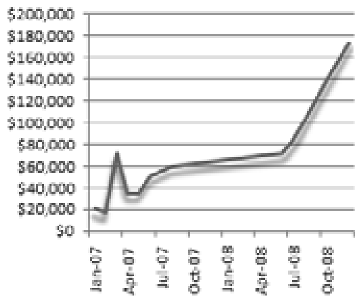

Figure 18-17:
Here, a shadow makes this line stand out from the chart background. Shadows work best with simple charts, like this one. If you had several different series on the same chart and gave each line its own shadow, the effect would be dizzying.

Formatting Data Series and Data Points

Adding labels is one way to distinguish important points on your chart. You can also use color, borders, and patterns. These techniques can't provide any additional information (like the value of the data point), but they're a great way to emphasize important information without cluttering up your chart. (Figure 18-18 shows two examples.)

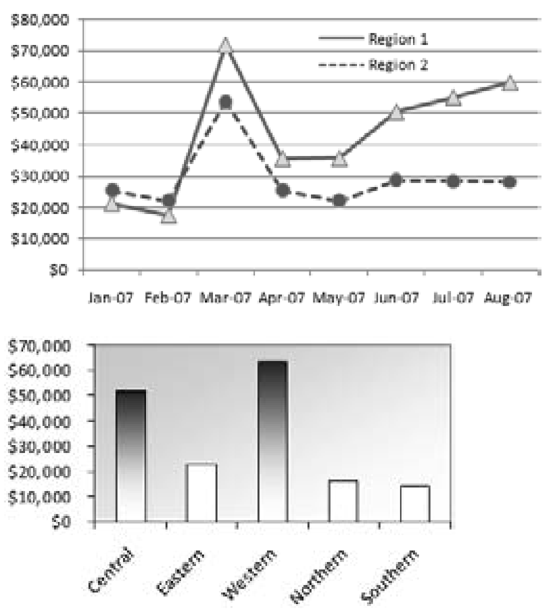

Figure 18-18:
These figures show two examples of formatting at work.

Top: Here's a line chart where the two lines are carefully distinguished from one another with different shaped markers and line styles. To get this effect, format each series separately.

Bottom: Here's a column chart where a few columns are emphasized with a gradient fill color. For this result, format data points individually.

You could have several reasons for formatting a data series or data point:

- You want to draw attention to specific data.
- You want to make sure that you distinguish between different series.
- You want to make sure your printout's legible, and that you can identify all the important information, even if your printout's black and white.

You already know the basic steps to format a data series because they're almost identical to the steps you use to format other parts of the chart. Start by selecting the area you want:

- If you select a data series, your changes affect all the data points in that series. Usually, you'll use this approach to help distinguish between different series.
- If you select a single data point (using the two-click technique explained in Figure 18-9), you affect only that one data point. Usually, you'll use this approach to highlight important values.

Tip: Your use of color and fills is limited only by your imagination, but excessive formatting can be distracting, so it's best to add extra flourishes only when they help you make a point. You could use different colors in a bar chart to help highlight the meaning of the results on a company's annual sales chart. Red-colored bars could represent losses, while black bars could show profits.

Then, use the familiar Chart Tools | Format→Current Selection→Format Selection command. You'll see a Format Data Series or Format Data Point dialog box that looks a lot like the Format Plot Area dialog box you started out with on page 509. Now, though, the formatting window includes a Format Series section where you can adjust some additional chart-specific details. You'll explore most of these options as you build better charts through the rest of this chapter.

Note: Use the Format Data Series or Format Data Point dialog box to implement any of these ideas.

Here are some formatting ideas:

- If you're using a column, bar, or area chart, you can adjust the fill colors and even use gradients and textures to make different series stand out. Go straight to the Fill section.
- If you're using a column or bar chart, you can adjust the width of the bars and the spacing in between. (In the hands of a crafty charter, this technique lets you build an overlay chart, as described on page 532.)
- If you're using a line or XY scatter chart, you can use the Line Color and Line Style section to change what the line looks like for each series.

- If you're using a line or XY scatter chart, you can use the Marker Options (Figure 18-19) and Marker Fill sections to choose the style (square, triangle, cross, and so on) and size of each data point that's placed on the line.

Tip: If you format a data point, and *then* format the series that contains that data point, the new formatting for the series takes over. Therefore, you need to reapply your data point formatting if you want a specific value to stand out from the crowd. To save time, you can use the helpful Redo feature to apply changes over and over again. First, format a data point the way you want it. Then, select a second data point, and press Ctrl+Y to reapply your formatting to the new data point. This technique can save you loads of time.

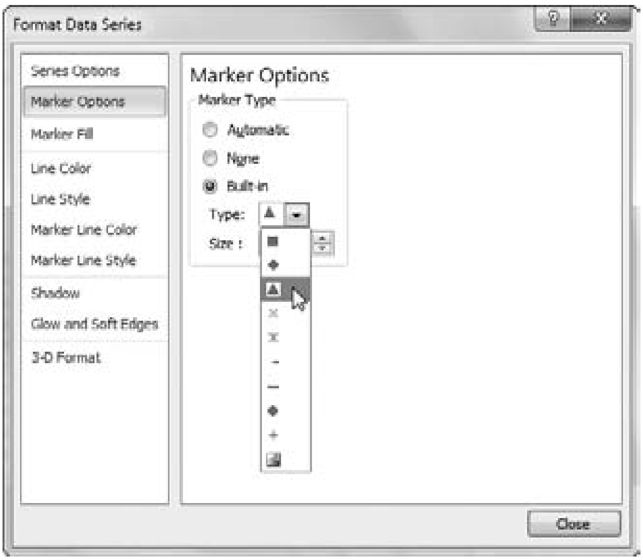

Figure 18-19:
It rarely makes sense to format an individual data point in a line or scatter chart. Instead, you'll usually format the whole series so that different lines stand out from one another. If you have more than two or three lines, you may need to get creative with different line styles and marker shapes to make sure that you really can tell them apart at a glance.

Reusing Your Favorite Charts with Templates

You can put a lot of work into creating the perfect chart. After you've slaved over your creation, it would be nice to have a way to reuse the formatting again in another workbook. Fortunately, Excel makes it possible through a template feature that lets you store your chart settings. Each chart template stores all the chart-formatting settings you've made, but none of the data.

Here's how it works. Once you've finished polishing up your chart, complete with all the formatting choices, choose Chart Tools | Design→Type→Save Template. Excel then prompts you to save a chart template file (which it identifies with the .crtx file extension). Choose a descriptive file name, like "Psychedelic Pie Chart," and then click Save.

By default, Excel offers to save the chart template in a Chart subfolder inside your personal template folder (page 455). Don't change this folder—the Chart Template folder's the only place Excel looks for templates, so if you place it somewhere else, you can't reuse it. (Of course, nothing's stopping you from *copying* the chart template file, perhaps to get it into the Charts folder on someone else's computer so they can benefit from all your hard work.)

To reuse your chart template, you need to pick it from the Create Chart or Change Chart Type dialog box. To create a new chart using your template, head to the ribbon's Insert→Charts section, and then click the dialog launcher (the icon of a square with an arrow in it) in the bottom-right corner. This shows the Create Chart dialog box, giving you a full selection of chart types. If you pick the Templates section at the top, you see all the chart templates that are stored in the Templates folder, complete with a mini preview (Figure 18-20).

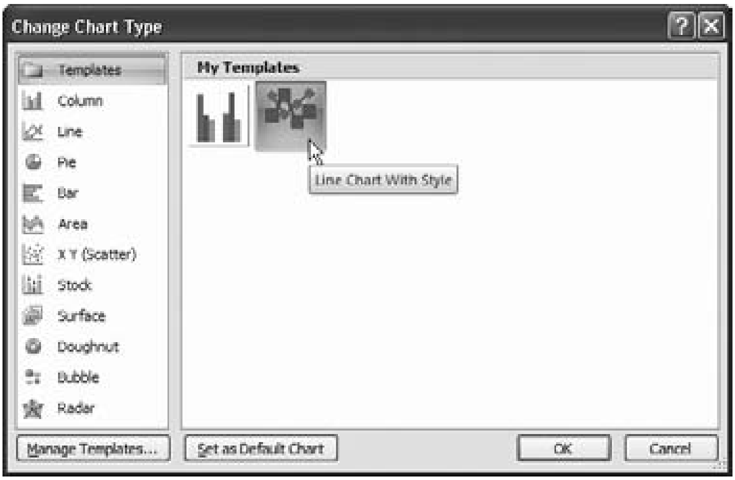

Figure 18-20: *This computer has two custom chart templates to choose from. You can hover over either one to see the descriptive file name (minus the .crtx file extension).*

When you select your template, and then click OK, Excel creates a new chart with the same formatting but using the data that's selected on your worksheet. Obviously, these options may not all apply to a new chart you create based on your template. Maybe your template includes formatting information for four series, but your new chart has only three. In this case, Excel just ignores any formatting information it isn't using.

Note: The formatting in the chart template is just a starting point. If you want to reuse some of the formatting but not all of it, you're free to use any of the formatting techniques you learned about in this chapter to further refine your new chart.

Improving Your Charts

So far, you've learned the key techniques you need to make sure your charts tell the right story. However, Excel lets you do plenty more, including adding trendlines, data tables, and error bars, and tweaking 3-D perspective and shapes. In the following sections, you'll learn even more about how to make the perfect chart.

Controlling a Chart's Scale

Many people don't think twice about the scale they use when they create a chart—instead, they let Excel set it automatically based on the values their chart has been built from. There's nothing wrong with this laissez-faire approach, but if you know how to take control of your chart's scale, you can make important data stand out and make it easier for people looking at your chart to spot relative differences in data and understand overall trends.

Usually, you'll be most interested in changing the scale of the value axis that runs on the left side of most charts. You can modify the scale of the value axis on most charts, including column charts, line charts, scatter charts, and area charts. (In a bar chart, the value axis actually runs horizontally along the bottom of the chart, although you can modify the scale in the same way as you do with these other chart types.) Pie and donut charts don't show a value scale at all.

Note: It's worth noting that quite a few unsavory individuals try to skew charts with crafty scale tricks. People often show two similar charts next to each other (for example, sales in 2006 and sales in 2008), and use a smaller scale in the second one to make it look like nothing's changed. Once you finish this section, you'll have a good idea how to spot these frauds. Some companies even have policies that enforce strict scale usage!

To change the scale, right-click the value axis, and then choose Format Axis. Or, if you find it hard to select the part of the chart you want, choose the value axis from the list in the ribbon's Chart Tools | Format→Current Selection section. Then, choose Chart Tools | Format→Current Selection→Format Selection.

When the Format Axis dialog box appears, choose the Axis Options section (shown in Figure 18-21). You have the choice of letting Excel automatically set the scale based on your data, entering the values you think are appropriate.

Note: When you set a scale value to Auto, Excel calculates it based on the current chart size and your current data. If you add more data, change the data values, or resize the chart (in which case there's more room to show intermediate values on the axis), Excel may modify the scale. But when you use Fixed, your numbers are hard-wired into the chart, and Excel never changes them (although you may, later).

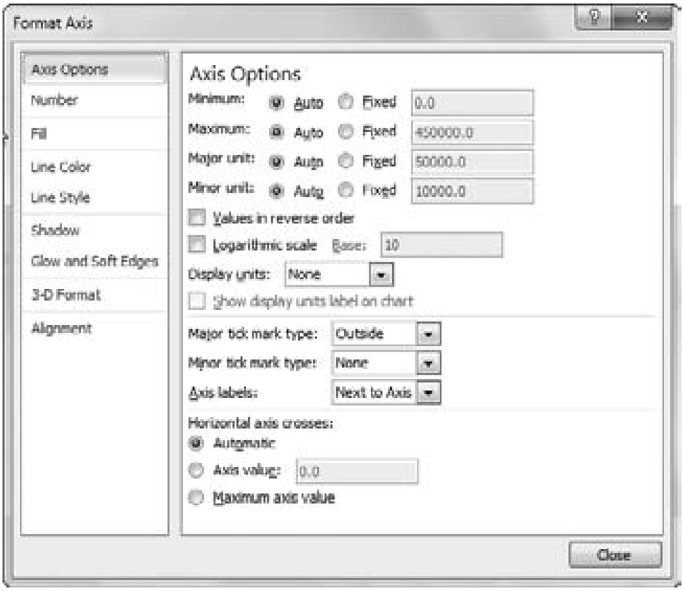

Figure 18-21:
The Format Axis dialog box lets you change the scale used on your chart. There's an "Auto" option next to each value. If selected (as in this example), Excel chooses the scale value (and the value it chooses appears in the text box on the right). In this example, the scale currently stretches from 0 to 180,000, with a major tick mark shown every 20,000 units. If you want to take control of the scale yourself, click Fixed instead of Auto (next to the value you want to change), and then edit the number in the corresponding text box.

Several settings determine the scale of your chart. These settings include:

- **Minimum** and **Maximum**. These values set the range of your scale. The axis starts at the minimum value (at the bottom of your chart), and ends at the maximum (at the top). Usually, Excel sets these values so that the minimum is 0 and the maximum is just a little bit above your largest data point. However, if your data points are very large and have only minor differences, you'll be able to help your chart's viewers make comparisons more easily if you compress the scale by having it start at a higher minimum value.
- **Major unit**. The major unit setting determines how many units the scale is divided into. If you have a scale from 0 to 1,000 and a major unit of 100, gridlines and axis labels appear every 100 units. Altogether, that makes for 11 labels.
- **Minor unit**. The minor unit setting determines how many *tick marks* are on the scale. (Tick marks are tiny lines added to the axis to help you judge the scale.) Usually the minor unit is less than the major unit. You may use a major unit of 100 and a minor unit of 10 in a chart that stretches from 0 to 1,000 (which generates 11 gridlines and 10 tick marks between each gridline). Five to ten major units on each axis, and five to ten minor units for every major unit, usually make for a readable chart.

Note: When you first create a standard chart, minor tick marks are turned off, so the minor unit setting doesn't have any effect. To set whether Excel shows major and minor tick marks, choose an option from the "Major tick mark type" and "Minor tick mark type" lists. Anything other than None does the trick. (The various options just determine exactly what the tick marks look like–for example, whether they're on the inside of the grid, the outside, or if they cross completely.)

Along with the settings just listed, you may also want to tweak the "Horizontal axis crosses" setting at the bottom of the dialog box. This number controls where the category axis line crosses the value axis. Usually, this line's right at the bottom of the chart, at the minimum value. However, you have two other choices. You can choose "Maximum axis value" to place the category axis at the top of the chart. The scale remains the same (meaning the minimum value's still at the bottom of the chart and the maximum value's at the top).

More interestingly, you can choose "Axis value", and then type in the exact value where the category axis should appear. This choice lets you put the axis somewhere in the middle of your chart. For example, you may want to plot a chart of test scores and draw the axis at a point that would indicate the minimum passing mark. Note that in a column chart, when a column has a value that's less than the axis, it points "downward," as you can see in Figure 18-22 (bottom).

Using these basic ingredients, you have a good deal of control over your chart's appearance. Figure 18-22 compares a few options that show how different scale choices can transform a chart, with the help of a little formatting.

The Format Axis dialog box also provides a few specialized options that aren't as commonly used but are still quite interesting. They include:

- **Values in reverse order**. This option turns your chart upside down. It places the category axis at the top of the chart (instead of the bottom), and changes the value scale so that it increases as it stretches down the side of the chart.
- **Logarithmic scale**. A logarithmic scale doesn't increase gradually. Instead, every major unit represents an increase by a power of 10. The values 0, 10, 20, 30, 40, 50 make up a typical scale. The values 0, 10, 100, 1000, 10000 make up a logarithmic scale. Logarithmic scales lend themselves to scientific and mathematical applications to help you see certain types of relationships and patterns in your data.
- **Display units**. You can use this option to shorten the text labels on your axis. This option's particularly useful if you have large dollar amounts. Imagine that your value axis stretches from $100,000 to $300,000. If you choose Thousands as your display unit, Excel adds the label Thousands along the value axis and changes the scale values to three digit numbers like $100 and $300. The bottom example in Figure 18-22 shows this space-saving trick at work.

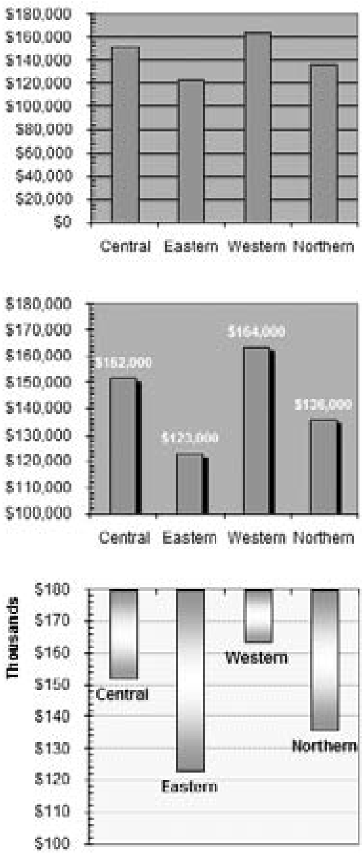

Figure 18-22:
These charts show the same data prepared three different ways.

Top: Here's a column chart the way Excel creates it, with the addition of minor tick marks in the value axis. This chart's problem is that the difference between the column values doesn't come across as very substantial–but, in fact, it is.

The middle and bottom examples solve this problem by setting the minimum value of the axis to $100,000, which shrinks the scale and emphasizes the differences between the columns' values.

Middle: This chart also includes data labels that show each column's value.

Bottom: This chart's data labels have been removed from the X-axis, and category labels have been added to the whole data series so the region's name shows up just below the bar. This example is further enhanced with gradient fills, dotted gridlines, and a vertical scale set to use thousands.

Note: If you're using a numeric or date-based category axis, you can format the scale of the category axis in the same way you format the scale of the value axis. You may want this option when you're creating an XY scatter chart or a line chart. If your category axis just displays labels, you can still format it, but you have fewer options. You can't change the scale, but you can reverse the order of categories, add tick marks, hide labels, and format the axis line's look.

Adding a Trendline

One of the main reasons that people create charts is to reveal patterns hidden inside the data. A gift card company may look at a historical record of sales to make an educated guess about the upcoming holiday season. Or a researcher might look at a set of scientific data to find out if potatoes really can cure the common cold. In both these examples, what's most important is spotting the trends that lurk inside most data collections.

One of the easiest ways to spot a trend is to add a *trendline* to your chart. A trendline's similar to an ordinary line in a line chart that connects all the data points in a series. The difference is that a trendline assumes the data isn't distributed in a perfectly uniform pattern. Instead of exactly connecting every point in a series, a trendline shows a line that best represents all the data on the graph, which means that minor exceptions, experimental error, and ordinary variances don't distract Excel from finding the overall pattern. Figure 18-23 shows an example.

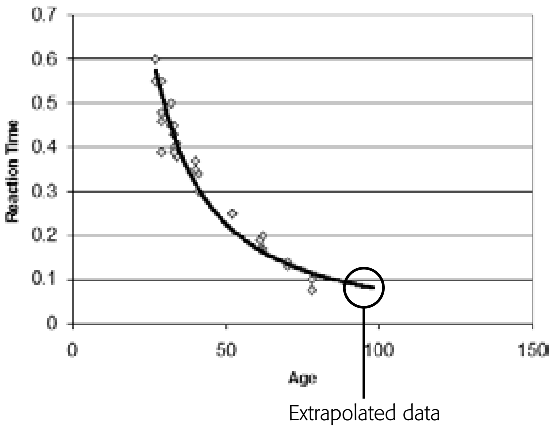

Figure 18-23:
This example shows the comparison between age and reaction time from the last chapter (Figure 17-21). The difference here is that the trendline has been extended beyond the actual data to predict results for people who are older than 80.

The other important point about trendlines is that they can predict values you don't have. The gift card company can use a trendline to get a good estimate of future sales, while a scientific experimenter can make educated guesses about data that wasn't recorded.

People most often use trendlines in XY scatter charts. Trendlines also makes sense in a column chart, and they can work in line, bar, and area charts in specialized circumstances.

To add a trendline, follow these steps:

1. **Select your chart.**

   If you're using an XY scatter chart (the most common choice), make sure you're using a subtype that doesn't draw lines. Otherwise, the trendline and the chart line may cause confusion. To change the subtype of a chart, right-click the chart, and then choose Change Chart Type.

2. **If your chart has more than one series, choose the series you want to use to create the trendline.**

   If your chart just has a single series, then you don't need to go to this trouble—just select the entire chart.

3. **Choose the type of relationship that you expect to find in your data from the Chart Tools | Layout→Analysis→Trendline list. To get even more types, choose Chart Tools | Layout→Analysis→More Trendline Options.**

   You can choose one of several types of trendlines:

   - **A Linear trend** varies regularly is the simplest possible relationship. If you're driving at a constant speed, the distance traveled increases linearly as the time increases.
   - **A Power and Exponential trend** starts off slowly and quickly ramps up. Logarithmic is the inverse—it starts off rapidly, and then levels out. These types of trends are more complex, but they're more common than linear trends in natural phenomena. Two example exponential relationships are the change in a population size with successive generations, and the distance traveled in a car if you continuously accelerate.
   - **A Linear Forecast trend** is similar to a Linear trend, but it extends a little bit further (the equivalent of two data points) to make a guess at where this relationship is going.
   - **A Polynomial trend** attempts to fit the data by creating an equation with a combination of different terms. This trend is the best choice if multiple factors are involved and your data doesn't fit a smooth line or curve.

   If there's more than one series in your chart, Excel asks you which one you want to use for the trendline. (You can repeat this process to add a trendline for each series.)

   After you add a trendline, you can right-click it, and then choose Format Trendline to change its line color, thickness, and style.

A standard trendline fits the data you have. However, you can also extend a trendline forward or backward to fill in values you don't know. This process of estimating data that you don't have (based on data that you do have) is *extrapolation*. A closely related concept is *interpolation*, which estimates unknown data values *between* known existing values. If a gift card company has information about sales from 2007 to 2010, you'd use extrapolation to predict sales in 2011. To make an educated guess at what sales were like in March 2008 (a month in which your firm lost its sales data), you'd use interpolation. All trendlines necessarily use interpolation, since there's always, in effect, "missing" data points between the data points you provide.

To extrapolate values in a trend, choose Chart Tools | Layout→Analysis→Trendline→More Trendline Options (if you're creating a trendline), or right-click the trendline, and then choose Format Trendline (if you want to extend the line after you've created it). Look for the Forecast box, and then specify the number of units that you want to add going forward or backward. Figure 18-23 shows an example.

Note: Don't always trust trendlines. It's quite possible that a relationship holds true only over a limited set of values. If you use a rising sales trendline as the basis for guessing future results, Excel's guesses don't, of course, take into account unexpected developments, like limited inventory or rising production costs. Similarly, if you extend the age versus reaction time comparison in Figure 18-23 too far, you'll wind up with ages and values that don't make sense (like a reaction time of 0 seconds, or a reaction time for a 300-year-old).

Adding Error Bars to Scientific Data

In a typical scientific experiment, you have two important sets of information: the actual results and an estimate that indicates how reliable these results are. This "reliability" number is the *uncertainty*. The uncertainty doesn't compensate for human error, faulty equipment, or invalid assumptions. Instead, it accounts for the limited accuracy of measurements. Think of the typical bathroom scale, which can give you your weight only to the nearest pound. That means there's an uncertainty of 0.5 pounds because any given measurement could be off by that amount. If a scientific experimenter weighs in at 150 pounds, he would record that measurement as 150±0.5 (150 pounds plus or minus 0.5 pounds). Any other calculations based on weight need to take this potential inaccuracy into account.

Because every type of measurement has a different range of accuracy, there's always a certain degree of imprecision that you need to watch out for before you make a dramatic conclusion. In a scientific chart, you can indicate the uncertainty using error bars, as in Figure 18-25. If you plot 150±0.5 on a chart, you should end up with a point at 150 and an error bar that stretches from the point up to 150.5 and down to 149.5.

To add scientific error bars to a chart, follow these steps:

1. **Click to select the appropriate data series on the chart.**

   If you have more than one data series, each series can have its own error bar settings.

2. **Make a selection from the Chart Tools | Layout→Analysis→Error Bars list.**

   You can use a standard choice to apply an automatic error bar setting (like one that assumes a five-percent error). If you find a good fit, your job is finished. But usually you'll need to choose exactly what you want by picking Chart Tools | Layout→Analysis→More Error Bars Options. Doing so shows the Format Error Bars dialog box (Figure 18-24).

3. **Choose how large your error bars should be.**

   Two simple and useful choices are "Fixed value" (which lets you specify the same measurement of uncertainty for every value), and Percentage (which lets you specify a percentage uncertainty). You can also supply a different fixed uncertainty for each value. In this case, add a new column with this information in your worksheet, choose the Custom option, and then click Specify Value to specify the range of cells with the uncertainty information (you need one uncertainty value for each data point). Statistics fans can also use two more-advanced options, Standard deviation and Standard error. For definitions of these two weighty concepts, consult your favorite statistics textbook.

4. **Click Close to add error bars to your chart.**

   Figure 18-25 shows a chart with error bars added.

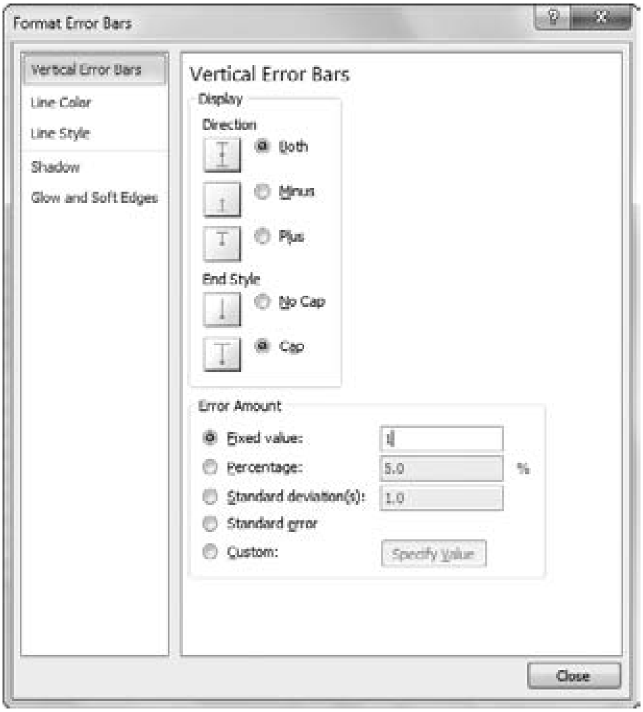

Figure 18-24: *You can set error bars to use a fixed value (as in this example), a fixed percentage of the value, or you can specify a range of cells in the worksheet that contains the uncertainty value for each measurement.*

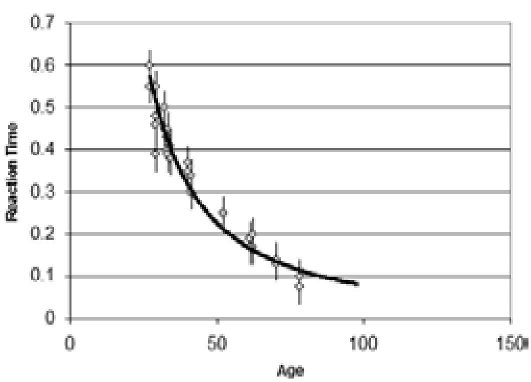

Figure 18-25: *This graph attempts to show a link between age and reaction time, with the error bars indicating that the reaction time couldn't be measured precisely. The error bars tell us that even though the trendline doesn't fit the data exactly, it lies within the range of measurement error. In this example, you can't distinguish all of the error bars, because sometimes the points and error bars overlap (consider the two data points around age 40). On the other hand, it's easy to see the single data point at age 50 and the error bars above and below it.*

Formatting 3-D Charts

As you learned in Chapter 17, many charts provide subtypes drawn in three dimensions. Some 3-D chart types are no different from their plainer 2-D relatives. In these charts, a 3-D effect simply gives the chart a more interesting appearance. But in true 3-D charts, it adds information by layering data from the front of the chart to the back, with each series appearing behind the other. (The column chart has seven subtypes. The first three are ordinary 2-D charts, the second three use a 3-D effect, and the last one's the only true 3-D chart. For more details about what subtype each chart supports, see the "Chart Types" section that starts on page 481.)

In true 3-D charts, you may need to take special care to make sure that data in the background doesn't become obstructed. A few tricks can help, like reordering the series and simplifying the chart so it isn't cluttered with extraneous data. You may also want to rotate or tilt the chart so that you have a different vantage point on the data.

To rotate or tilt a chart, follow these steps:

1. **Right-click the chart, and then select 3-D Rotation.**

   The Format Chart Area dialog box appears with the 3-D Rotation section selected, as shown in Figure 18-26.

Figure 18-26: *The 3-D Rotation section of the Format Chart Area dialog box lets you twist and turn your 3-D chart. As you make your changes, Excel updates the chart on your worksheet to give you a preview of what your settings will look like.*

2. **Use the arrow buttons to tilt the chart in various ways, changing the rotation and perspective.**

   You can also edit the individual settings by hand:

   - The X value lets you turn the chart from side to side.
   - The Y value lets you tilt the chart up and down.
   - The Perspective value lets you make the chart seem closer or farther away.

   You also have a few more-specialized settings you can use to fine-tune your chart:

   - Turn on the "Right angle axes" checkbox if you're tired of turning your chart this way and that, trying to find a good vantage point. When you switch this setting on, the Perspective value has no effect. Instead, Excel "straightens" out the chart by squaring it up, face-forward.
   - Turn off the Autoscale checkbox if you want to set the height of the chart using the Height box. The Height box specifies a percentage that's compared against the length of the bottom of the chart (the X-axis). In other words, a Height of 100 percent means the chart will be just as high as the base is wide. A height of 200 percent means the chart will be twice as high as it is wide. Keep in mind that Excel compares the height against the width of the base, not the width of the entire chart. If you rotate the chart, the perspective makes it appear as though the base isn't as wide as it really is. In this situation, Excel still makes the chart just as tall as if you hadn't rotated it.
   - Use the Depth box to set how deep the chart is (from front to back). As with the Height box, you use a percentage that's compared to the base of the chart. That means a depth of 200 percent makes a chart twice as deep as the base is wide. Sometimes, a deeper chart creates a more dramatic 3-D effect.
   - Click the Default Rotation button to return your chart to the just-slightly tilted way it began life.

3. **Click Close when you're finished.**

Changing the Shape of a 3-D Column

Excel provides several subtypes of column and bar charts that use a 3-D effect. Excel also provides the Cylinder, Cone, and Pyramid chart types, which are always three-dimensional (see the description of chart types, beginning in "Column", for more information). It doesn't really make much difference whether you use columns, pyramids, or cones—the overall effect is pretty much the same. However, you create a much more dramatic effect by putting more than one shape in a single chart. If you want to compare two series, then you could represent one with columns and the other with cones, as shown in Figure 18-27.

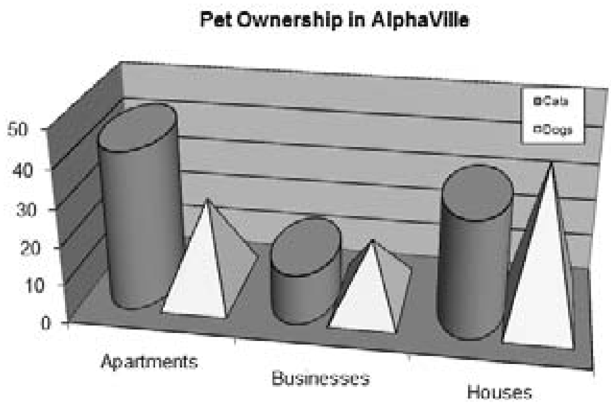

Figure 18-27:
To make a dramatic point, you can combine more than one shape in a chart. In this example, cylinders represent the cat population and pyramids represent the corresponding dog figures. Creating this chart is refreshingly easy–you just need to change the standard choice in the Shape tab of the Format Data Series dialog box.

To try this out, follow these steps:

1. **Create a column chart with two or more data series.**

   Make sure you pick one of the chart subtypes that use a 3-D effect.

2. **Right-click one of the data series, and then choose Format Data Series.**

   The Format Data Series dialog box launches.

3. **Choose the Shape section. Click one of the pictures to change your columns into cylinders, cones, pyramids, or something else. Then click Close.**

   Repeat steps 2 and 3 for any other data series that you want to change.

Advanced Charting

At this point, you've seen a wide range of tips and tricks for improving almost any Excel chart. But it's worth learning a few advanced charting techniques. Most of these work only with specific chart types, like exploding pie slices and overlaying bar charts. You'll also learn about one of the most interesting chart tricks known to Excel-kind: using a combination chart to fuse together two different chart types into one object.

Exploding Slices in a Pie

Data labeling and formatting help make individual slices of a pie chart stand out. However, to really accentuate important information in a pie or donut chart, you can separate a piece from the rest of the pie by *exploding* the slice (to use the rather dramatic technical term). You've already seen this feature at work with the pie or donut chart subtypes that explode *all* the pieces. This technique shows you how to explode just a single slice, as shown in Figure 18-28.

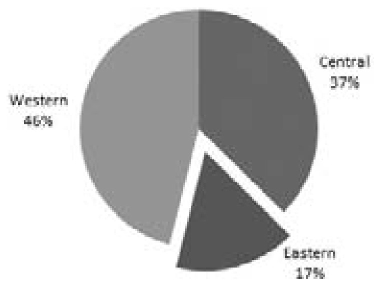

Figure 18-28:
This example shows a pie with one exploded piece, representing the sales in the Eastern division.

To explode only one slice of a pie chart, just click the slice, and then drag it away from the pie. You can pull it as far away as you want, and you can repeat this process to explode several pieces (although the emphasis works best when you remove just a single slice). Be careful that you don't select the whole data series before you start dragging, or you'll end up exploding the whole pie.

Tip: You can also help separate your slices by adding a thin white border around each slice. Right-click any slice, and then choose Format Data Series. Then, in the Border Color section, choose to create a solid white line. You can use the options in the Border Styles section to make the line thicker so the slices stand out more.

Grouping Slices in a Pie

An even more interesting pie chart feature breaks down a single slice into more detail. You may want to create a pie chart that shows your personal budget by category (including food, living expenses, clothes, and so on). You could then further subdivide a single slice, like food, to show what portion goes to groceries, restaurants, the local hot dog vendor, and ice cream splurges. Adding all this information into a single pie would result in too many slices, making the chart less effective.

Creating a pie chart with a single slice breakdown is pretty easy, but it's not the most intuitive operation in the world. First, you need to set up your data correctly. Your table needs to include *all* the information you're putting into the chart in a single series. Imagine you want to create a pie with these slices:

Food	$13,911
House	$18,342
Clothes	$4,790
Fun	$7,980
Computer	$34,010

Now, assume that the slice you want to subdivide is the Computer slice. That means you need to remove the Computer entry and replace this information in the chart with more detailed information. Here's what you might end up with:

Food	$13,911
House	$18,342
Clothes	$4,790
Fun	$7,980
Computer Software	$7,500
Computer Hardware	$6,500
Missing Manual Books	$20,010

Notice that the Computer Software, Computer Hardware, and Missing Manual Books entries add up to the whole computer slice ($34,010), which has been removed from the chart.

Note: The only limitation to breaking down an individual slice is that you can perform this trick on only one slice.

Now that you've organized your data in the right way, create a new pie chart, and choose either the "Pie of Pie" or "Bar of Pie" subtypes. Excel automatically chooses some slices from the bottom of your chart to group in a separate pie or bar chart. Unless Excel makes a lucky guess, this selection doesn't include the same slices that you want. Fortunately, it's easy to change the separation just by dragging.

If you have a "Bar of Pie" chart, you can add a slice to the bar by dragging it from the main pie to the column. If you want to take a slice that's in the column and put it back in the pie, drag it from the column back to the main pie. Continue this process until you've grouped the Computer Software, Computer Hardware, and Missing Manual Books slices into the standalone bar.

Figure 18-29 shows a perfected "Pie of Pie" chart.

Gaps, Widths, and Overlays in a Column Chart

Column and bar charts have their own options for fine-tuning. You can adjust how far each column group's spaced from the next and how much space appears between each column in a group. With a little imagination, you can use these settings to create an innovative *overlay chart*, which layers two different sets of data on top of each other.

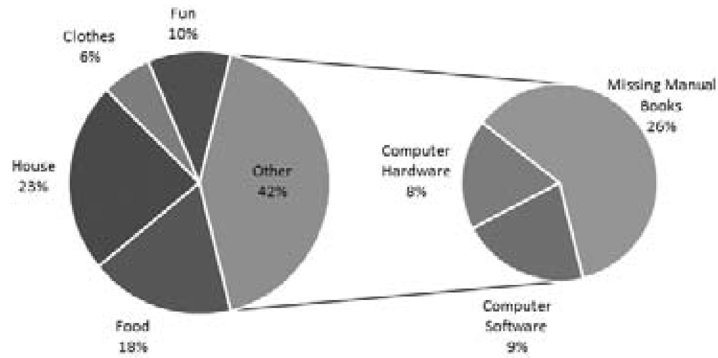

Figure 18-29: *This example shows a pie chart of expenses, with the Other slice broken down further into a second pie chart. Excel pie charts often need a fair bit of tweaking before they look respectable. Usually, you'll want to remove the legend and use category labels, as shown here. If you drag the label a short distance away from the midpoint of the slice, Excel adds a line pointing from the label to the slice. These lines often help to make a pie chart more readable.*

To see these extended column chart settings, right-click a data series in the chart, and then select Format Data Series. Then, select Series Options, which is the first section in the list. You get two settings:

- **Series Overlap** is the amount of overlap between columns in the same category. This setting takes effect only if you have more than one series. The standard option is 0, which means that clustered columns touch each other but don't overlap. With a value of 100, the columns in the same category overlap completely, while negative values put a space between the columns.
- **Gap Width** is the space between each category. In a chart with one series, this setting is the space between each column; in a chart with multiple series, it's the space between each group of columns. The standard value is 150, which leaves a space that's roughly equal to the width of 1.5 columns.

Note: In a 3-D column chart, you'll find an extra setting: Gap Depth. This is the space between the columns at the front and the columns that appear behind them.

One interesting trick is to use the overlap setting to compare two sets of data. People often use this technique to compare projections against actual results. Here are the steps you need to follow:

1. **Create a bar or column chart with two series.**

   Figure 18-30 shows one possible example of a bar chart that compares projected sales against actual sales.

2. **Decide which series should be on top of the other.**

   The series that's on top needs to be transparent, so the other series can show through. In Figure 18-30, the projected value should be placed on top.

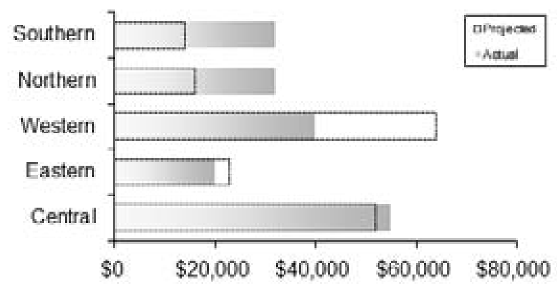

Figure 18-30: *This example shows an overlay bar chart. The projected sales values have dotted borders and a transparent fill. The actual sales numbers show a gradient fill. Using this model, you can clearly see that the Southern and Northern locations exceeded their expectations, while the Western region fell far short of its predictions, even though it had the second-best overall sales.*

3. **Find this topmost series in the chart. Right-click it, and then choose Format Data Series.**

   The Format Data Series dialog box launches.

4. **Select the Series Options section, and then set the Series Overlap to 100.**

   This setting ensures that two series are layered one on top of the other.

5. **Select the Fill section, and then choose "No fill".**

   This tells Excel to make this data series transparent, with only the border visible.

6. **Click Close.**

   Next, you want to make sure the right series is on top.

7. **Right-click the chart, and then choose Select Data. Make sure the series that's supposed to be on top is at the bottom of the list.**

   To move a series, select it, and then click the down arrow button. Excel draws the series from top to bottom in this list, so the one it draws last ends up in front.

8. **Now select the other series, right-click it, and then choose Format Data Series. Select the Border Color section, and then choose "No line".**

   You can also choose any fill color or fill effect. You may need to try several combinations in order to have enough contrast for the overlay chart to really work.

9. **Click Close.**

Creating Combination Charts

Sometimes, you want to use a chart to compare trends in different but related data. Imagine you create two charts, one to show how many hours you've worked in the last few months, and the other to show how much money you've spent. After you create these two charts, you start to wonder if there's a link—in other words, do you spend more money when you have a greater workload? Unfortunately, because these two measurements use different units (one records the number of hours, while the other counts the number of dollars), you can't put them on the same chart—or can you?

Combination charts are a well-kept Excel secret. Using combination charts, you can compare trends across different sets of data, even if the units are wildly different. Combination charts are also useful when you need to compare more than one piece of information to tell the complete story. A chart that shows quarterly sales could make it look like your company is meeting wild success. But using a combination chart, you can contrast the sales against another factor (like changing currency exchange rates, or increased business expenses brought about by an irrational 80-percent-off coupon campaign) to tell a different tale.

Tip: There's a screencast—an animated online tutorial—that shows the process of creating a combination chart, online at *www.missingmanuals.com/cds*. It was created using Excel 2007, but the process is the same for Excel 2010. Surf to the Missing CD page for *Excel 2007: The Missing Manual* to check it out.

You make a combination chart by creating a new secondary value axis. This axis appears on the right side of the chart. In other words, you'll end up with two value axes, one that applies to the first series of data, and the other that applies to the second series. Figure 18-31 shows how this process works.

To create a combination chart, follow these steps:

1. **Create a chart that includes both sets of data.**

   Initially, the smaller set of data doesn't appear correctly, because the scale is wrong. To solve this problem, you need to plot this data against another axis.

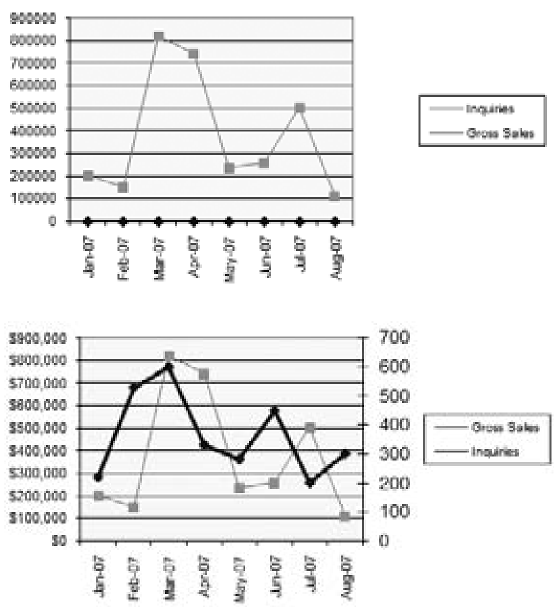

Figure 18-31:
These charts both compare the number of customer inquiries with the amount of sales made in each month.

Top: Both lines use the same scale, which actually doesn't make much sense because the units tracked in each data series are different.

Bottom: But once a secondary axis is added at the right side, the numbers are a bit easier to compare. For an even better combo chart solution, see Figure 18-32.

2. **Right-click the data series that doesn't fit, and then choose Format Data Series.**

   If the scale difference is really extreme, you may have trouble selecting the data series you want to change, because it may be squashed against the bottom chart axis. In this case, select the series by name from the Chart Tools | Format→Current Selection list.

3. **Select the Series Option tab, choose Secondary Axis under the heading Plot Series On, and then click Close.**

   Excel creates a new scale on the right side of your chart and uses this scale to plot the data series. Excel automatically chooses the best scale, although you're free to change it by formatting the axis.

Combination charts don't just let you compare different units of data. They also help you fuse together two different types of charts. You could plot one series using a scatter chart, and then add columns to represent the values in the second series. When done right, this combo produces an attractive chart with a series of columns and a line above them, which you see commonly in shareholder reports and marketing documents.

Tip: Some chart types just don't mix. To mix properly, the arrangement of axes should be the same in both chart types. Thus, a combination of a column chart and line chart works great, but a combined line chart and pie chart doesn't make much sense.

To create a combination chart that uses more than one chart type, right-click the *series*, and then choose Change Series Chart Type. You can now choose a new chart type and subtype, as usual. The difference is that this change is applied only to the series that you've selected. Figure 18-32 shows a good example.

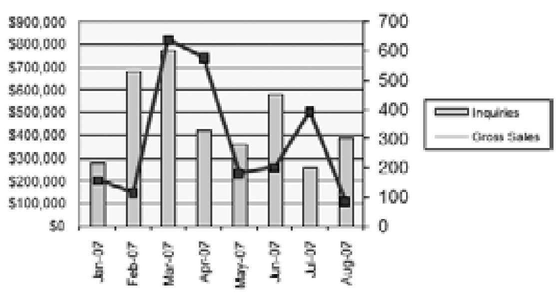

Figure 18-32:
This example compares the inquiry and sales numbers using two different chart types. The inquiry numbers are depicted as columns, and the sales figures use a line chart. This technique is a great way to highlight the fact that you're comparing different types of data.

CHAPTER

19

Inserting Graphics

Most Excel fans concentrate on numbers, formulas, and charts when they create worksheets. But Excel has another dimension—graphics. In fact, Excel includes a slew of drawing tools that may seem like they belong more in an artsy illustration program than the rigid confines of a spreadsheet. Some of these drawing features are shameless frills that just take up space on the Excel ribbon. Others are genuinely useful, letting you add useful touches and highlight important information with real pizzazz.

Excel's drawing capabilities fall into the following four categories:

- **Picture Handling**. These features are what most people think of when they think about graphics. Picture handling lets you take the picture files that you have on your computer and insert them into an Excel worksheet. For instance, if you've created the perfect company logo in another program, you may want to place it in a blank spot on your worksheet.
- **Clip Art**. Clip art graphics are usually cartoonish, themed pictures (like a stack of dollar bills or a drawing of a globe). While you probably don't need this stuff in most worksheets, Excel's clip art features are still quite impressive. Instead of limiting you to a small selection of preinstalled image files, Excel lets you search an online collection of thousands of images using subject keywords.
- **Shape Drawing (and fancy text)**. Excel's shape drawing tools let you create images directly on your worksheet. These shapes include arrows, circles, stars, banners, floating text boxes, and captions with zany gradients and 3-D effects. Shape drawing is occasionally useful when you want to highlight information on your worksheet, but it's most practical when you use it with charts.

- **SmartArt**. SmartArt lets you add complete diagrams that are made up of a combination of shapes and text. You can use SmartArt to create organizational charts, simple lists and flow charts, and pyramid-style diagrams. Under the hood, SmartArt uses the shape drawing features to create its diagrams, which means you can use all the same graphical flourishes and fancy effects.

This chapter covers all four kinds of drawing in Excel.

Adding Pictures to a Worksheet

In the previous two chapters on charts, you learned how Excel places charts in special floating boxes that hover above your worksheets. Pictures work in a similar way—they're distinct, floating objects that you can place anywhere. And just like charts, a picture box may hide data underneath it, but it'll never disturb the data.

Common examples of graphics in worksheets include minor embellishments, like a company logo next to the title, or an exclamation mark icon that highlights a worst-case scenario. You shouldn't go overboard with pictures because they tend to clutter up the real data. A few careful touches, however, can go a long way to making your spreadsheet more readable and more memorable.

Inserting a Picture

To insert a picture file that exists on your computer, follow these steps:

1. **Choose Insert→Illustrations→Picture.**

   The Insert Picture dialog box appears (Figure 19-1). It's essentially the same as the trusty Open dialog box you use to open old workbooks, except the "Files of type" list is set to show all the types of image files you can use in Excel. You can change this option to show only the file type that you're interested in (for example, bitmap files, JPEG files, GIF files, or something else).

2. **Browse to the picture you want to insert, select it, and then click OK.**

   Excel lets you use a wide range of image file formats, including files with the following extensions: .bmp, .gif, .jpeg, .tif, .png, .wmf, and .emf.

Note: Excel includes an odd feature that lets you insert a snapshot of a currently running program instead of a picture file. To try this out, choose Insert→Illustrations→Screenshot, and then pick the window you want from the list of thumbnails. This quirky feature is interesting, but few people are likely to use it–and those who really need screen captures probably won't be satisfied with the way Excel rudely chops off the edges and corner of the window frame. (For a more professional result, you can use dedicated screen capture software like WinSnap–we do.)

   When you insert the picture, Excel places it in a new floating box. Figure 19-2 shows the result.

Figure 19-1: *Depending on what mix of content exists in the current folder, the Insert Picture dialog box may or may not show thumbnail previews for each picture it finds. If you don't see the previews but you'd like to have them, click the Views button (circled), and then choose Tiles.*

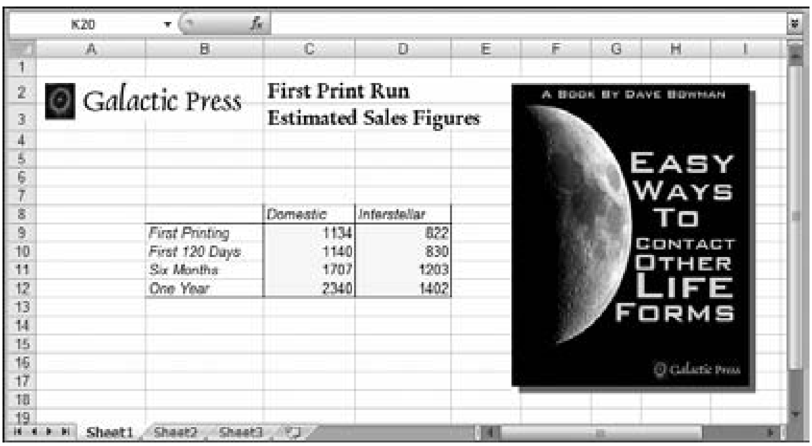

Figure 19-2: *This worksheet has two picture objects: a logo in the top-left corner, and a book cover image on the right. The picture boxes that Excel creates are similar to chart boxes, except Excel doesn't draw a border around a picture unless you insist. Whenever you select the picture, the ribbon sports a new tab named Picture Tools | Format.*

Note: When you insert a picture this way, Excel copies all the picture data into your worksheet. Even if you change or delete the original picture file, it doesn't have any effect on your worksheet.

FREQUENTLY ASKED QUESTIONS

Inserting Really Big Pictures

I want to insert a few really big pictures that I took with my digital camera. But I don't want to have a workbook that's 30 MB!

Inserting big pictures into a workbook is often a bad idea. If you do so, you'll have trouble emailing your workbook to other people or just moving it around. Sometimes you can slim down pictures using the Compress Pictures command described on page 545. But slimming down often isn't enough.

Fortunately, Excel has another picture trick up its sleeve. Rather than inserting a full picture, you can insert a *link* to your picture file. This way, your Excel workbook file stays small (and if anyone changes the picture, you'll see the difference the next time you open your workbook). The obvious disadvantage is that if you move, rename, or delete the image file, it disappears from your spreadsheet. (It's generally a good idea to keep the picture in the same folder as your Excel workbook.)

To insert a picture link, choose Insert→Illustrations→Picture, and then browse to the picture file, just as you normally would. But instead of clicking the Insert button, click the down-pointing arrow on the Insert button's right side. This click pops open a menu with extra choices. Choose Link to File to add a linked picture.

You have another option. You can choose Insert and Link to insert a picture normally *and* add the link to the original file. This way, if the picture changes, your workbook reflects the new look, but if the picture file disappears (say you mistakenly delete it), your workbook still has its own internal copy. The drawback of Insert and Link is that your workbook still stores a copy of the picture file, so it's still just as large as if you inserted the picture in the usual way.

Positioning and Resizing a Picture

The easiest way to move a picture once you've inserted it is to click anywhere on the picture surface, and then drag it to a new location.

Tip: Sometimes you want to move your picture just a small amount so that it lines up perfectly with some other part of your worksheet. To get really fine-grained control over picture positioning, click the picture to select it, and then use the arrow keys. For example, if you press the left arrow key, Excel nudges the picture ever so slightly to the left (one pixel, to be exact).

You can nearly as easily change a picture's size. First, click the picture once so that the *resizing handles* appear. These handles look like small circles and squares, and they show up at each corner and in the middle of each side. Drag one of these handles to one side to expand or shrink the picture.

Depending on which handle you drag, the type of resizing Excel performs changes:

- Use the squares that appear in the middle of each side to stretch the picture in one direction (possibly stretching it beyond all recognition at the same time).
- Use the circles that appear in each corner to resize the picture bigger or smaller without changing its proportions. See the box below for an explanation of the difference.

UP TO SPEED

Resizing Without Distortion

Excel doesn't impose any limits when it comes to resizing pictures, and, if you're not careful, you can completely mangle your image. Watch out for two problems when resizing a typical image.

First, you need to resist the urge to expand or compress your image dramatically. When you enlarge an image, Excel needs to use *interpolation* to guess the information it should add. (Interpolation is the process by which Excel looks at the existing pixels in your picture, and then uses them to calculate extra pixels that it should add between. If Excel finds a blue dot next to a yellow dot, then it may add a blue-yellow dot in between.) If you expand an image too much, Excel needs to make far too many guesses, and you end up with a poor-quality image (usually the edges of lines and shapes appear blocky or jagged).

Similarly, when you shrink an image, Excel needs to decide what information to discard. In order to minimize the damage, Excel tries to smooth out the new picture. But if you shrink an image too much, you end up with a picture that looks blurry or fuzzy.

To avoid either of these problems, take a careful look at your worksheet (and print it out) after you make your changes to make sure the images remain acceptable.

The second issue to be aware of is *aspect ratio*—the ratio of an image's width to its height. A company logo may be twice as wide as it is tall, giving it an aspect ratio of 2:1. When you resize the picture, you need to keep this sense of proportion in mind. If you change the height of the logo without adjusting its width correspondingly, the image becomes distorted. To avoid these problems, just use the resizing circles at the *corners* of the image (rather than the squares that appear in the middle of each side). These resizing handles let you change the size of the image without altering its aspect ratio. Instead, the height and width change in lockstep, keeping the right proportions.

If your image is what's known as a *vector graphic*—created by a program like Adobe Illustrator that uses formulas rather than bit-by-bit information to draw its pictures—you're in luck. You can resize vector images without causing any distortion. Excel's clipart and shape drawing features (both of which you'll use later in this chapter) use fully resizable vector graphics.

If you're not too handy with a mouse, or if you just want to size your picture with exact precision (possibly because your worksheet contains several pictures and you want them all to be consistent), head to the ribbon's Picture Tools | Format→Size section. You'll see two text boxes, which provide the current height and width of the picture box. (The units depend on the computer, but inches and centimeters are two common possibilities.) You can click either of these text boxes, and then type in a new value by hand.

Note: Excel gives you one more way to set the size of a picture by hand. From the Picture Tools | Format→Size section, you can click the dialog launcher (the tiny arrow-in-a-square icon at the bottom-right corner). This opens the Size and Properties window and shows the Size tab. The Size tab not only lets you set the size (using exact measurements or percentages), it also lets you fill in some cropping settings (page 547), and it lets you rotate the picture by entering an angle in the Rotation box. (Try 180° to flip a picture on its head.)

Pictures, like charts, are anchored to specific cells. Excel may move and resize the picture when you insert, remove, or resize rows and columns. The difference between picture boxes and chart boxes is that you can explicitly control how Excel moves and sizes the picture, to make sure it isn't inadvertently moved when you want it to remain firmly in one spot.

To change the *way* a picture's positioned, follow these steps:

1. **Select the picture.**

   The Picture Tools | Format tab appears.

2. **In the ribbon's Picture Tools | Format→Size section, click the dialog box launcher (the tiny arrow-in-a-rectangle icon).**

   The Size and Properties dialog box appears.

3. **Choose the Properties tab.**

4. **Under the "Object positioning" section, choose one of the options.**

   The "Move and size with cells" option gives pictures the same behavior as charts. If you insert new rows above the picture, the whole picture shifts down. If you insert rows under a picture, the bottom edge stretches down. This behavior usually isn't what you want because stretching a picture could distort it.

   The "Move but don't size with cells" option anchors the top-left corner of the picture. If you insert new rows above the picture, it shifts down. If you insert rows under a picture, Excel doesn't stretch the picture. When you first insert a new picture, Excel uses this option—which makes the most sense if you have a picture you want to position near some related data (but you don't want the picture dimensions to get mangled when you add or remove content).

   The "Don't move or size with cells" option doesn't anchor the picture at all. In this case, Excel doesn't move the picture or resize it no matter where you insert or delete rows. This option makes the most sense if you want a picture to stay put, regardless of wherever the content moves. This choice is typically a good one for a graphical header or company logo.

Tip: The Properties tab has another useful option: You can turn off the "Print object" checkbox in order to tell Excel to leave the picture out of your printouts. This option makes sense if you want to include rich graphics that don't look right on your black-and-white printer (or just waste too much ink).

5. **Click OK.**

   Ta da! Marvel at the glory of your picture.

POWER USERS' CLINIC

Transfer Pictures Quickly with Copy and Paste

You can also insert a picture by copying it from within another program, and then pasting it into Excel. After you've copied the image in the program where you're viewing it, select Home→Clipboard→Paste→Paste Special in Excel. When you use the Paste Special command, a Paste Special dialog box appears, with a list of different choices. Choose the option that reflects the image format ("Bitmap" if you're pasting a bitmap file), and then click OK.

You must use Paste Special instead of Paste so that you insert the right type of content. Depending on the program you're copying from, Excel may paste the picture as a bitmap or as a *linked* or *embedded object* (page 777). If you use the Paste command, Excel decides which option it thinks is best. If you use Paste Special, you get to decide.

Either way, the picture looks the same in your spreadsheet. The difference is what happens when you select the picture. If it's an ordinary picture, you can manipulate it using Excel's picture tools. If it's a linked object, you can double-click it to edit it with the program that created it. But if you open the worksheet on a computer that doesn't have the required program, you can't edit it. One quick way to tell whether you have an object or a picture is to right-click the object. If you see the option Format Picture, you've selected a picture. If you see the command Format Object instead, you've selected an object.

So which choice is best? It you don't intend to change an image, it's always best to paste it as picture data. This choice also ensures you can share your workbook files with other people without any complications. On the other hand, if you decide that you absolutely need the ability to modify the image using the original program, you can paste a linked object instead. Head straight to Chapter 27 to learn how linked objects work and how to manipulate them.

Picture Touch-Up

Once you get your picture into Excel, you may decide you want to polish it up by changing colors or applying special effects. Depending on the result you're after, you could use a dedicated graphics program (and if you want more features than Excel provides, that's the best choice). But you'll probably be surprised to see Excel's sophisticated built-in picture-tweaking features. The ribbon's Picture Tools | Format tab, which appears whenever you select a picture, is the starting point for these features.

So what can you do with the Picture Tools | Format tab? First, you should explore the Adjust section, which lets you adjust colors, contrast, and other details. Here are the buttons you can use:

- **Brightness** lets you increase or decrease the overall brightness of your image. As you increase brightness, all colors get brighter. As you decrease brightness, all colors get darker. As a side effect, increasing or decreasing brightness often reduces the contrast. Brightness is always adjusted in percentages, so –30% makes the image 30% darker, and +30% makes it that much lighter. If you aren't happy picking one of the standard percentages from the list (which go in increments of 10 from –40% to 40%), you can choose Picture Correction Options (at the bottom of the list) to pop open a dialog box where you can enter an exact percentage.

- **Contrast** lets you increase or decrease the contrast. The contrast is a measure of how much difference exists between the brightest and darkest colors. As you increase contrast, bright colors get brighter and dark colors get darker. As you decrease contrast, all colors start to converge toward a middle-of-the-road gray. As with brightness, contrast is set in percentages. You can choose one of the ready-made percentages (which go in increments of 10 from –40% to 40%), or you can choose Picture Correction Options to pop open a dialog box where you can enter an exact percentage.

Tip: Choose a 0% brightness or contrast to remove your previous brightness or contrast setting and return the image to normal.

- **Recolor** lets you adjust the color in the selected picture. When you click this button, you'll see a gallery with thumbnails showing different possibilities. At the top of the list, under the Color Modes heading, you see four preset options: Grayscale, which changes every color to a shade of gray (and gives you a good idea of what a picture will look like on a black-and-white printer); Sepia, which does much the same thing as Grayscale but adds a slight brown tone reminiscent of old photographs; Washout, which fades the picture colors (and helps save ink when you print the worksheet); and Black and White, which changes every color to either black or white (a process that ruins all but the simplest of pictures).

  Underneath these basic options are other options that tint your picture using one of the accent colors from the current theme (page 174). Use the choices under Dark Variations to apply a dark tint of a given color, or use the choices under Light Variations to lighten the picture as you apply the new tone. If none of these options float your boat, take a look at the More Dark Variations command at the bottom of the list, which lets you pick any color you want.

- **Recolor→Set Transparent Color** lets you make certain portions of an image transparent. If there's worksheet data underneath the transparent regions, it shows through. When you click Set Transparent Color, the mouse pointer changes into a crosshairs symbol, which looks like a small plus sign (+). Next, click the color in the image that should become transparent. If you want a white background to be invisible, click a white portion of the image. Figure 19-3 shows an example.

  Usually, you'll use Set Transparent Color to make a background transparent. You may want a company logo to blend right into the background color you've defined on a cell. In this case, when you create the company logo picture in your graphics program, you should make sure that you choose a color for its background that's not used anywhere else in the logo. That way, when you make that color transparent, it doesn't affect any other portion of the picture.

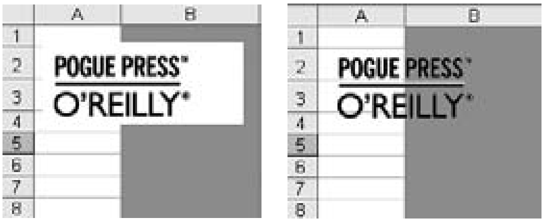

Figure 19-3:
Using the Set Transparent Color feature, you can make any color in an image become transparent. Here are two versions of a graphic, one with no transparent color (left), and one where the background white color is transparent (right). As you can see, in the transparent version, the shaded content of the cell clearly shows through. If there were any data in those cells, it would also show through.

- **Compress Pictures** lets you shrink your workbook file by discarding extra picture details you don't need. When you click this button, the Compress Pictures dialog box appears, with several options for reducing the data size of your pictures. The next section describes these options in more detail.
- **Change Picture** replaces the currently selected picture by popping a different image into the current picture box. When you click this button, Excel opens the Insert Picture dialog box (covered on page 539) so you can choose the new picture file that you want to use. Any picture changes you made to the previous image (like changing the contrast or colors) are lost.
- **Reset Picture** restores the picture to its original form. The picture looks the exact way it did when you first imported it, and Excel discards any color changes, cropping (page 547), and other modifications you've made. Excel also returns the picture to its natural size, although it stays in its current location.

Note: The modifications you make to a picture affect only how it *appears* on the worksheet (and in your printouts). Behind the scenes, Excel actually stores the full-size original picture data (unless you've used the Compress Pictures command). If at any point you realize that you've applied an edit you don't want, you can revert to the original version of the picture by clicking Reset Picture. However, this action resets *all* the changes you've made. You can't roll back just a single change this way (instead, use Ctrl+Z to undo a change right after you make it).

Compressing Pictures

Pictures increase the size of a spreadsheet file, and if you create a worksheet with dozens of graphics, the file's size can grow significantly. Most of the time, you won't worry too much about the size of your Excel files. However, if you plan to send it through email or put it on an old-fashioned diskette, then you may need to pare it down in size. You can cut down the picture data, for instance. In a spreadsheet with a number of high-quality bitmap pictures, the images can take up a significant amount of disk space. (On the other hand, vector drawings like clip art and shapes don't use much space at all.)

Note: Don't compress pictures if you want to change them later. Why? Because compressing a picture discards the original picture information. If you shrink a picture, compress it, and then enlarge the picture back to its original size, you end up with a lower-quality image.

To compress a picture, follow these steps:

1. **Select the picture in your worksheet that you want to compress.**

   If you want to compress more than one picture at once, hold down Ctrl while you click each picture. If you want to compress all the pictures in your file, the Compress Pictures dialog box (in the next step) has a shortcut you can use, so just select one picture for now.

2. **Choose Picture Tools | Format→Adjust→Compress Pictures.**

   The Compress Pictures dialog box appears (see Figure 19-4).

Figure 19-4:
The Compress Pictures dialog box lets you choose just how compressed your pictures will be. The Print setting (used here) keeps 220 pixels for every square inch of print space.

3. **If you want to change all the pictures in your workbook, clear the check box next to the "Apply only to this picture" setting.**

   This option is useful in a workbook that has plenty of big pictures, because it slims them all down with a single step.

4. **If you want to remove the unused portion of a cropped picture, turn on the "Delete cropped areas of pictures" checkbox.**

   Cropping lets you cut out a smaller part of a larger picture. (There's more on cropping in the next section.) When cropping a picture, Excel ordinarily keeps the full-size original in case you want to change the cropping later on. To throw out that extra information and save space, use this option.

5. **Choose an option in the "Target output" section.**

   When you compress a picture, Excel *resamples* the picture based on its current size. In other words, if you've reduced the size of a picture (by dragging the picture handles), Excel saves only enough information for the current, smaller version of the picture.

Exactly how much information Excel stores depends on the resolution option.

- **Print** tells Excel to keep enough information for a decent printout at the current size. However, you may notice a little blurriness if you resize the picture larger later on.
- **Screen** tells Excel to store a lower-quality picture that's sufficient for screen display (in a PowerPoint presentation or an online web page, for instance) but produces a poorer printout.
- **E-mail** tells Excel to store an even smaller picture, which makes sense if you're planning to send your workbook in an email and you need to pare it down as much as possible to avoid bloating your recipient's Inbox.
- **Use document resolution** tells Excel to use the standard image compression setting for this document. By default, this is 220 dpi (the same as Print). However, you can change the document resolution at any time using the Excel Options window.

Note: To change the default document resolution, choose File→Options. Then, click the Advanced category (on the left) and scroll down until you see the Image Size and Quality section. Finally, choose a new resolution from the "Set default target output to" list.

6. **Click OK to close the Compress Pictures dialog box and apply your new compression settings.**

   You can now save your new, leaner spreadsheet with the File→Save command.

   Once you compress your pictures, there's no turning back. (Of course, if you keep the original version of a picture file somewhere else on your computer, you can always re-insert it later if you need it.)

Cropping and Shaping a Picture

Ordinarily, Excel puts the whole picture you select in your worksheet. However, in some cases you may decide that you want to highlight just a small part of the picture and forget about all the rest. In this situation, you can clip your picture down to size using another program, or you can crop it right inside Excel.

To crop a picture in Excel, follow these steps:

1. **Select the picture.**
2. **Choose Picture Tools | Format→Size→Crop.**

   After you click this button, cropping handles appear on each corner of the picture and in the middle of each side. (If your picture is dark, you may need to look carefully to see the cropping handles, because they're black.)

3. **To crop your picture, click one of the cropping handles and drag it inward.**

   As you drag a cropping handle, Excel hides the outlying part of the picture (Figure 19-5).

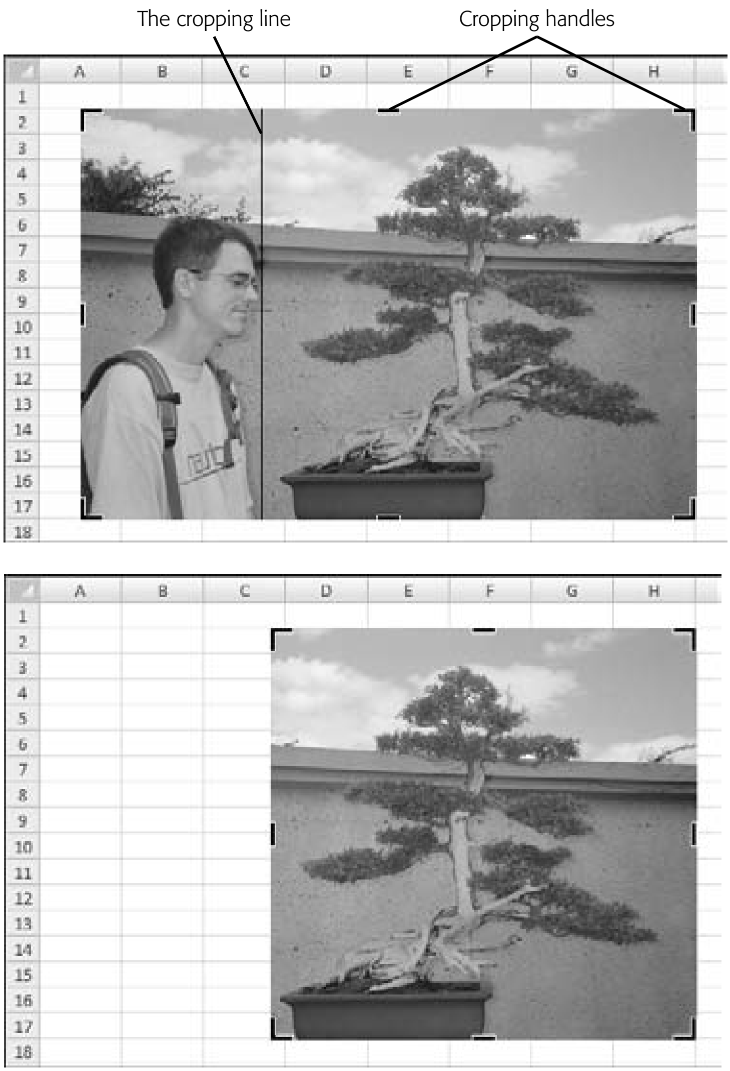

Figure 19-5: *Top: You can use just one cropping handle to remove part of the picture on one side, or you can drag all the cropping handles in until you're left with just the part of the picture you want to see.*

Bottom: The final result.

4. **Once you're finished, choose Picture Tools | Format→Size→Crop again to turn off the cropping handles.**

   Excel keeps the picture data you've cropped out in case you want to return the picture to its original size (either by recropping the picture or using the Picture

Tools | Format→Adjust→Reset Picture command). If you're sure you don't need this option, then you may want to use the Compress Pictures option (explained on page 545) to discard this extra data and cut down on the size of your workbook file. This idea is a good one if you've cropped a tiny bit out of a very large picture (in which case your workbook file's storing a lot of excess picture).

After you finish cropping a picture, you may also want to apply a different edge effect to shape it. You can choose from a huge gallery of simple shapes by choosing Picture Tools | Format→Picture Styles→Picture Shape. Figure 19-6 shows one possible example.

Figure 19-6:
Top: This picture began its life as a group portrait.

Bottom: Excel's cropping and shaping features get to the heart of the matter.

Picture Borders, Effects, and Styles

If Excel's picture coloring, cropping, and shaping features don't keep you busy, you'll be happy to learn that the graphical fun doesn't stop there. Along with the features you've already seen, Excel also lets you apply a picture border and a picture effect.

When picking a border, it's up to you to pick the color, thickness, and style. You set all these details by selecting the picture in question, and then using the Picture Tools | Format→Picture Styles→Picture Border list. Here's what to do:

- If you don't want to use basic black, pick a color from the Picture Border list. It's always a good idea to use theme colors (page 174) so your pictures blend in with the scenery.
- To make your border appear, choose a thickness from the Picture Border→Weight submenu. 1/4pt is sleek; 6pt is thick and heavy.
- If you don't want a solid border, choose another line style from the Picture Border→Dashes submenu. You'll see a variety of different types of dashed and dotted lines.
- To get rid of a border you don't like anymore, choose Picture Border→No Outline.

Picture effects are more exotic, but just as easy to discover. To get picture effects, you need to use the Picture Tools | Format→Picture Styles→Picture Effects list. You'll see submenus for applying shadow, reflection, 3-D rotation, a soft edge, or a glowing edge. Each of these submenus has a gallery of common options with thumbnail previews. Figure 19-7 shows one example.

Figure 19-7:
A 3-D rotation is a dramatic way to make the statement, "I have too much free time."

Tip: The best way to learn about all Excel's wacky picture effects is to experiment. As you move your mouse over the different picture effects, Excel's live preview changes the picture on your worksheet accordingly. To get a better look, click to apply the change, and then hit Ctrl+Z to undo it if it isn't to your liking.

If you're in a hurry, you don't need to fiddle with the picture shaping, border, and effect settings separately. Instead, you can choose a preset style that applies a combination of these settings from the style gallery in the ribbon's Picture Tools | Format→Picture Styles section. You'll find an option that makes your picture look like a postcard, a scrapbook clipping, or a wavy piece of glass that's fallen on its side (which is the current frontrunner for the "Feature Least Likely to Ever Appear in a Real Spreadsheet" award).

Excel's Clip Art Library

One of the challenges in using pictures is finding the right image. If you need to use a company logo, you probably have that on hand already. However, if you're looking for a picture on a specific subject—say a drawing of French fries for your analysis of fried versus steamed carbohydrate calories—it can take hours of web surfing to find an image that fits your requirements and is legal to use.

Excel answers this challenge with an innovative online library of clip art that, at last count, contained a staggering 150,000 images. Best of all, anyone who owns Excel can search this clip art library and use any of its images for free.

To insert an image from the clip art library, follow these steps:

1. **Select Insert→Illustrations→Clip Art.**

   A Clip Art pane appears on the right side of the Excel window (see Figure 19-8). As described in the next steps, you use it to search for images.

Note: The Clip Art pane provides a few extra options at the bottom of the window. You can click the "Find more at Office.com" link to browse Microsoft's online clip art using your web browser instead of searching from within Excel.

2. **In the "Results should be" list, you can choose the media types you want to search for.**

   The Illustrations category includes common vector graphic file types like Windows Metafiles (.wmf and .emf). *Vector graphics* are images that are stored as a series of instructions (as in "draw blue square, add green circle…"). The key benefit to vector graphics is that you can resize them without losing quality.

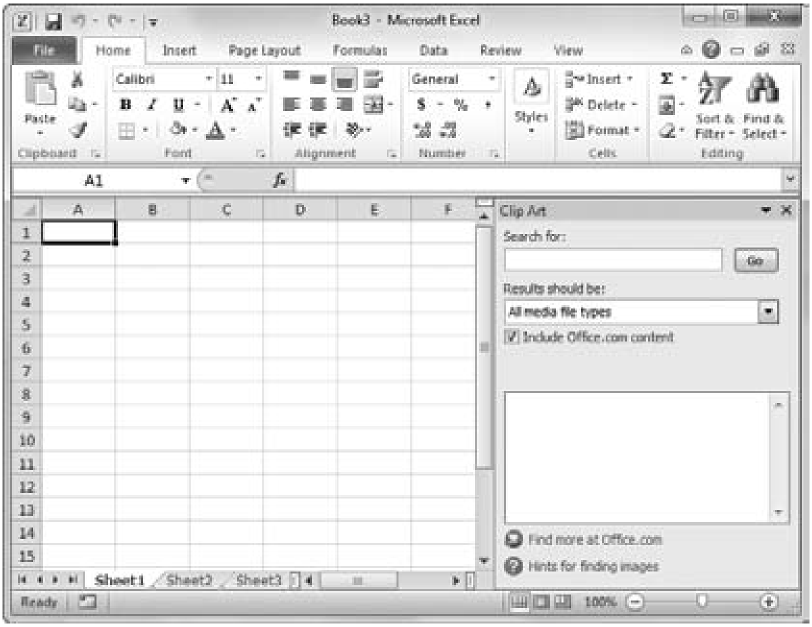

Figure 19-8: *The Clip Art pane is your doorway to the vast resources of clip art on your computer and on Microsoft's Office Online website. You can search for images using topic keywords, and then drag the pictures into your worksheet.*

The Photographs category includes common bitmap file types like Bitmaps, JPEGs, and GIFs. Bitmap graphics are stored as a series of individual points, or pixels, arranged in a rectangle. Bitmaps require more disk space than vector graphics, and you can't resize them without some compromise in quality. However, they can provide photorealistic detail.

The Videos and Audio categories include animated files and audio recordings.

Tip: The majority of Microsoft's catalog consists of ordinary clip art. However, if your search returns movie or audio files (like when you search with the keyword "scream"), you should probably remove these categories from your search and try again. Unless you want to startle your coworkers, there's really no point to dropping a sound object into your worksheet.

3. **In the "Search for" text box, enter a few words that describe the images you want to find.**

   Microsoft stores a list of descriptive words for every image in its clip art library. When you perform a search, Excel tries to match the words you enter against the description for each image. If it finds the words in the description, it shows you the picture as a match.

   You can include one word or a combination of words. If you use the search words *mad monkey*, you'll find only pictures that have the words *mad* and *monkey* in their description. On the other hand, if you search for *monkey*, you'll end up with a much larger list of results.

   For even more results, you can search for pictures that have *either one* of two or more words in their description. Just use a comma to separate the words. The search *mad, monkey* gives you all pictures that have the word *mad* in their description, and all pictures that have the word *monkey* in their description.

   If you're searching for a file in one of your computer's clip art folders, and you know the file name, you can use that as a search term. If you don't know the exact file name, you can use wildcards like ? (for any single character) and * (for any series of characters). You'd type in *car?.jpg* to locate file names like car1.jpg or car2.jpg. Or enter *car*.jpg* to locate file names like cardboard.jpg or carton.jpg. If you want to use the asterisk (*), you must place it at the end of the search string, but before the file extension. It can't appear at the beginning or in the middle of the string (**board.jpg* or car*d.jpg don't fly).

Tip: Make sure you keep the "Include Office.com content" setting switched on if you want to perform a search that includes the huge library of art on Microsoft's Office Online website.

4. **Click Go to start the search.**

   The results appear as a list of thumbnail images in the Clip Art task. Because Excel needs to download results from the Internet, it can take a little bit of time before all your results appear. You can click Stop at any point to cancel a search that's in progress.

5. **If you see an image you want to use, double-click it to insert it into your worksheet (Figure 19-9).**

   Once you add a picture to your worksheet, you can treat it like you would any other picture, so go ahead and resize it, crop it, change the contrast, and so on. When you save your worksheet, Excel adds the picture data to your file, so you don't need to download it again.

Figure 19-9: *In this example, a search using the word "dollar" has just finished. The search turned up dozens of matching images, which it shows as thumbnail-sized images, including the collection of currency that has been added to the worksheet as a picture object.*

Drawing Shapes

If the stock graphics provided in the clip art collection don't satisfy your inner art critic, you can create your own pictures. Excel's drawing features make this process a lot easier than you might expect. In fact, you can create everything from simple shapes to complex art without leaving your worksheet.

The starting point for all drawing activity is the ribbon's Insert→Illustrations→Shapes section, which is filled with potential shapes (Figure 19-10).

Before you can really get started drawing anything, you should understand the basic shape categories. They include:

- **Lines**. This category includes straight lines, curved lines, and arrows.
- **Rectangles**. Albeit ordinary, rectangles are great for storing bits of text or just wrapping themselves around groups of other shapes.

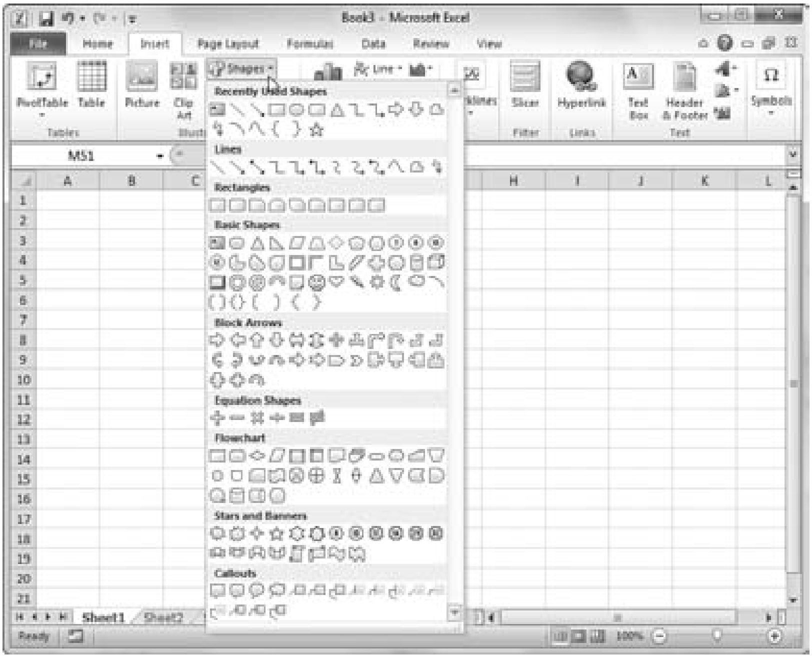

Figure 19-10: *Depending on how large your Excel window is, you may see a few shapes in the Insert→Shapes section. But if you click the drop-down arrow, you get a complete list, sub-divided into logical sections.*

- **Basic Shapes**. This category includes geometric shapes like the square, circle, rectangle, octagon, and more. Leave it to Microsoft to also include not-so-basic shapes like rings, lightning bolts, suns, moons, and even a happy face.
- **Block Arrows**. This category includes a variety of one-way and two-way arrows, as well as shapes with arrows attached to them.
- **Equation Shapes**. This category includes large mathematical symbols, like the multiplication, division, and equal signs.
- **Flowchart**. This category includes shapes that are often used in flowcharts, like the rectangle (which represents a step in a process) and the diamond (which represents a decision).
- **Stars and Banners**. This category includes the common five-pointed star and other starburst shapes. It also includes different types of banners, like award strips and unfurled scrolls. These shapes look best if you put some text inside them.
- **Callouts**. Callouts are designed to add information to a worksheet. Most Excel callouts are shapes with a connected line. The line points at something important, and the shape contains any descriptive text you want to write.

Note: Most Office programs feature these shapes. Once you learn to use them in Excel, you can also use them in the same way in Word or PowerPoint. This fact also explains the existence of some of the shapes that don't make much sense in Excel spreadsheets—they're really intended for other Office applications.

Drawing a Shape

Excel lets you draw a wide range of shapes, from simple lines and circles, to banners and three-dimensional arrows. To insert a new shape, follow these steps:

1. **Find the shape you want in the Insert→Illustrations→Shapes section, and then click it.**
2. **Click your worksheet in the spot where you want the shape to appear, and then drag to make the shape as big as you want.**

   Usually, Excel inserts the image as soon as you release the mouse button. However, some shape types, like the curved line or freeform line, have an extended drawing mode. With these shapes, every time you click the worksheet, Excel adds a new curved line segment. When you want to finish the drawing, you need to double-click the last point.

   Once the shape appears, Excel selects it, and then adds the Drawing Tools | Format tab to the ribbon.

Note: The Drawing Tools | Format tab is a lot like the Picture Tools | Format tab you learned about earlier. It includes similar buttons for applying borders and effects, and arranging and resizing your shape.

3. **Pick a color for your shape from the Drawing Tools | Format→Shape Styles→Shape Fill list.**

   This color fills the inside of all shapes except for lines. You can also choose No Fill to make the shape transparent so that other shapes (and your worksheet data) show through. You can use a circle with no fill to point out some important data on your worksheet, for example.

   Along with the standard color choices, you can also use a fancy texture, an existing picture, or a gradient. In fact, shapes offer exactly the same options that you saw when you colored in chart elements (page 510).
4. **Pick a border color, thickness, and dash style from the Drawing Tools | Format→Shape Styles→Shape Outline list.**

   To pick a shape border, you follow the same process you did to add a border to an ordinary picture from a picture file (page 550).

5. **If you want a fancy shape effect, like shadow or 3-D rotation, choose the effect from the Drawing Tools | Format→Shape Styles→Shape Effect list.**

   The effects that you can use with shapes are mostly the same as the effects you can use with pictures (although they usually make more sense with shapes). They include:

   - **Shadow** adds diffused gray shading behind your shape, which makes it look like it's floating over the page.
   - **Reflection** adds a faint copy of part of the image just under the bottom edge, as though it's being mirrored in a pool of water or piece of shiny glass.
   - **Glow** adds a blurry edge in a color you choose.
   - **Soft Edges** adds a blurry edge that softens your border.
   - **Bevel** shapes the surface of the image so that part of its surface appears raised or indented.
   - **3-D Rotation** turns the image around in three dimensions. This trick works best with images that have some depth to them—the thick block shape is a better choice than the flat square.
   - **Preset** lets you choose from some ready-made options that combine more than one effect.

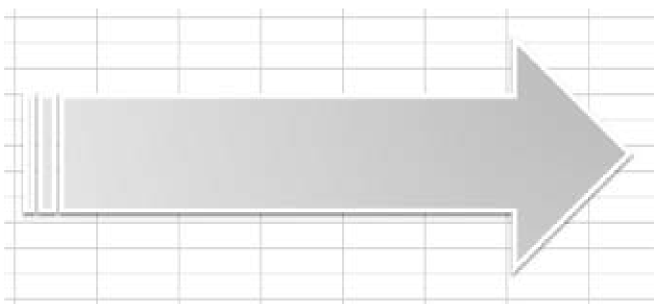

Figure 19-11:
This eye-catching arrow sports a thick white border, gradient fill, and shadow.

Tip: If you don't want to pick a separate fill color, border, and effect, you can use one of Excel's preset styles. Just make your choice from the gallery of options (each of which has a tiny thumbnail preview) in the Drawing Tools | Format→Shape Styles section of the ribbon.

6. **Now that you have perfected your shape, you can drag it to the position you want, and then resize it.**

   When you select a drawing, Excel not only displays the usual resize handles, but it also gives you one or more yellow diamonds and a green circle, as shown in Figure 19-12. You can drag the green circle to rotate the image. You can use the yellow diamonds to change the proportions on the shape. You can change the amount of curve in a curved banner, the width of each point in a star, or the length of a line in a callout. As you drag, Excel superimposes a light copy of the shape to indicate how the shape will change.

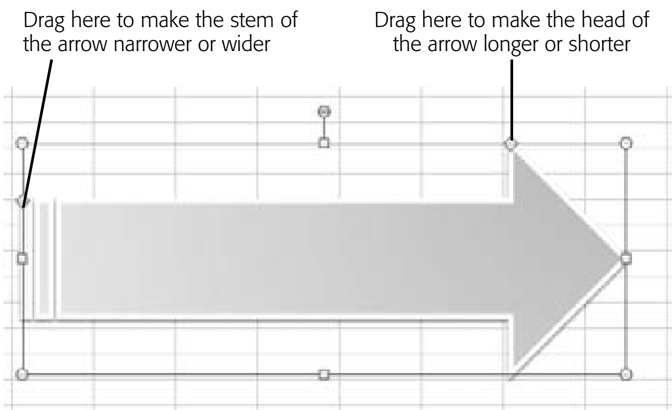

Figure 19-12: *Look for a drawing's yellow diamonds, which let you alter the shape in all kinds of weird and wonderful ways. Different shapes support different types of alterations.*

POWER USERS' CLINIC

Hard-Core Shape Manipulation

If you're feeling very punk rock, you can make dramatic changes to a shape in Excel. In fact, Excel lets you edit a shape like it's a diagram in an illustration program. To do so, select the shape, and then choose Drawing Tools | Format→Insert Shapes→Edit Shape→Convert to Freeform. Although your shape still looks the same, this action changes it to a collection of straight lines, curves, and points.

You can now modify each point. First, choose Drawing Tools | Format→Insert Shapes→Edit Shape→Edit Points to show the points on the shape. Then you can drag a point to move it somewhere else. As you move the point, Excel "pulls" the rest of the shape along with it. Drag a few points and you can take an ordinary shape like an arrow and transform it into a strangely distorted blob.

Most people find that this feature is an effective way to ruin a perfectly good shape. But your art skills could make the difference in improving a design.

Adding Text to a Shape

You can add text to almost any shape. It doesn't matter whether you've got a circle, a box, an arrow, a banner, a starburst, or even something weird. Shapes that don't have any interior space, like lines, are the only exceptions.

When you add text to a shape, the text wraps itself to fit neatly inside. Figure 19-13 shows a few examples.

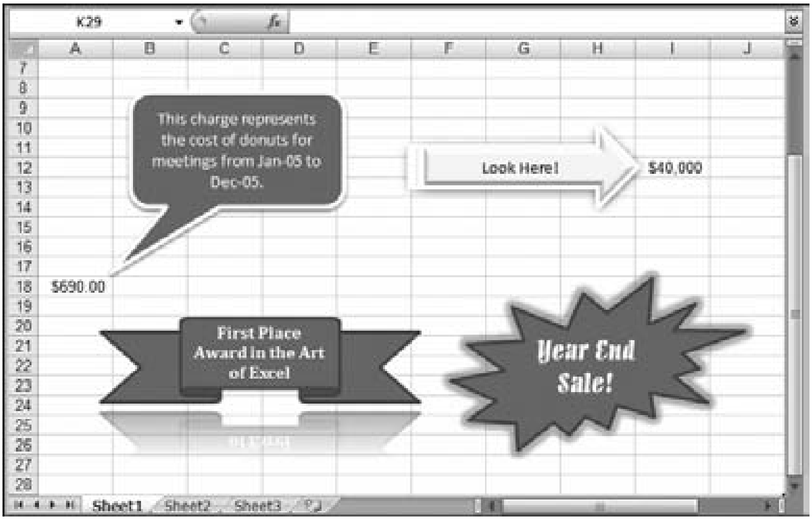

Figure 19-13: *All these shapes share one thing in common: they contain descriptive text. Used creatively, text and shapes can add a little life to your worksheets.*

To add text to a shape, follow these steps:

1. **Right-click the shape, and then choose Edit Text.**

   Your cursor moves inside the shape, and a box appears around the current text, if there is any.

   You can also just click the shape, and then start typing, without bothering to choose Edit Text.

Tip: If you want to add floating text that isn't inside a shape, choose Insert→Text→Text Box, draw the text box somewhere on your worksheet, and then start typing. Or, use one of fancy presets from the Insert→Text→WordArt gallery if you want to add a text box that already has some wild formatting in place.

2. **Type the text you want to use.**

3. **If you want to format your text, use the mini formatting bar or the buttons on the Drawing Tools | Format→WordArt Styles section of the ribbon.**

   To use the mini bar to make basic formatting changes, select the part of the text you want to change, and then choose a new font, size, color, and so on from the small toolbar that appears just above.

   To use the WordArt feature to apply eye-catching effects, select the whole shape, and then head to the ribbon's Drawing Tools | Format→WordArt Styles section. You'll find buttons that let you independently apply a fancy fill, border, and various effects. You've already used these features with shapes and pictures. Or, you can choose a ready-made combination of formatting settings from the Drawing Tools | Format→WordArt Styles→Quick Styles list. Figure 19-14 shows a shape with some WordArt-enhanced text.

Figure 19-14:
This text uses a dazzling reflection effect to distract spreadsheet readers from abysmal sales numbers elsewhere in the worksheet.

Note: There's a little-known trick you can use to put a cell reference into the text of shape. Any time you're typing text into a shape (or in a floating text box), click the formula bar. You can then type in a simple cell reference (like *=B3* to show the value that's in cell B3) or a more complicated formula, complete with all the usual Excel functions. This handy-dandy trick lets you pull the current information out of your worksheet, avoiding duplication and making sure your graphics are always up-to-date.

Selecting and Arranging Shapes

If you add enough shapes, you may start to run into trouble manipulating and layering all these different objects. Here are some potential headaches you could face:

- **Some shapes are difficult to select**. If you don't click exactly on a line, you end up selecting the worksheet cell *underneath* the line.
- **Some shapes may obscure other shapes**. What if you want to put a starburst shape inside a circle? Depending on the order in which you've added the shapes, when you move the starburst over the circle, it could actually disappear *underneath* the circle.

Excel has a handy tool to help you out. It's called the Selection and Visibility pane (shown in Figure 19-15), and you call it into action by choosing Page Layout→Arrange→Selection Pane. Or, if you have a shape that's currently selected, you can get the same feature using the Drawing Tools | Format→Arrange→Selection Pane command.

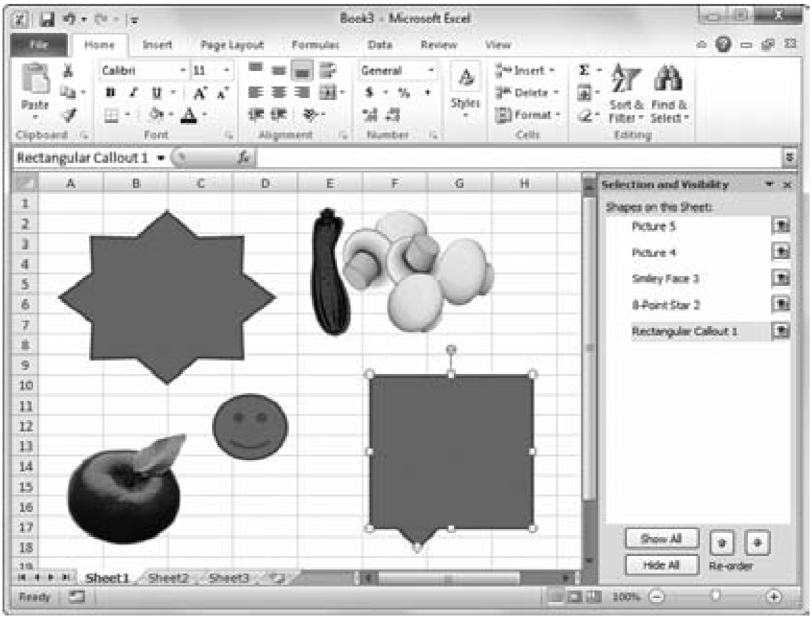

Figure 19-15: *The Selection and Visibility pane lists all the floating objects you've added to your worksheet. These objects include shapes, pictures, text boxes, and charts. The name is determined by the shape type, and the sequence in which you added it. The 10th text item you add may have a name like TextBox 10. For pictures, you see the word "Picture," and for charts you see whatever name you set in the ribbon's Chart Tools | Layout→Properties→Chart Name box.*

The Selection and Visibility pane lets you do two things: select difficult-to-reach objects, and change the way they're layered.

To select an object, simply click it in the list. This method works even if your shape is buried underneath another shape and therefore impossible to click with the mouse. Once you've selected your shape, you can move it, resize it, or format it using the ribbon.

Tip: To quickly jump from the currently selected shape to the next shape, just press Tab. The resizing handles appear around the newly selected shape. You can also use Shift+Tab to move back to the previously selected shape.

Sometimes, you'll want to select several shapes at once. To do so, just hold down the Ctrl key while you click each shape in the list. Once you've selected several shapes, you can move or format them as a group. You still use the ribbon in the same way, but now your changes affect every selected shape.

Tip: If you plan to use a group of shapes as a single unit, you can *group* them together. When you do so, Excel treats them like one shape object when you select or move them. To group shapes, select them all, and then choose Drawing Tools | Format→Arrange→Group→Group. The only disadvantage to having a grouped shape is that you can't modify the individual shapes separately unless you first choose Drawing Tools | Format→Arrange→Group→Ungroup to remove the grouping.

The Selection and Visibility pane also shines when you need to change the *layering* in your worksheet (the way that different images overlap one another). Technically, each image on your worksheet exists in its own private layer. Whenever you add a new shape, Excel creates a new layer at the top of your worksheet, and then puts the new shape in this layer. That means new objects are layered on top of older ones—which may not be what you want.

To change the way Excel layers objects, you need to change the order of items in the Selection and Visibility pane list. Objects at the top of the list appear on top of other objects further down the list (Figure 19-16). To move an item, select it, and then click the up or down arrow button. Figure 19-17 explains how to create transparent shapes—good for when you want the cells beneath your shapes to remain visible.

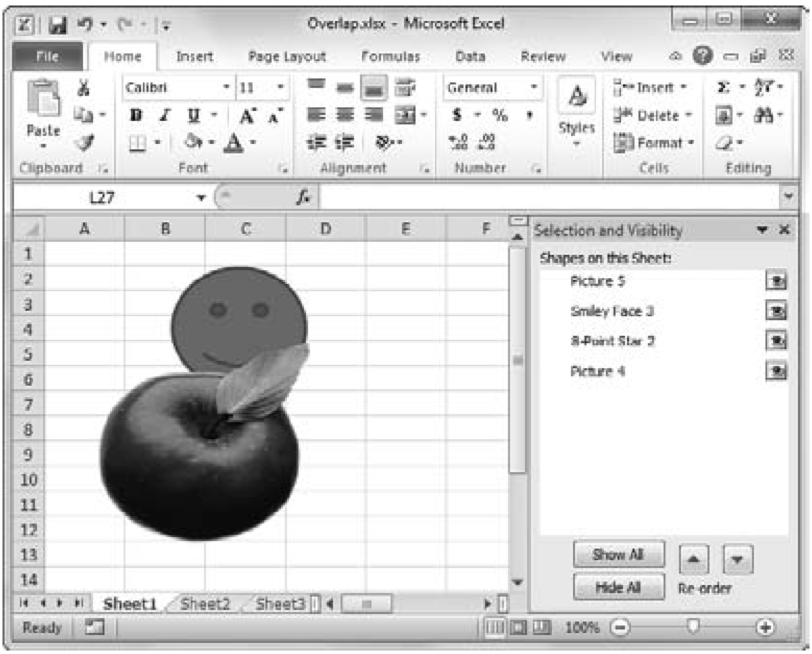

Figure 19-16: *In this example, the apple (Picture 5) appears on top of the happy face (Smiley Face 3) because it's higher in the list.*

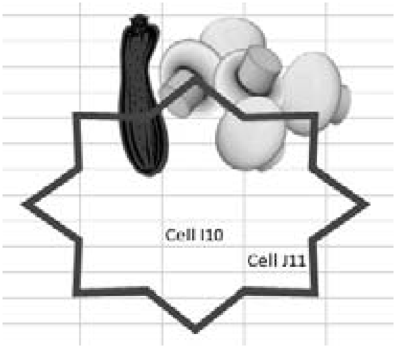

Figure 19-17:
The starburst shown here is transparent, so the data and shapes underneath show through. To create this effect, select the shape and choose Drawing Tools | Format→Shape Styles→Shape Fill→No Fill.

Note: You can also change the layering order using the Bring to Front and Send to Back buttons in the ribbon's Drawing Tools | Format→Arrange section. However, the Selection and Visibility pane is much easier to use.

If you have an extraordinarily complex worksheet that's dense with shape objects, you may find that it helps to temporarily hide the ones that you aren't interested in. You can do this in two ways:

- **To hide everything**, click the Hide All button in the Selection and Visibility pane. Now, whenever you click to select an image in the list, it appears. When you're done making your changes and you're ready to see all your shapes again, click Show All.
- **To hide just a few items**, click the eye icon next to each item in the Selection and Visibility pane list. Click the eye again to make the item reappear (or use the Show All button to show everything).

TIMESAVING TIP

Lining Up Shapes

When creating a complex piece of Excel art (like a diagram that's entirely made up of shapes), you need a way to line up shapes with exact precision. Simply dragging each shape into place with the mouse may be too difficult or just take too long. Happily, Excel has an automatic alignment feature that can really help you out.

To use Excel's alignment feature, begin by selecting all the shapes that you want to line up. (Hold down the Ctrl key, and then select each one on your worksheet or in the Selection and Visibility pane.) Then, make a choice from the Drawing Tools | Format→Arrange→Align section of the toolbar. You can line up shapes along their left, right, top, or bottom edges, or center them so their mid-points line up. And if you've selected more than two objects, you can use the Distribute Horizontally and Distribute Vertically commands to space them out evenly, with a consistent amount of space between each shape.

Connecting Shapes

On the one hand, you may have noticed that the shapes in the Lines category are perfect for connecting other shapes. The creators of the Office shape-drawing model also recognized this fact, and they made it easier for you to snap your lines into place.

On the other hand, you may be wondering if there's really any point to connecting two shapes. Why not just drag your line anywhere on the border of a nearby shape? After all, even if you don't hit a connection point, it still looks like your line is connected to the shape. Connections' real benefit appears when you *move* the connected object. Imagine you have a line that links together two squares. If you've used connections, when you drag one of the squares to a new place, the line follows. If you haven't used connections, you have to move the square, and then resize the line every time.

Here's how to connect shapes. Every shape has predefined *connection points*, which are ideal places where you may want to connect a line. A basic rectangle has four connection points—one in the middle of each side—and a typical circle has about eight, arranged in even intervals along the border. To use the connection point feature, you simply click to select a line, and then drag one end over another shape. As you get close, Excel shows you all the connection points using small red squares (see Figure 19-18). When you drop the line in one of these places, you've created a connection.

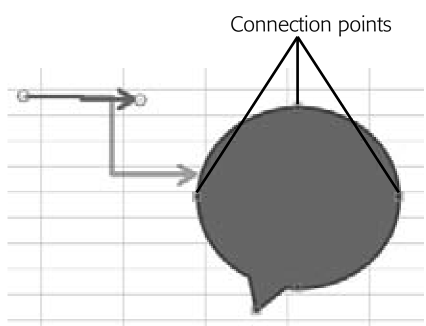

Figure 19-18:
As you drag the line with the arrow over this callout (page 555), Excel shows the possible connectors.

Connections let you take separate shapes and build more impressive diagrams. When you use connections, you may also want to consider using grouping (page 562). And if you don't want to connect everything on your own, you may be interested in the SmartArt feature (discussed on page 565), which gives you premade diagrams that include numerous shapes and connecting lines.

GEM IN THE ROUGH

Drawing Graphic Objects in Charts

You may find that shapes are more trouble than they're worth, because they can quickly gunk up a worksheet. But graphic objects become extremely useful in one area: your charts. With the right shapes, you can break out of Excel's limiting rules for labeling and highlighting data and add eye-catching arrows and shapes.

Excel lets you draw on a chart object in the same way that you draw on a worksheet–using the tools found on the Insert→Illustrations section of the ribbon. Best of all, once you draw a shape in a chart box, it stays locked into that box. That means if you move the chart, the shape follows along, remaining in the appropriate position. Figure 19-19 shows some of these techniques.

Here are a few ways that Excel's drawing features can enhance your charts:

- Use arrows to point to important places on a chart. This technique works well if you need to highlight a single data point.
- Use circles or squares around an important region on the chart. This technique works well if you need to highlight a section containing multiple data points.
- Use callouts to add descriptive text explaining why a chart line takes a sudden dive or turns upward suddenly.
- Add picture objects, like logos or a themed background (for example, show a picture of a beach in a chart that tracks favorite vacation destinations).

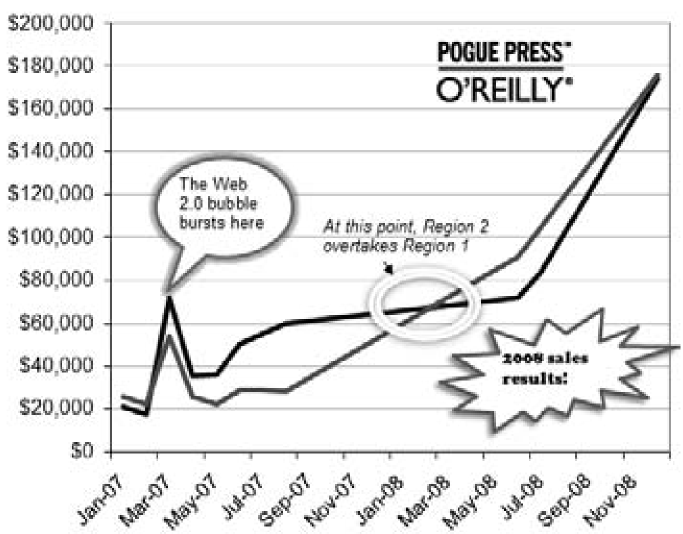

Figure 19-19:
This chart features half a dozen graphical elements. A logo floats in the top-right corner, while a starburst announces the results in the bottom-right. A callout points to a sudden change in the data, and a combination of a textbox, arrow, and ring highlights where the two lines cross.

SmartArt

SmartArt is a feature that lets you create business graphics and place them in your Excel worksheet. Figure 19-20 shows a few examples of SmartArt diagrams.

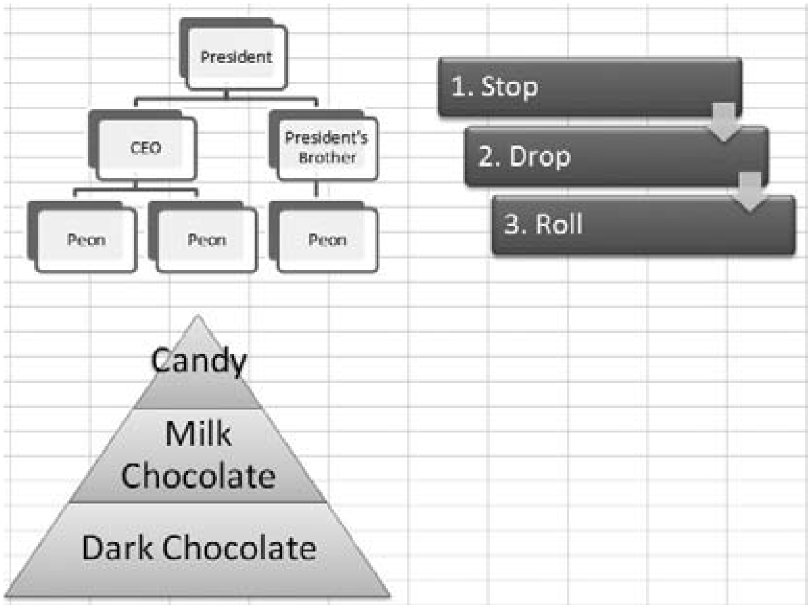

Figure 19-20: *You can use SmartArt graphics to show a company's organizational structure (top left), a sequence of steps (top right), a pyramid of good eating choices (bottom), and much more.*

SmartArt and Excel have a slightly awkward relationship. Although the SmartArt graphics are unarguably attractive (and easy to build), they don't make sense in most Excel workbooks. After all, most people expect to use Excel to record reams of numbers, and analyze them with number-crunching formulas and sophisticated charts. Diagrams make more sense in the company report (a Word document) or a budget presentation (a PowerPoint document). Truthfully, you're more likely to use SmartArt in both these programs than in Excel.

Furthermore, SmartArt diagrams have the same fundamental limitation that all shapes have in Excel—you can't use cell references as part of your text. That means if you need to have summary numbers or the result of a complex calculation in a SmartArt diagram, you need to copy the values yourself (and remember to update them when the value changes in the worksheet).

Even with all these considerations, SmartArt can still help you create a professional-caliber diagram in a hurry.

Here's how to create a SmartArt diagram:

1. **Choose Insert→Illustrations→SmartArt.**

   The Choose a SmartArt Graphic dialog box appears (Figure 19-21).

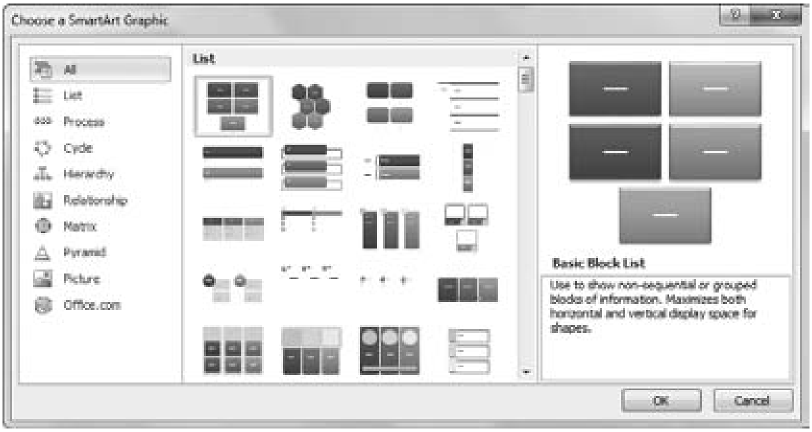

Figure 19-21: *You can choose from a gallery of about 80 diagram types.*

2. **Choose the diagram you want to use, and then click OK.**

   The Text pane appears where you can enter the text that's presented in the diagram (Figure 19-22). If the Text pane doesn't appear automatically, choose SmartArt Tools | Design→Create Graphic→Text Pane.

3. **Type the text for your diagram.**

   The Text pane works like a bulleted list. You press Enter to create more lines, each of which is called a *point*. You hit Tab to indent your point one level, which makes it a subpoint. (Press Shift+Tab to turn a subpoint back into a normal point.)

   Each point is always a separate shape. However, different diagrams have different ways of presenting points and subpoints. In some diagrams (like the process diagrams), the subpoints become bulleted points inside each shape. In others (like the organizational diagrams), the subpoints become their own distinct shapes, as shown in Figure 19-22.

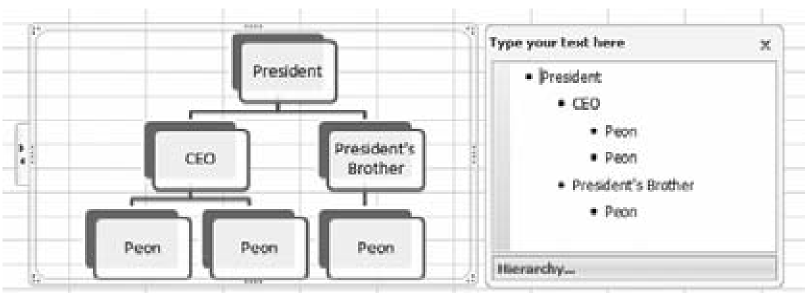

Figure 19-22: *You don't need to fiddle with individual shapes to create a SmartArt diagram. Instead, you type in everything you need in a special Text pane (shown here on the right). Excel uses this text to create the diagram.*

4. **When you're finished, drag your diagram into place, and then resize it as you see fit.**

   You can click anywhere on your worksheet to return to Excel.

Once you've created a diagram, you can format and fine-tune it much the way you format shapes. First, select your SmartArt graphic, at which point two new tabs appear in the ribbon: SmartArt Tools | Design and SmartArt Tools | Format.

Here are some tricks you may want to try out:

- Make a choice from the SmartArt Tools | Design→Layouts gallery to switch to a different type of diagram. Excel automatically transfers the text you entered to the new diagram.
- Make a choice from the SmartArt Tools | Design→SmartArt Styles gallery to apply different color, border, and shape effects.
- Click one of the shapes inside the SmartArt graphic (like a single text box). You can then format it differently so it stands out from the rest using the Smart Tools | Format tab. Among the changes you can make, include changing the type of shape, applying WordArt effects, and moving or resizing the shape.
- Choose Smart Art Tools | Design→Reset→Reset Graphic to clear your formatting and return everything to normal. The list of points and subpoints that you've typed in remains.

CHAPTER

20

Visualizing Your Data

Ordinary numbers are fun, but on their own they lack a little pizzazz. That's why when numbers need to advertise just how important they are, they usually turn to some sort of *data visualization*—a graphical technique that turns part of their information into a picture.

You've already learned about Excel charts, which are the most obvious type of data visualization. The graphics you put together yourself can also serve the same purpose, although they're not nearly as flexible. (The problem is that, when your data changes, your custom graphics might not apply any longer. To keep them up to date, you'll be forced to perform constant edits and fine-tuning.)

In this chapter, you'll learn about several additional types of data visualization:

- **Data bars, color scales, and icon sets.** These three features extend Excel's conditional formatting feature, which you learned about in Chapter 6. All of them modify cells based on their values—for example, by drawing a shaded bar, changing the background color, or adding a tiny picture.
- **Sparklines.** These tiny chart-like graphics show basic information (trends, patterns, and changes) without a lot of fuss. Even better, they integrate seamlessly into your worksheet, unlike massive floating charts.

Used with a little bit of thought and experience, both of these features can extract meaning out of the densest jumble of worksheet data and turn otherwise drab spreadsheets into clear and powerful message-makers.

Data Bars

Data bars—a feature that places a shaded bar in the background of every cell you select—are one of the simplest and most useful forms of conditional formatting. The trick is that the data bar's length varies depending on the content in the cell. Larger values generate longer data bars, while smaller values get smaller data bars.

To see how this works, consider the worksheet shown in Figure 20-1, which shows a boring grid of numbers with no formatting.

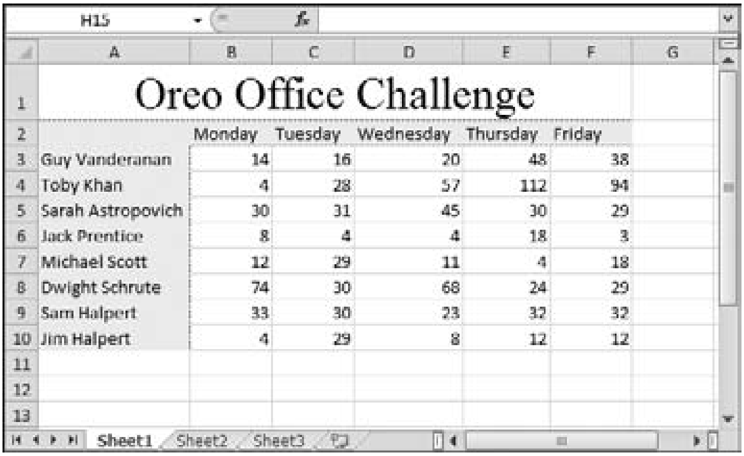

Oreo Office Challenge

	Monday	Tuesday	Wednesday	Thursday	Friday
Guy Vanderanan	14	16	20	48	38
Toby Khan	4	28	57	112	94
Sarah Astropovich	30	31	45	30	29
Jack Prentice	8	4	4	18	3
Michael Scott	12	29	11	4	18
Dwight Schrute	74	30	68	24	29
Sam Halpert	33	30	23	32	32
Jim Halpert	4	29	8	12	12

Figure 20-1: *This worksheet shows the weekly results of the company Oreo-eating competition. Right now, it's a densely packed grid of information, so it's hard to see who set the most impressive totals.*

To use data bars, select the cells you want to format (in this example, that's cells B3 to F10), and then choose Home→Styles→Conditional Formatting→Data Bars. You see a gallery with several different data bar choices, arranged into two categories:

- **Gradient Fill.** Choose this and you get attractively shaded data bars that grow lighter as they travel from left to right. Not only is this a nice detail, but it also has a practical purpose. As you know, Excel aligns numbers against the right side edge of a cell. With gradient data bars, the right side of the bar is much lighter than the left side, making your superimposed numbers easy to read.
- **Solid Fill.** Choose this and you get solid, filled-in data bars with no outline. They're plain but functional. Solid data bars work well if you plan to show data bars in one column and the numbers in another column. (To pull off this trick, you need the Show Bar Only setting described on page 574.)

Once you decide whether to have a gradient fill or a solid fill, you simply need to pick the color you want (blue, orange, purple, and so on). Figure 20-2 shows a worksheet that uses gradient data bars.

I15

	A	B	C	D	E	F
1	Oreo Office Challenge					
2		Monday	Tuesday	Wednesday	Thursday	Friday
3	Guy Vanderanan	14	16	20	48	38
4	Toby Khan	4	28	57	112	94
5	Sarah Astropovich	30	31	45	30	29
6	Jack Prentice	8	4	4	18	3
7	Michael Scott	12	29	11	4	18
8	Dwight Schrute	74	30	68	24	29
9	Sam Halpert	33	30	23	32	32
10	Jim Halpert	4	29	8	12	12

Sheet1 Sheet2 Sheet3

Figure 20-2: *The data bars let you quickly pick out the largest and smallest values. For example, there's no doubt that Toby's Thursday and Friday totals and Dwight's performance on Monday set the office standard.*

When using data bars, Excel finds the largest value and makes that the largest data bar (so it fills the entire cell). Excel fills every other cell proportionately. That means a cell with a value that's half the maximum gets a data bar that fills half the cell, and a cell with a value of 0 has no data bar at all.

The best part about data bars is that Excel keeps them synchronized with your data. In other words, if you change the worksheet shown in this example by filling in the numbers for the next week, Excel automatically adjusts all the data bars.

Data bars work best with groups of numbers that are spread out evenly. This feature makes sense for recording the Oreo cookie extravaganza, because the Oreo consumption of most employees falls into the same basic range. However, when you have one or two values that are dramatically higher or lower than the rest, they skew the scale and make it hard to see the variance between the other values.

Incidentally, data bars also work with negative numbers. To deal with this situation, Excel draws a dashed vertical line to represent the zero mark in each cell. Positive values are drawn using a shaded bar that stretches from this line to the right, while negative values are drawn from this line to the left and colored red (Figure 20-3).

The position of the dashed zero-line depends on the scale. For example, if your values range from –10 to 10, the dashed line falls in the middle of every cell. But if your values range from –20 to 10, the dashed line is closer to the right side, because the zero value is placed proportionately.

Tip: It's best not to use red, pink, or purple data bars if your data includes negative numbers. If you do, it becomes more difficult to separate the bars that represent positive values from those that represent negative numbers.

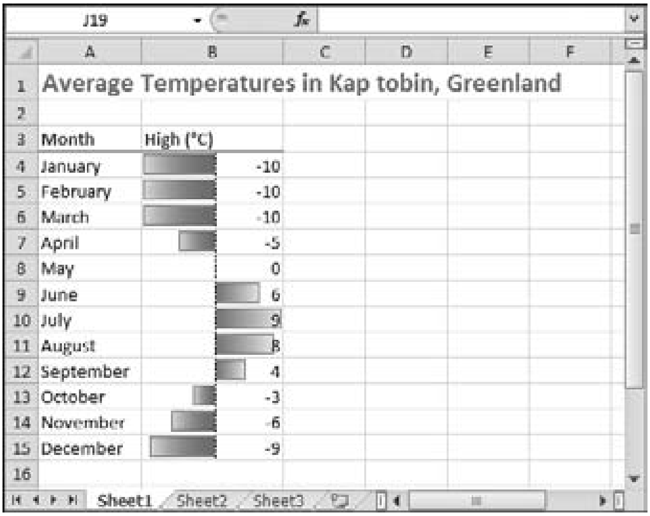

Figure 20-3:
Excel uses a dashed line to show where zero falls. It places this line proportionately between the lowest negative value and the highest positive value and keeps the position the same in every cell, so it's easy to compare numbers at a glance.

TROUBLESHOOTING MOMENT

Conditional Formatting in Older Versions of Excel

Some conditional formatting tricks—namely, data bars, color scales, icon sets, conditional formatting that overlaps, and cells with more than three conditional formatting rules—don't work in versions of Excel before Excel 2007. If you use one of these tricks, export your spreadsheet to the Excel 97–2003 standard (as described on page 41), and then open it in an older version of Excel, you won't see your conditional formatting. Instead, your cells will have the normal, unformatted appearance.

Unless you start fiddling with the conditional formatting settings in an older version of Excel, all the formatting should return the next time you open the spreadsheet in Excel 2010. That's because the information is still there in your spreadsheet file. Older versions of Excel just can't deal with it.

Editing a Formatting Rule

When you apply data bars, color scales, or icon sets to a group of numbers, Excel creates a new conditional formatting rule, which tells Excel how to format the group of cells you've selected.

To fine-tune the way conditional formatting works, you can tweak your rule. Here's how:

1. **Select Home→Styles→Conditional Formatting→Manage Rules.**

   The Conditional Formatting Rules Manager appears, displaying any conditional formatting rules you previously created. (You learned about this window on page 192.)

2. **Select the rule you want to edit, and then click Edit Rule.**

   The Edit Formatting Rule dialog box appears, which looks exactly the same as the New Formatting Rule dialog box you saw earlier. The top portion of the window lets you change the type of rule (which isn't what you want to do), and the bottom section lets you refine the current rule (which makes more sense).

3. **Modify the settings in the "Edit the Rule Description" section, and then click OK to apply them.**

   You can also click Preview to see what the result of your change will be on your worksheet.

4. **Click OK again to close the Conditional Formatting Rules Manager.**

   The options in the Edit Formatting Rule dialog box depend on the type of rule you're modifying. The following sections guide you through your choices.

Fine-Tuning Data Bars

If you're editing a data bar (using the steps described in the previous section), you can change several details:

- Use the Show Bar Only checkbox to hide cell value and show just the data bar.
- Use the Minimum and Maximum settings to control how Excel determines the length of each bar.
- Use the Bar Appearance section to set the exact color that you want for the data bars and the border that's drawn around them.
- Click the "Negative Value and Axis" button to control the color of data bars that represent negative values and the position of the dashed line that represents the zero point (which can be at the midpoint of the cell or placed proportionately between the smallest and largest value).
- Use the Bar Direction setting to make right-to-left data bars, instead of the standard left-to-right ones.

The following examples show two ways you can use these settings to create fine-tuned data bars.

Putting data bars in a separate column

Using data bar customization, you can put your numbers and the corresponding data bars in separate columns. Some Excel purists love this trick for the clear, no-nonsense way it presents their data. Figure 20-4 shows how you can use this technique on a worksheet that shows test scores for a top culinary school.

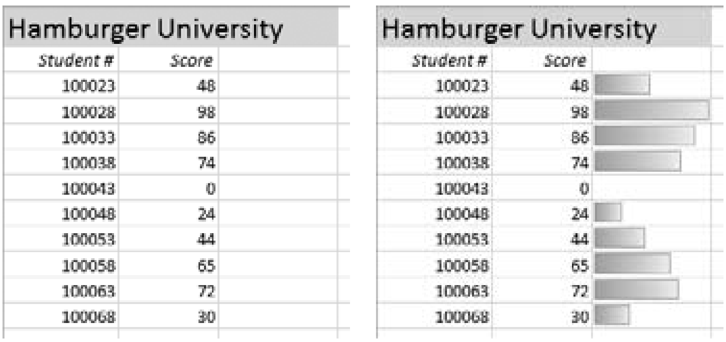

Hamburger University

Student #	Score
100023	48
100028	98
100033	86
100038	74
100043	0
100048	24
100053	44
100058	65
100063	72
100068	30

Hamburger University

Student #	Score
100023	48
100028	98
100033	86
100038	74
100043	0
100048	24
100053	44
100058	65
100063	72
100068	30

Figure 20-4:
Left: Initially, the values are in the Score column.

Right: Now, the data bars are in the adjacent column. This makes it easier to compare the data bars at the glance, and it lets you use solid color data bars without obscuring the numbers.

To create the spreadsheet shown in Figure 20-4, you need two tricks: a column of cell references and the Show Bar Only checkbox. Assume that you start with the values in column B. Here's how to create the data bars in this example:

1. **The first step is to copy your values. Select the cells on your worksheet (in this case, that's cells B3:B12), and then press Ctrl+C.**

   You can't create data bars unless you put some data in the cells first. So the first step is to transfer the data you need to the right cells—but with a twist.

2. **Move to the first cell where you want to put your data bars (in this case, cell C3).**

   Don't paste your cells just yet! The last thing you want to do is duplicate a bunch of numbers, because the next time you change them you'll wind up with inconsistent data. Instead, you need to use a pasting trick you learned on page 108.

3. **Choose Home→Clipboard→Paste→Paste Link.**

   This pastes a series of cell references. For example, cell C3 gets a formula that points to cell B3, so when you change cell C3, cell B3 refreshes with the new value automatically. Presto: the same values in two places, with no possibility for inconsistency.

4. **Now it's time to add the data bars. Select the newly pasted cells.**

   This is where you want to place the data bars.

5. **Choose a data bar option from Home→Conditional Formatting→Data Bars.**

   Because you're going to hide the cell values, feel free to choose a solid fill if you like.

6. **Now you need to edit the rule to hide the numbers. With the same cells selected, choose Home→Styles→Conditional Formatting→Manage Rules.**

   The Conditional Formatting Rules Manager appears.

7. **Select the data bar rule you just created, and then click Edit Rule.**

   The Edit Formatting Rule dialog box appears.

8. **Put a checkmark in the Show Bars Only setting, and then click OK. Click OK again to close the Conditional Formatting Rules Manager.**

   Although your pasted cells still have the values in them, the values no longer appear in the cells. Even better, because the cells in column C are linked to the cells in column B, when you change the visible data, the data bars are also updated.

Changing the data bar scale

Ordinarily, Excel decides what scale to use when it draws your data bars. In its standard scale, the smallest value is 0 (unless you're using negative numbers, as explained on page 571), and the highest value is the highest number in your selection of cells. However, this isn't always the best approach. For example, if most of your values fall into a tight range but include a few super-high values, the outliers will distort the scale. Figure 20-5 shows the problem.

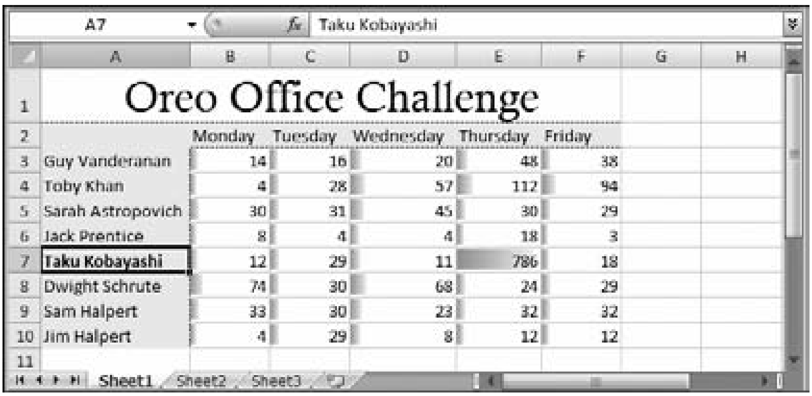

Figure 20-5: *One challenge when using data bars is how to handle off-the-charts values (like Kobayashi's Thursday performance) that make everyone else's numbers look the same. The difference between the smallest value (4) and the largest value (786) is much larger than the difference between most values. See Figure 20-6 for an elegant solution.*

To fix this problem, you need to take control of the data bar scale. By manually setting a smaller maximum value, you can focus the scale on the range of values that's most important for your data.

To accomplish this, begin by editing your data bar rule (as described on page 572). Then, in the Maximum section, change the Type to Number, and then set the Value to 100 (Figure 20-6).

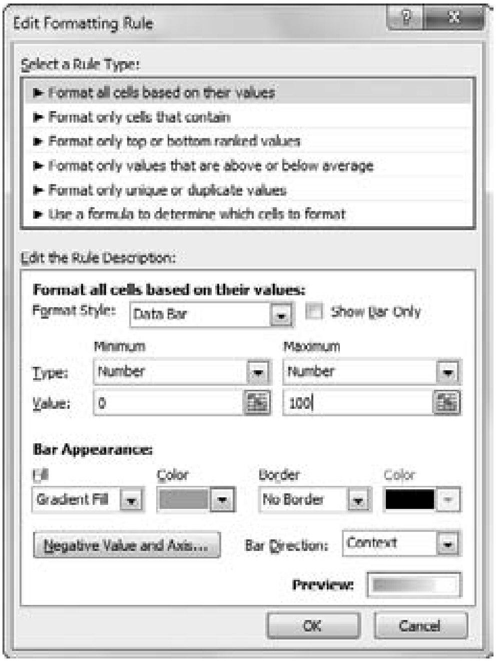

Figure 20-6:
The data bars in this example use an explicit maximum. Excel treats 0 as the smallest value and 100 as the largest. All values in between are filled proportionally (so 50 is a half-length data bar). Any value less than the minimum gets the minimum bar size, and any value larger than the maximum gets the maximum bar size. Figure 20-7 shows the result.

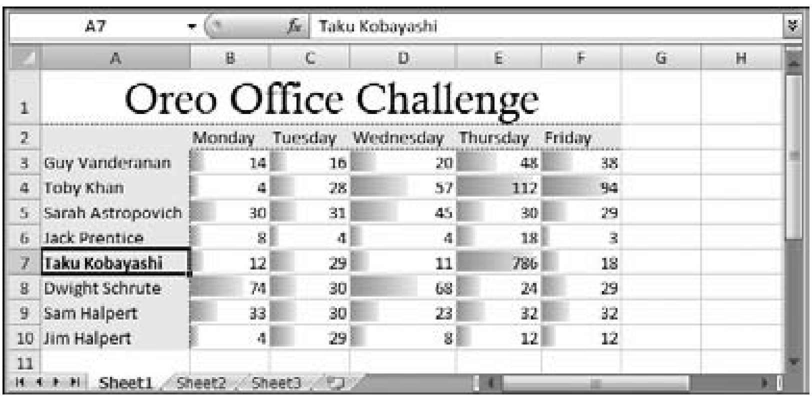

Figure 20-7:
Now it's easy to see the variation between midrange values. The only tradeoff is that every value above 100 gets the same data bar, so you lose the ability to distinguish between high and extremely high values.

UP TO SPEED

Setting the Scale

When you edit a data bar rule, you can set the minimum and maximum of the scale in several different ways. It all starts with the Type list, which controls how Excel determines these values. You can set the type of both the minimum and maximum, and they don't need to match. Your options include:

- **Automatic.** Excel uses its standard approach to determine scale, making 0 the minimum and using the highest value for the maximum. (One exception: if there are negative values, the lowest value becomes the minimum.)
- **Number.** You supply the number you want to use for the minimum or maximum. Because data bars can't get shorter than the minimum value or longer than the maximum, you can use this setting to "cap off" outlying values, as shown in Figure 20-7.
- **Percent.** This option lets you supply the minimum and maximum values as *percentages* instead of fixed numbers. Ordinarily, the lowest value is 0 percent and the highest value is set to 100 percent. So, if you want to cut off the ends of the scale, you might start the smallest bar at 10 percent and cap the highest at 90 percent. If your actual values range from 0 to 500, the bottom value (at 10 percent) becomes 50, and the top value (at 90 percent) becomes 450.
- **Percentile.** This option works similarly to Percent, but in a way that's more satisfying for mathematically-minded statisticians. This number arranges all the values in order from lowest to highest, and then slots them into different *percentiles*. In a set of 10 ordered values, the 40th percentile is always the fourth value, regardless of its exact number. In other words, when you're using percentiles, Excel isn't all that interested in how high or low the exact value is; instead, it pays attention to how that value falls in relationship to everything else. If you set a data bar minimum to a percentile value of 10, the bottom 10 percent of values get the shortest bar. If you set the maximum to 90, the top 10 percent of values get the longest bar. The nice thing about percentiles is that you always get a good range of short, medium, and long data bars, even if your numbers unevenly spread.
- **Lowest Value (for Minimum) or Highest Value (for Maximum).** Setting Lowest Value for the minimum exaggerates the differences between your data bars, especially if none of your values are close to zero. Setting Highest Value for maximum has no real effect, as that's Excel's default anyway.
- **Formula.** This option is an advanced trick that lets you use a formula that tells Excel what the highest or lowest value should be. Chapter 8 has more about formulas.

The previous example showed how data with a wide spread can cause problems. Another related problem happens if none of your numbers are close to 0. Once again, the scale tends to wipe out fine distinctions. Figure 20-8 shows the problem and explains the solution.

Figure 20-8:
Left: These sales figures are close together but far from zero. The result is a set of data bars with little obvious difference. If this isn't what you want, you can adjust the data bar scale. But this time, tweak the lower bound of the scale rather than the upper bound.

Right: Here, the lower bound is set at 115,000 and the difference between the bars is obvious–even exaggerated.

Color Scales

Color scales let you format different cells with different colors. As with data bars, Excel automatically chooses the right color for each cell. Excel assigns a predefined color to the lowest value and another predefined color to the highest value, and it uses a weighted blend of the two for all the other values. For example, if 0 is blue and 100 is yellow, the value 50 gets a shade of green.

To apply a color scale, select your cells, choose Home→Styles→Conditional Formatting→Color Scales, and then choose one of the color combinations shown in the gallery. To test out different color scale options, hover over them with the mouse, and then take a look at the live preview on your worksheet.

Color scales aren't used as often as data bars, because they tend to create a more visually cluttered worksheet. If you do decide to use color scales, you may want to customize how they work so they're applied only to specific cells, or so that they use a less obtrusive pair of colors (like white and light red). You'll learn to apply these techniques in the next section.

Fine-Tuning Color Scales

You can fine-tune color scales in much the same way you adjust data bars. Here are the settings you have to play with:

- In the Format Style list, choose 2-Color Scale or 3-Color Scale, depending on how many colors you want to use.
- Choose the colors for the minimum and maximum (and, in a three-color scale, the midpoint).
- Set the Type and Value of the Minimum, Midpoint, and Maximum to control the scale explicitly. These settings work exactly the same as they do for data bars, giving you the ability to set the scale with fixed values you supply, percentages, or percentiles, as described on page 577.

One reason you might fine-tune a color scale is to use less obtrusive formatting by setting white as your minimum value color and choosing a light color for the maximum. You can also use percentiles to make a more reserved-looking worksheet, where most values get no color and the highest values get just a tinge (see Figure 20-9).

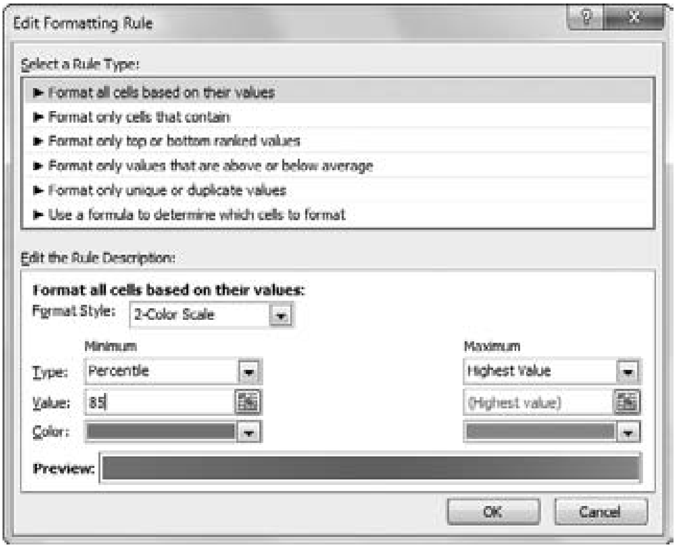

Figure 20-9:
A two-color scale blends from one color to another, while a three-color scale has a specific midpoint color. Excel blends values in the bottom range between the lowest value color and the midpoint color, and values in the top range between the midpoint color and the highest value color. In this example, the minimum type is set to Percentile and the minimum value is set to 85, which means only the top 15 percent of values have any color.

Icon Sets

So far, you've seen how to highlight different values with a different graphical representation using shaded bars and colors. Both these tricks are called *data visualizations*. Excel has one more data visualization tool that uses conditional formatting: icon sets.

The idea behind icon sets is that you choose a set of three to five icons. Excel then examines your cells and displays one of these icons next to each value, depending on the value. To do this, Excel begins by determining the range of values you have. Excel then divides that range into equal portions based on the number of icons you're using.

For example, if you choose an icon set with three icons, Excel carves your scale into three segments. If your values range from 0 to 100, cells that have a value from 0 to 33 get the first icon, cells with a value of 34 to 66 get the second icon, and cells with a value of 67 to 100 get the third icon. Similarly, if you pick an icon set with four icons, the first range is 0 to 25, the second range is 26 to 50, the third range is 51 to 75, and the fourth range is 76 to 100. (If this automatic scale setting doesn't work for you, it's possible to set the ranges for each icon manually. You'll learn how to pull off this trick on page 582.)

To see the available icon sets, choose Home→Styles→Conditional Formatting→Icon Sets (Figure 20-10).

Sadly, you can't create your own icon sets in this version of Excel. However, you can accomplish a lot by using the existing icon sets with a little imagination. Figure 20-11 shows two options with the table of Oreo eaters. The top example ranks each competitor's performance each day, while the bottom one calculates per-person totals (using the SUM() function described on page 268), and then adds icons to those totals.

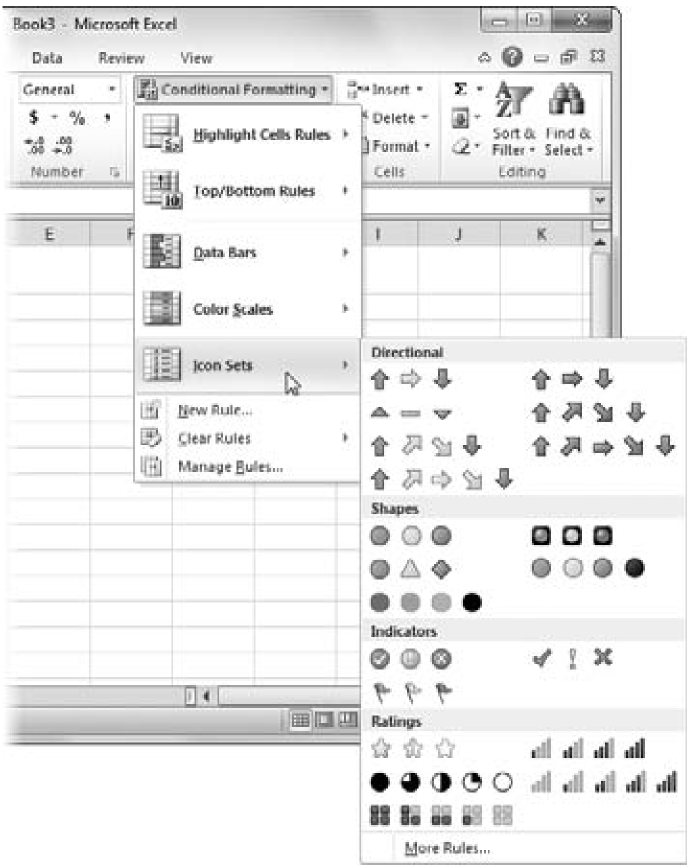

Figure 20-10:
Excel has several useful icon sets. Popular choices are red-yellow-green shapes, arrows pointing in different directions, Xs and checkmarks, moon phases, and partly filled circles.

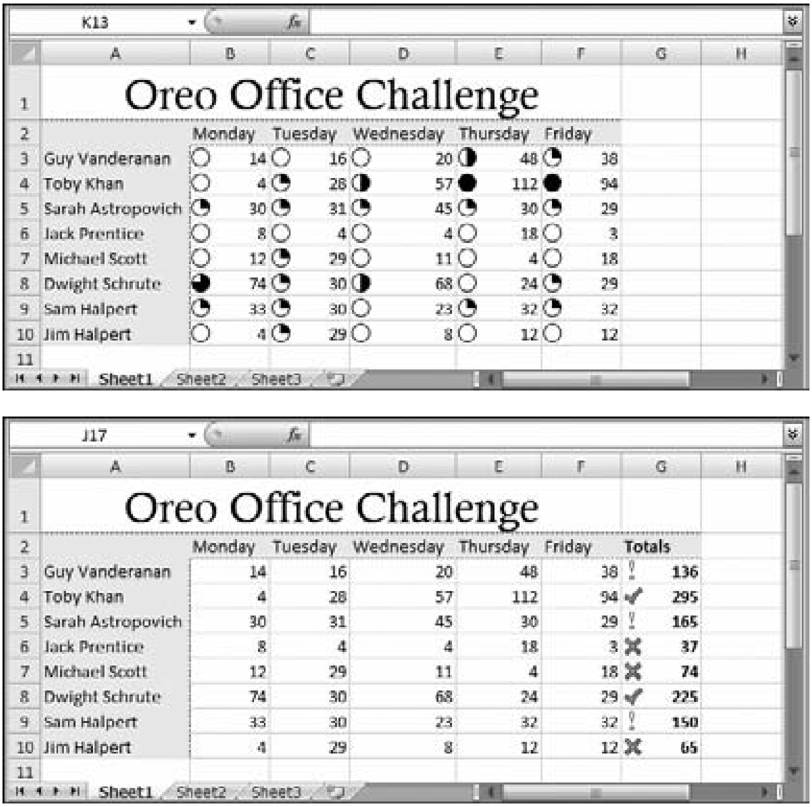

Oreo Office Challenge

	Monday	Tuesday	Wednesday	Thursday	Friday
Guy Vanderanan	14	16	20	48	38
Toby Khan	4	28	57	112	94
Sarah Astropovich	30	31	45	30	29
Jack Prentice	8	4	4	18	3
Michael Scott	12	29	11	4	18
Dwight Schrute	74	30	68	24	29
Sam Halpert	33	30	23	32	32
Jim Halpert	4	29	8	12	12

Oreo Office Challenge

	Monday	Tuesday	Wednesday	Thursday	Friday	Totals
Guy Vanderanan	14	16	20	48	38	136
Toby Khan	4	28	57	112	94	295
Sarah Astropovich	30	31	45	30	29	165
Jack Prentice	8	4	4	18	3	37
Michael Scott	12	29	11	4	18	74
Dwight Schrute	74	30	68	24	29	225
Sam Halpert	33	30	23	32	32	150
Jim Halpert	4	29	8	12	12	65

Figure 20-11: *Top: The moon phase icon set (known by the seemingly contradictory name 5 Quarters) shows how different employees stack up in the Oreo-eating competition. A full moon, for example, indicates a competition-leading performance. Empty moons are clear signs of Oreo-haters.*

Bottom: It's time to go to round two. Employees with a checkmark icon qualify to move on, while those with an X icon can stay home and drink milk. Employees with an exclamation mark are in the middle ground and can fight for the last remaining playoff spaces in another qualifying round.

Fine-Tuning Icon Sets

With icon sets, you have several settings to play with:

- Use the Icon Style list to change the icons you're using. You have all the options you saw in the Home→Styles→Conditional Formatting→Icon Sets gallery.
- Use the Reverse Icon Order checkbox to arrange icons in reverse order (so the icon that you previously used for the top values now applies to the bottom ones).
- Use the Show Icon Only checkbox to hide cell values and display just the icon. You can also use this trick with data bars, but it's way more convenient with icon sets. It lets you create a dashboard-like display that indicates the *significance* of different values (for example, good, bad, or neutral) rather than the exact value (which may be of much less interest to the spreadsheet reader).

- Use the Icon section to mix and match your icons. For example, you can use icons from different icon sets to represent different values, and you can even choose not to show an icon at all for certain values.
- Choose the range of values that Excel uses for each icon. As with data bars and color scales, you can set these ranges using fixed numbers, percentages, percentiles, or a formula.

The following sections show the two most common ways to fine-tune icon sets.

Giving your icons added significance

As you learned on page 579, Excel splits your values into icon groups automatically. If you want icons with no hassle, this seems pretty convenient. But if you want your icons to actually mean something—for example, you want a checkmark to signify a passing grade or an up-arrow to flag a profitable product—you need to define the scale yourself.

Fortunately, it's an easy job. Begin by applying the icon set you want to your data. Then, edit the rule, as described on page 572. Finally, change the ranges that assign icons to cell values, as shown in Figure 20-12.

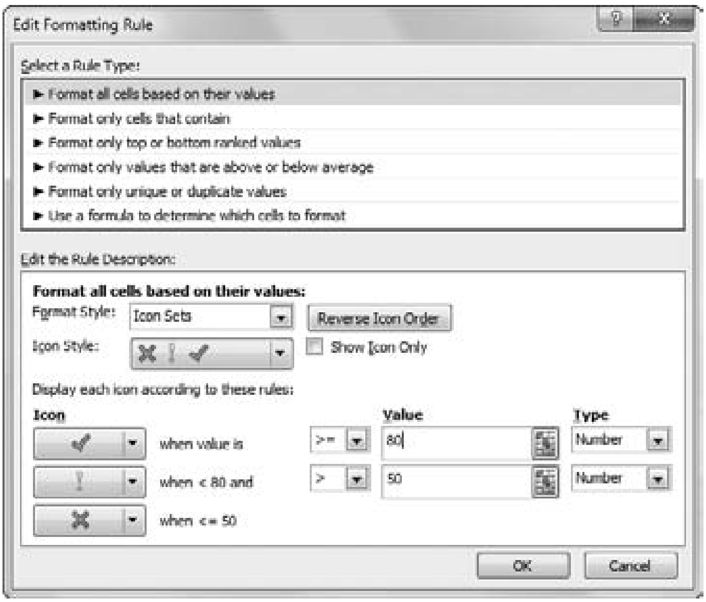

Figure 20-12: *Ordinarily, Excel splits the range of values into equal groups based on equal percentages. By changing the ranges, you can reserve some icons for narrower ranges of numbers. For example, you could decide to use the X icon for all test scores that fall under 50, and keep the green checkmark for the best performing above-80 scorers (as shown here).*

Tip: Usually, it's easiest to get the ranges you want with fixed numbers. That means you should set the Type setting for each range to Number (rather than Percent, Percentile, or Formula).

Hiding icons for some values

Icon sets bring a lot of clutter. If you're not careful, your worksheet can dissolve into a jumble of pictures. One way to control the distraction is to use icons for certain ranges, but not for others. (In fact, design experts suggest that icons work best when no more than *10 percent* of your cells use them.)

For example, consider the worksheet with student scores shown earlier. To keep it clear, you want to show an X icon next to all the failing grades, but you don't want to bother showing any icon next to the others. Figure 20-13 shows the difference.

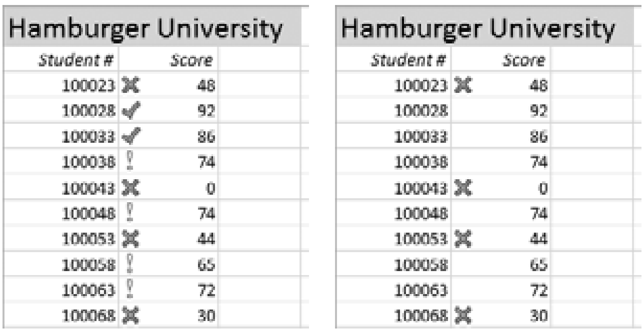

Hamburger University

Student #	Score
100023	48
100028	92
100033	86
100038	74
100043	0
100048	74
100053	44
100058	65
100063	72
100068	30

Hamburger University

Student #	Score
100023	48
100028	92
100033	86
100038	74
100043	0
100048	74
100053	44
100058	65
100063	72
100068	30

Figure 20-13:
Left: Tiny icons show the difference between good marks, bad marks, and really bad marks–if you don't go cross-eyed.

Right: If you need to keep track of the really bad marks, the best bet is to hide all the other icons.

You can accomplish this easily by editing the rule. Look at the Icon section in the bottom-right corner of the Edit Formatting Rule dialog box. There, you can click the drop-down arrow next to an icon to pick something different. Or, choose No Cell Icon to leave the icons out of this range. To create the worksheet shown in Figure 20-13, set the middle and top range to No Cell Icon, while keeping the X icon for the bottom range.

Sparklines

Sparklines are tiny graphics that look like miniature charts, but fit in a single worksheet cell. Compared to charts, they're simplified and stripped down, with no gridlines, borders, legend, or callouts. Interestingly, these limitations make sparklines so much more versatile than their graphing big brothers. While a worksheet might feel overburdened with more than a couple of full-fledged charts, you can easily fill its cells with a dozen sparklines, while still keeping it clear and clean.

As you'll soon see, using sparklines is easy. Using them well isn't too difficult either, as long as you have a solid understanding of their proper role in life. The key detail you should understand is that sparklines are *not* a poor man's chart. They aren't intended to replace, simplify, or compete with Excel's charting feature at all. In fact, sparklines have a closer relation to data bars, color scales, and icon sets—the conditional formatting features you explored in the first half of this chapter.

Much like these features, sparklines are a way to give you an at-a-glance overview of your data and its significance. For example, a properly placed sparkline can highlight plummeting profit, a seasonal spike in orders, or an uncanny winning streak in the company bowling league. They might not give you all the details, but they can alert you to trends, patterns, and dramatic changes, without forcing you to perform a painstaking review of all the numbers.

Currently, sparklines come in just three flavors:

- **Line**. These sparklines chart a series of values with a zig-zagging line that stretches from left to right. It's the same as the line you see in the familiar line chart.
- **Column**. These sparklines create a series of vertical bars, one for each value, with the bigger columns indicating the bigger values. As you've no doubt guessed, it's the sparkline version of the ordinary column chart.
- **Win/Loss**. These sparklines create a series of squares, and each square can be in one of two positions. If the value is positive the square is placed at the top of the cell (win), and if the value is negative the square is placed at the bottom (lose). Unlike the other sparklines, this one simplifies your data. For example, you could use it to take a series of golf scores and show, for each game, whether the golfer beat par.

Figure 20-14 shows all three varieties.

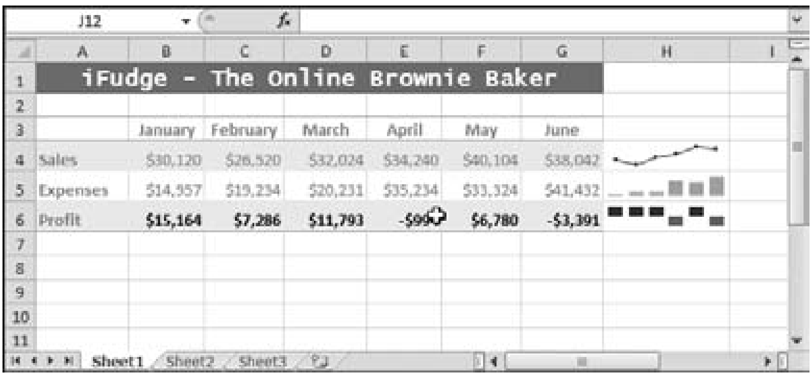

Figure 20-14: *In this worksheet, a line sparkline (cell H4) shows the changing sales of a web-based baker. Underneath, a column sparkline (H5) shows the worrying increase in expenses. (These two types are essentially interchangeable, so you could use a column sparkline for sales and a line sparkline for expenses.) Finally, a win/loss sparkline (H6) shows the profit in each quarter. A profit counts as a win, while a loss is, well, a loss.*

Creating a Sparkline

The easiest way to understand how sparklines work is to create one. Consider the sample spreadsheet shown in Figure 20-15.

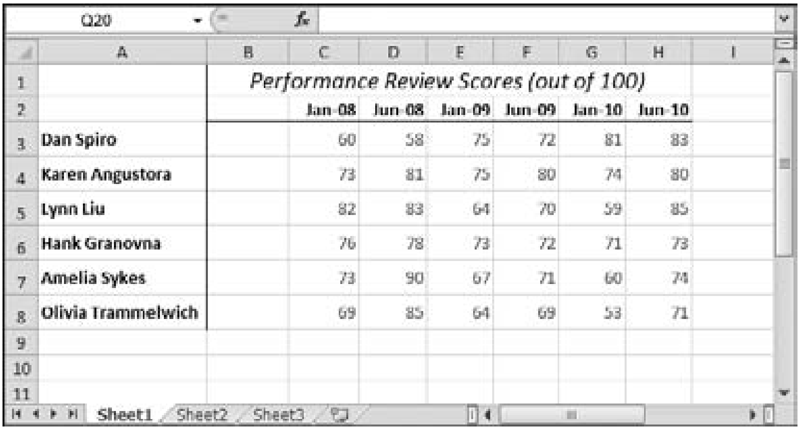

	A	B	C	D	E	F	G	H	I
1		Performance Review Scores (out of 100)							
2			Jan-08	Jun-08	Jan-09	Jun-09	Jan-10	Jun-10	
3	Dan Spiro		60	58	75	72	81	83	
4	Karen Angustora		73	81	75	80	74	80	
5	Lynn Liu		82	83	64	70	59	85	
6	Hank Granovna		76	78	73	72	71	73	
7	Amelia Sykes		73	90	67	71	60	74	
8	Olivia Trammelwich		69	85	64	69	53	71	
9									
10									
11									

Figure 20-15: *This spreadsheet shows a morass of data: the twice-a-year performance review scores for a handful of long-time employees. All the information you could possibly need is here, but it might take some time to dig through it.*

In this example, column B (cells B3:B8) is empty and ready for a set of sparklines. If you're working with your own worksheet, you might need to move some cells around to make space. Remember, each sparkline needs just a single cell, so the example in Figure 20-15 has room for a sparkline next to each employee.

To create a sparkline for this example, follow these steps:

1. **Move to the cell where you want to place the sparkline.**

   Often, this is next to the data you're charting, but it doesn't need to be. In this example, you can start at cell B3.

2. **Go to the Insert→Sparklines section of the ribbon, and then choose one of the three sparkline types.**

   In this example, click Line. The Create Sparklines dialog box appears (Figure 20-16).

3. **Specify the cells for the Data Range.**

   You can type the range in by hand (in this example, it's C3:H3). Or, click the Data Range box, and then drag on the worksheet to select the cells.

4. **Click OK.**

   A sparkline appears in the cell.

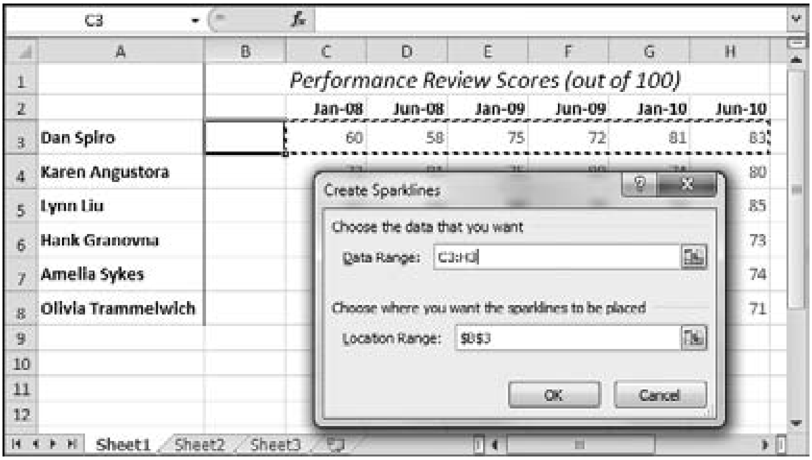

Figure 20-16: *The Create Sparklines dialog box needs two pieces of information. First is the Data Range: the range of cells with the data that'll be used to create the sparkline, Second is the Location Range: the cell where you're placing the newly created sparkline. (This box is automatically filled in with the current cell address.)*

5. **If you want to add a series of sparklines to a table, you can copy the newly added sparkline now.**

   In the current example, that means you can copy the sparkline in cell B3 (press Ctrl+C), select cells B4 through B8, and then paste it (press Ctrl+V). Much as when you copy a formula, each copy uses adjusted cell references. So, while the sparkline in B3 points to the data in C3:H3, the copy in B4 points to C4:H4, and so on. Figure 20-17 shows the result.

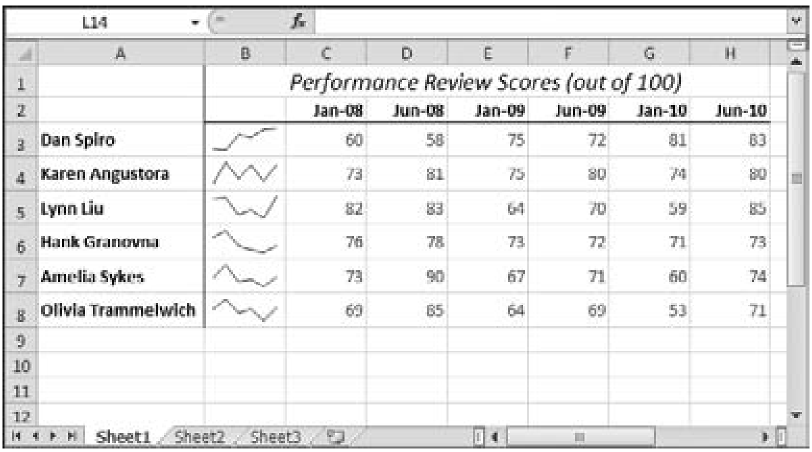

Figure 20-17: *These sparklines don't tell you much about the individual scores, or whether the scores were satisfactory. But they do show you the overall* trend*—that, is how each employees scores have varied over time. For example, you can quickly tell that Dan Spiro is on the up-and-up.*

Tip: If you're in a real hurry, you can create several sparklines in one fell swoop. It's all in how you fill in the values in the Create Sparklines dialog box. In the current example, you supply the full table (cells C3:H8) for the Data Range and the empty cells where you want to create the sparklines (cells B3:B8) for the Location Range. As a side effect, this approach places all the sparklines into a group, as described on page 588.

6. **Fine-tune the width and height of the cells that hold your sparkline.**

   Each sparkline gets the height and width of its containing cell. If it includes a lot of data points, you might want to widen the column to stretch the line out longer. Or you can heighten the row to better see the variations in value between the points.

Tip: To make all the rows taller by the same amount, begin by selecting all the row buttons—in this case, that's rows 3 to 8. Then, right-click the selection, and choose Row Height. You can enter a new value in the Row Height dialog box that appears. Figure 20-16 uses a row height of 20 pixels, rather than the standard 15.

7. **Set the color of your sparkline.**

   Whenever you select a cell that has a sparkline, the Sparkline Tools | Design tab appears in the ribbon. To change the color of the sparkline, choose another color from the Style→Sparkline Color list. If you want to change all your sparklines at once, select all the cells that have sparklines, and then choose a new color.

   You can also use the Sparkline Tools | Design tab to modify the data your sparkline is using (choose Sparkline→Edit Data), change the sparkline type (make another selection in the Type section), show markers (page 590), or show a horizontal axis line (page 588).

It's worth noting that, like data bars, color scales, and icon sets, sparklines still leave room for data in the cell. Most of the time, you'll probably keep your worksheet simple and leave the sparkline cells blank, as in Figure 20-17. But for a different effect, you can try adding formatted text, as shown in Figure 20-18.

Figure 20-18:
A cell that has a sparkline can still accommodate cell data, like a number, formula, or piece of text. In this example, the cell value is right-justified and bottom-aligned, so it doesn't get in the way of the sparkline.

As a side effect of this design, moving to a sparkline cell, and then pressing the Delete key won't remove the sparkline, only the data inside. And if there's no data inside, hitting Delete still won't remove the sparkline. To get rid of it, move to the sparkline cell, and then choose Sparkline Tools | Design→Group→Clear.

Changing the Axis

One of the most important principles of sparklines is this: they can tell you something about your data, but not everything. In order to get the most out of sparklines, you need to decide what it is you want them to tell you.

For example, the employee review example (Figure 20-17) uses sparklines to show *change*. You can tell, at a glance, whether an employee has been on an upward trend or a downward spiral. However, you can't tell if one employee is doing better than another, because each sparkline uses its own distinct scale. This scale is set so that the maximum value (the best performance review score) is plotted at the top of the cell, and the minimum value (the worst score) is plotted at the bottom. As a result, Hank Granovna's slight trend down (from 76 to 71) looks nearly as impressive as Dan Spiro's rise up (from 58 to 83).

To correct this, you can configure a group of sparklines to use the same scale. (This change makes sense only for line-type or column-type sparklines, because win/loss sparklines don't use any scale.)

Here's how to do it:

1. **Select all the sparklines that you want to give the same scale.**

   In the employee review example, that's cells B3 to B8.

2. **Choose Sparkline Tools | Design→Group→Group.**

   Now your sparklines are considered part of one group, and you can set their axis options as a single unit. Whenever you move to a sparkline cell that's part of a group, Excel displays a blue border around all the grouped sparkline cells. You can remove the grouping by choosing Sparkline Tools | Design→Group→Ungroup.

3. **Now it's time to give all the sparklines the same scale.**

   Choose Sparkline Tools | Design→Group→Axis. Under the Vertical Axis Minimum Value heading, choose Same for All Sparklines. Under the Vertical Axis Maximum Value heading, choose Same for All Sparklines.

   This sets the entire group of sparklines to use the same maximum and minimum value for their scale (Figure 20-19). That means the highest score of any employee (90) will become the new maximum, and the lowest score (53) will become the new minimum.

   Alternatively, you can set the scale minimum and maximum manually by choosing the Custom Value option. For example, you could use this technique in the current example to set the scale minimum of 0.

In this example, the same-scale sparklines are the best choice. They prevent people from drawing misleading conclusions about relative performance. However, in many cases same-scale sparklines don't make as much sense.

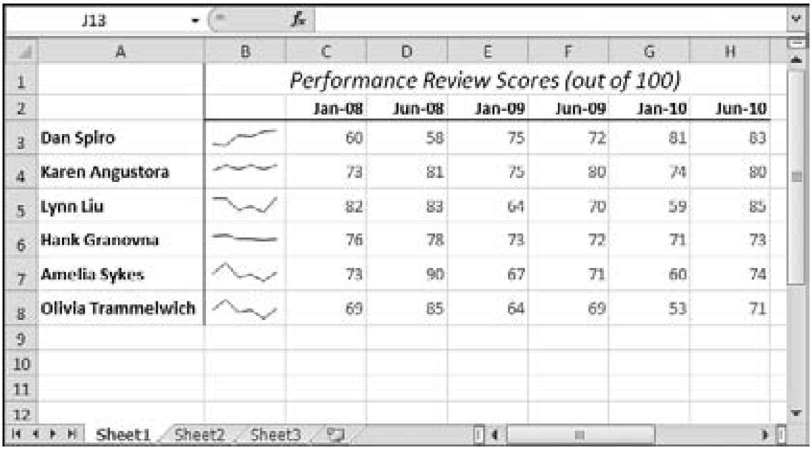

	A	B	C	D	E	F	G	H
1		Performance Review Scores (out of 100)						
2			Jan-08	Jun-08	Jan-09	Jun-09	Jan-10	Jun-10
3	Dan Spiro		60	58	75	72	81	83
4	Karen Angustora		73	81	75	80	74	80
5	Lynn Liu		82	83	64	70	59	85
6	Hank Granovna		76	78	73	72	71	73
7	Amelia Sykes		73	90	67	71	60	74
8	Olivia Trammelwich		69	85	64	69	53	71

Figure 20-19: *Now that the sparklines use the same scale, it's clear that Hank Granovna has nothing to worry about.*

POWER USER'S CLINIC

Even More Axis Options

Although setting the scale is the single most important sparkline change you can make, several other configuration options make a difference in the right situation:

- **Reverse your sparkline.** If you want your sparklines to plot their values in the reverse direction, choose Sparkline Tools | Design→Group→Axis→Plot Data Right-to-Left.
- **Show the horizontal axis.** The horizontal axis is drawn to represent the zero point. To show it, choose Sparkline Tools | Design→Group→Axis, and then select Show Axis under the heading Horizontal Axis Options. Although you won't actually see the horizontal axis unless you have negative values in your sparkline, some Excel gurus get around this by performing an extra calculation, as shown in Figure 20-20.
- **Use a date axis.** Ordinarily, Excel spaces out the values in a sparkline evenly. This makes sense in the employee review example, because the reviews are performed at regular six-month intervals. However, it's not hard to imagine a different sort of situation, where batches of employees are reviewed in different months. In this case, you can space the values out proportionately by date, as long as you have the corresponding dates in your worksheet. For example, the employee review has the dates that correspond to its review scores in cells C3:H8. To use these dates to space out the chart, choose Sparkline Tools | Design→Group→Axis→Date Axis Type, select these cells, and then click OK.
- **Deal with empty cells.** Ordinarily, sparklines leave a gap in the line to represent a missing value. If this isn't what you want, you can tell Excel to treat blank values as zeros, or just draw the sparkline straight through to the next value. To change this option, choose Sparkline Tools | Design→Sparkline→Edit Data→Hidden & Empty Cells. You also see the "Show data in hidden rows and columns" checkbox, which lets you configure whether sparklines include data that's in hidden cells.

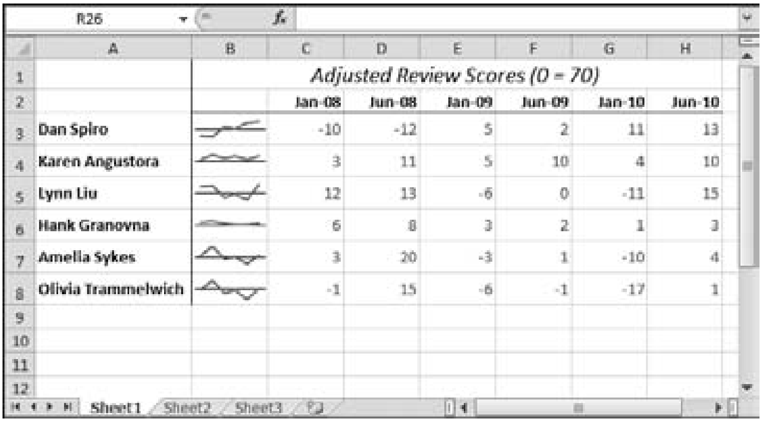

Figure 20-20: *Here, the horizontal axis is a baseline that indicates whether employees measure up to the minimum threshold (a 70 on each performance review). To get the threshold in the right place, you need to subtract 70 from each performance review score. Some Excel hotshots perform this calculation in a separate set of cells, and then hide those cells so they don't clutter the worksheet. Just remember, if you try this trick, you need to explicitly tell your sparkline to show hidden values, by choosing Sparkline Tools | Design→Sparkline→Edit Data→Hidden & Empty Cells, and then switching on the "Show data in hidden rows and columns" checkbox.*

For example, if you're creating a worksheet that tallies the sales performance of different stores for your line of thermal undergarments, you already know that bigger stores sell more underwear. What you really want to see is how sales and profit figures have changed at each location. In this situation, using the same scale will hide the important trend information. And if the stores are dramatically different sizes, the scale will expand so much that some of your sparklines will start looking like horizontal lines. In this situation, relative-scaled sparklines are often more useful than grouped sparklines that share the same scale.

Markers

The line type of sparkline supports *markers*—tiny squares that highlight important data points on the line. The catch is that you need to decide where to add them.

To show markers on your sparkline, select the cells that have the sparklines you want to change, and then go to the Sparkline Tools | Design→Show section of the ribbon. You see a set of checkboxes that let you show different markers. You can show a marker for the greatest data value (select High Point), the smallest (Low Point), all values less than zero (Negative Points), the first value (First Point), the last value (Last Point). Or all values (Markers). Once again, it's a case of deciding what story your sparklines should tell. Figure 20-21 shows a stock chart that uses markers.

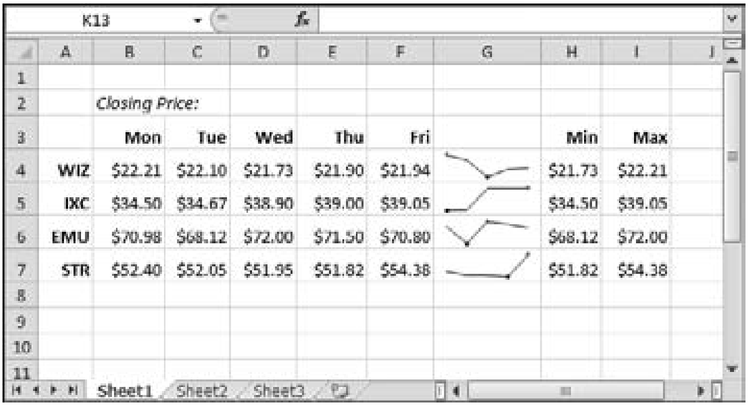

Figure 20-21: *These sparklines show fluctuating stock prices. Each one highlights the highest and lowest value with a marker. The worksheet drives the message home by indicating the exact minimum and maximum numbers in the cells on the right, These cells get their numbers using the MIN() and MAX() formulas you met on page 271.*

Note: You can use markers with column sparklines, but they're applied differently. If you choose to show red markers for the largest and smallest values, for example, the largest and smallest columns will be red in color.

You can use any combination of markers. But if you mix and match, you'll probably want to differentiate markers by color. To do so, head to the Sparkline Tools | Design→Style→Marker Color menu. Here you find different submenus for each marker type, where you can pick the exact color you want. Often, Excel pros like to use red markers for bad news (like negative values or the smallest value) and green markers for good news (like the greatest value). If you don't customize the colors, all the markers will use the same color as your sparkline.

CHAPTER 21

Scenarios and Goal Seeking

Excel formulas are perfect when you need to analyze information from the past, like sales figures, census data, and experimental observations. But life becomes more complicated when you need to plan for the future. Some of the problems you can face are:

- **Missing data**. Excel formulas require exact values—there's no room for guessing or estimates.
- **Multiple possibilities**. An Excel formula can give you only a single answer. If you need to make projections that take best-case and worst-case scenarios into consideration, you have to create multiple versions of the same worksheet.
- **Calculating backward**. Often, when creating plans or projections, you need to *reverse-engineer* your information. In other words, you need to start with the conclusions and work backward to find out what conditions need to happen in order for you to get the desired result. For example, you may ask what score you need on a final exam to get a certain grade in a class, or what interest rate you need to make a profit of $500.

To help you tackle these challenges, Excel includes a few specialized tools that you'll learn about in this chapter. First, you'll see how you can use Excel *scenarios* to consider multiple possibilities within the same worksheet. Next, you'll learn about Excel's *goal-seeking* tool, which provides a simple approach for calculating backward and filling in a missing value. Finally, you'll take a look at one of Excel's most sophisticated add-ins—*Solver*, which is an all-purpose tool for making predictions and optimizing plans and projections.

Using Scenarios

People often use spreadsheets to try to predict the future. For instance, in order to plan for the upcoming year, a company may need to make reasonable assumptions about expenses, sales, and profits. However, part of the planning process is considering different possibilities. A responsible company—one that doesn't want to tumble suddenly into bankruptcy, or sell more products than it can make—uses *contingency planning* to prepare for the best- and worst-case scenarios.

You can certainly use Excel to create a workbook that includes best- and worst-case information on your own. However, to do so, you need to create duplicate copies of your data. The problem with this approach is that duplicated data is difficult to maintain. If you change the profit calculation formula in a worksheet that calculates the best-case scenario, you need to make sure you track down the same calculation in every other worksheet in order to make sure your data remains consistent.

Excel provides an easier solution with its *scenarios* feature. Scenarios help you perform multiple calculations on the same set of data—all within the same worksheet. A ski resort may hope for the best and plan for the worst with the scenarios "Projected 2011 Results for a Long Winter" and "Worst Case 2011 Results for a Short Season." You can choose and view the scenarios separately, or even generate a summary report with the click of a button.

Creating a New Scenario

You can create a scenario for any worksheet. Start by identifying all the cells specific to your scenario. For instance, if you're creating a business plan that estimates revenue from international sales, you may create different scenarios based on different possible currency exchange rates. In this example, the cell containing the exchange rate's value is a scenario-specific cell. That means each scenario uses a different value for this cell. As you change from one scenario to the next, Excel modifies the value in this cell but doesn't change any of the other cells (although if they use formulas based on the currency exchange rate, the program recalculates their results with the new information).

To get a better understanding of how this whole process works, consider the worksheet shown in Figure 21-1.

To try out different possibilities for any formula, you can add scenarios to a worksheet. The following example uses Figure 21-1 as the basis for showing how the scenario tool works.

1. **Select the cells that will change in your scenario.**

   These cells contain the assumptions that your scenario makes. You can select any range of cells you want, and you can hold down the Ctrl key to select a noncontiguous range. In the sales projection example, these cells are C12 (the exchange rate) and, optionally, C13 (the profit margin).

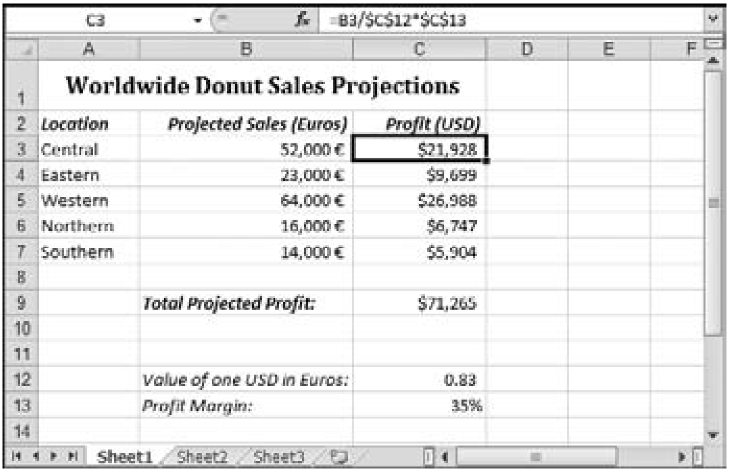

Figure 21-1:
This worksheet shows sales predictions for the coming year for five regional divisions of a company. Assuming that the projected sales are correct, it would be interesting to modify two cells in order to see their effect on the projected profits listed in cells C3 through C9: the exchange rate (shown in C12) and the profit margin (C13). How would fluctuations in either of these values affect the profits listed in column C? You can find the answers to those questions by using Excel's Scenarios feature.

2. **Choose Data→Data Tools→What-If Analysis→Scenario Manager.**

   The Scenario Manager dialog box appears.

3. **Click the Add button to create a new scenario.**

   The Add Scenario dialog box appears, as shown in Figure 21-2.

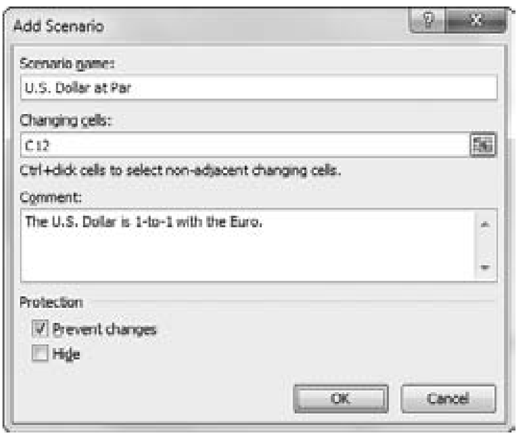

Figure 21-2:
In the Add Scenario dialog box, you choose a name for your scenario. If you like, you can enter a more complete description of the scenario in the Comment text box. You don't need to edit the reference in the "Changing cells" text box, because Excel automatically enters whatever cells are selected on the worksheet. In addition, you can make your scenario tamper-proof with document protection by turning on the two checkboxes at the bottom of the dialog box. These settings work only if you use the worksheet protection features, as explained in Chapter 24.

4. **Enter a descriptive name for the scenario you're creating.**

   If you plan to create a scenario that explores what happens if the U.S. dollar rises against the Euro (thereby causing your profits to plummet), you may use a name like *High U.S. Dollar, Worst Case*, or *Most Likely Case*.

5. **If you need to change the cell references used for this scenario, edit the "Changing cells" text box.**

   Excel automatically fills this text box with references to the cells you selected *before* you started creating the scenario. If you're creating multiple scenarios at once, you may need to change this information. You can edit it directly as long as you remember to separate each cell reference with a comma. Use a value of *C12, C13* if you want to include both cells C12 and C13. You can also point and click your way to success. Just click inside the "Changing cells" text box, and then click the worksheet to select the cells you want (holding down Ctrl to select several cells at once).

Note: You can set your scenario values in two ways. You can enter them in all the changing cells before you create the scenario (because Excel uses the current cell values when you create the scenario), or you can adjust them in the Add Values dialog box when you create the scenario. Both approaches work equally well—it's just a matter of which one you prefer.

6. **Click OK.**

   This action closes the Add Scenario dialog box and opens the Scenario Values dialog box, as shown in Figure 21-3.

Figure 21-3:
This particular scenario includes only a single changing value: the currency exchange rate. To examine what happens if the U.S. dollar rises, this currency exchange rate is adjusted up to 1 from its original value of 0.83.

7. **Set the values of all the changing cells according to your scenario.**

   Excel automatically inserts the current value of each cell in the Scenario Values dialog box. You can adjust these values or, if they're already correct for your scenario, you can keep them without making any changes. In the sales projection example, you can adjust the currency rate to make it higher or lower.

8. **Click OK.**

   This action returns you to the Scenario Manager dialog box, which now shows your newly created scenario in the list.

9. **If you want to create more than one scenario at a time, you can repeat steps 3–8 for each new scenario.**

   There's no limit to the number of different scenarios you can create. Figure 21-4 shows several scenarios in the Scenario Manager dialog box.

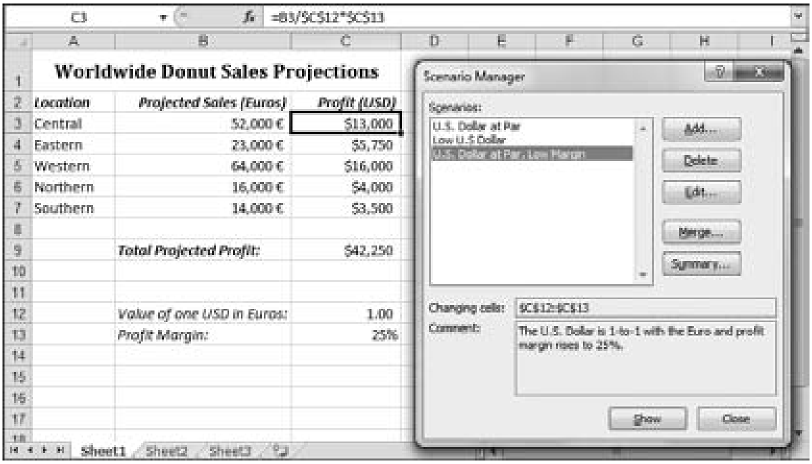

Figure 21-4: *Scenarios are extremely useful and easy to use. Apply a scenario by choosing it from the list shown in the Scenario Manager dialog box, and then clicking Show. The worksheet updates automatically. Using the scenarios in this example, you can see quickly how the total profit falls from $71,265 to $42,250 in the worst-case scenario.*

Note: Different scenarios don't always need to have the same changing cells, although it simplifies life if they do. In the sales worksheet, you can create some scenarios that change the exchange rate value, and others that change both the exchange rate and the profit margin percentage. However, this approach can be a little confusing if you switch rapidly from one scenario to another. If you switch to a scenario that updates both cells, and then switch to a scenario that updates just one of the two cells, the other cell still has the value from the previous scenario–which may not be what you want.

Managing Scenarios

Once you've created your scenarios, you'll probably want to put them to work to compare the different possible cases. You can switch from one scenario to another using the buttons in the Scenario Manager dialog box. (If the Scenario Manager dialog box isn't visible, choose Data→Data Tools→What-If Analysis→Scenario Manager.)

Here are the tasks you can perform in the Scenario Manager:

- **To switch from one scenario to another**, select the scenario you want to view in the list, and then click Show. Excel immediately updates the changed cells (and any cells that reference them).
- **To change a scenario's assumptions**, select it, and then click Edit. The Edit Scenario dialog box that appears looks exactly like the Add Scenario dialog box. It lets you change the name, modify the changing cells, and edit the description. Click OK to move ahead and adjust the actual values for the changing cells.

- **To remove a scenario completely**, select it in the list, and then click Delete.
- **To return to your worksheet**, click Close at any time. The values from the scenario you applied the previous time remain in effect.

Creating a Summary Report

Scenarios are great for exploring different possibilities, but you're still limited to viewing one scenario at a time. If you'd rather have an at-a-glance look at *all* the scenarios you've defined, you can generate an automatic summary report.

To create a summary report, follow these steps:

1. **Choose Data→Data Tools→What-If Analysis→Scenario Manager.**

   The Scenario Manager dialog box appears.

2. **Click Summary.**

   The Scenario Summary dialog box appears, as shown in Figure 21-5.

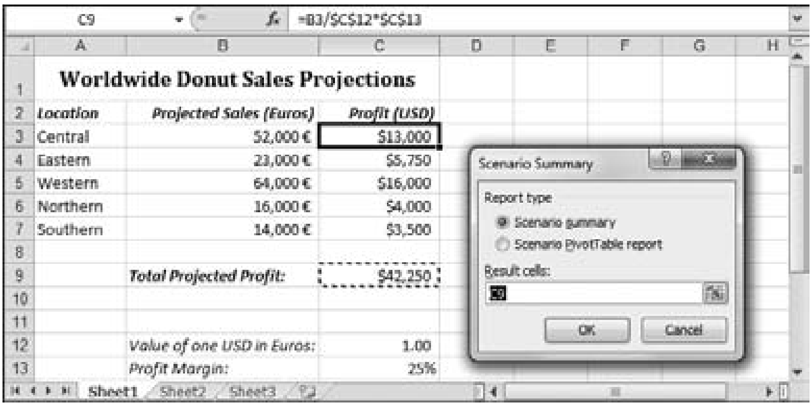

Figure 21-5: *Once you pick the report type, you need to specify (in the "Result cells" box) which cells to display in the summary report. Your choices include cells C3 to C7 and C9 (since these cells have values that change depending on the scenario you're considering). You can include all of these cells, but you may just be interested in the final total in C9.*

3. **Choose the type of summary you want. In most cases, the first option ("Scenario summary") is what you want.**

   You can also create a summary *pivot table* instead of a summary report by selecting "Scenario PivotTable report". For more information about pivot tables, which are an advanced feature for summarizing large amounts of data, check out Chapter 22.

4. **In the "Result cells" text box, specify the cells you want to include in the report.**

   These cells fluctuate based on the scenario values you've used. When Excel creates the summary, it includes all the values you've created for the changing cells, and then shows the resulting value for each of the result cells.

5. **Click OK to create the summary report.**

   The summary report always opens in a new worksheet named Scenario Summary. Excel automatically formats the scenario summary to be readable and gives it basic grouping (Chapter 15) that lets you collapse the summary to show just the portions that interest you. Figure 21-6 shows an example.

Note: Once Excel creates it, the summary report is completely independent from your data. If you change any of the scenario values or calculations, the summary report doesn't update (although you can recreate it using the same sequence of steps).

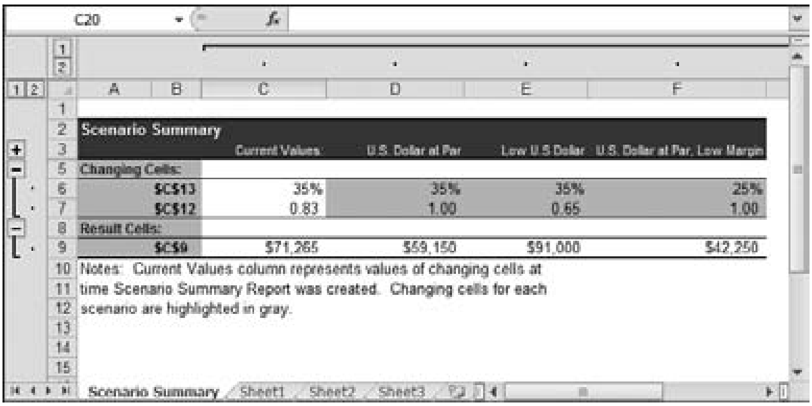

Figure 21-6: *This report compares different scenarios for the sales projection worksheet. Rows 6 and 7 show the scenario-specific cells. Row 9 gives the bottom line–the calculated result that appears in cell C9 based on the scenario-specific values.*

Using Goal Seek

There are some problems that formulas just can't crack. One of these is solving questions that need trial-and-error guesswork.

In most Excel spreadsheets, you begin with a set of data and use formulas to analyze that data and calculate some conclusions. When you create plans and projections, however, you sometimes invert this process and start with the conclusions. You may begin with a profit projection and calculate "backward" to find out the sales you need to make. Or, you may calculate the rate an investment needs to increase in value in order to meet a set target.

POWER USERS' CLINIC

Scenarios to Go

When you first create a scenario, it's attached to a single worksheet. However, the Scenario Manager also lets you use a little-known trick to *merge* scenarios. This technique lets you transplant your scenarios (and their accompanying cell values) into another worksheet in the same workbook, or another worksheet in a different workbook.

Merging scenarios isn't often useful, but it does come in handy if you need to perform the same analysis with different worksheets that contain the same arrangement of data. If you have several sales projections worksheets, and the exchange rate and profit margin cells are in the same place in every worksheet, you could copy the scenarios to each worksheet.

To transfer scenarios, make sure both worksheets are open (the one containing the scenario and the one where you want to copy it to). Then go to the worksheet where you want to place the copied scenarios, choose Data→Data Tools→What-If Analysis→Scenario Manager, and then click Merge in the Scenario Manager dialog box. You can choose the source workbook and worksheet from the lists that Excel provides. Click OK to complete the deal.

When you merge a scenario into a new worksheet, Excel copies the scenario names and the scenario values (the values for all the changed cells). Of course, the merged scenario isn't of much use unless the new worksheet has *exactly* the same structure as the original. Otherwise, when you select the scenario, Excel inserts the changed values in the wrong places.

Goal seeking is a fairly simple tool in Excel that can help you answer these sorts of questions, provided your problem meets the following guidelines:

- **There's only one result cell and this cell contains a formula**. You want to make sure the result of this calculation meets a target you specify.
- **There's only one variable cell**. This is the value that Excel adjusts to meet your target in the result cell.
- **There's a valid solution**. If you ask Excel how many years it'll take to pay off a $100,000 loan by making 25-cent monthly payments, you'll end up with an answer—a negative number of years—that makes no sense (since your loan is accumulating interest faster than you're making payments).

To better understand goal seeking, it helps to think about a simple example. Consider the worksheet shown in Figure 21-7, which shows a list of student grades. Grace Dewitt has the chance to resubmit her assignment, and she wants to determine what grade she needs on the assignment to increase her final grade from 72 to 80 out of a possible 100.

To use goal seeking in this scenario, just follow these steps:

1. **Choose Data→Data Tools→What-If Analysis→Goal Seek.**

   The Goal Seek dialog box appears (Figure 21-8).

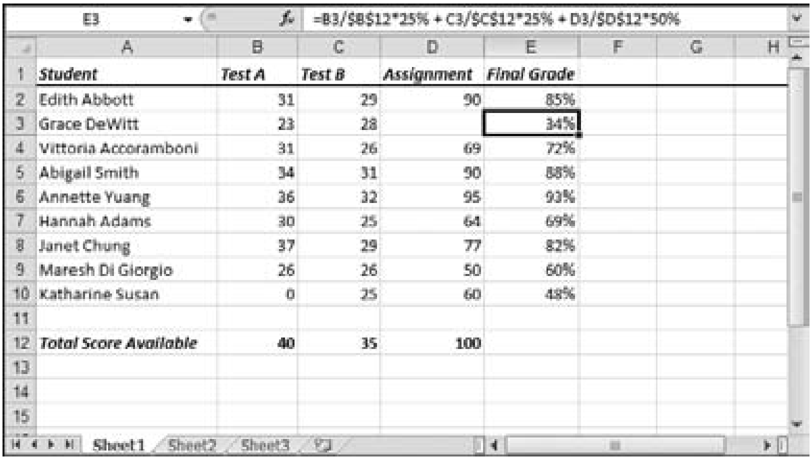

Figure 21-7: *Grace Dewitt is trying to determine the grade she needs on an assignment to achieve a final grade of 80 percent (which is the target). In this case, the result cell (the one she wants to optimize) is E3. The variable cell (the one that Excel will adjust) is D3.*

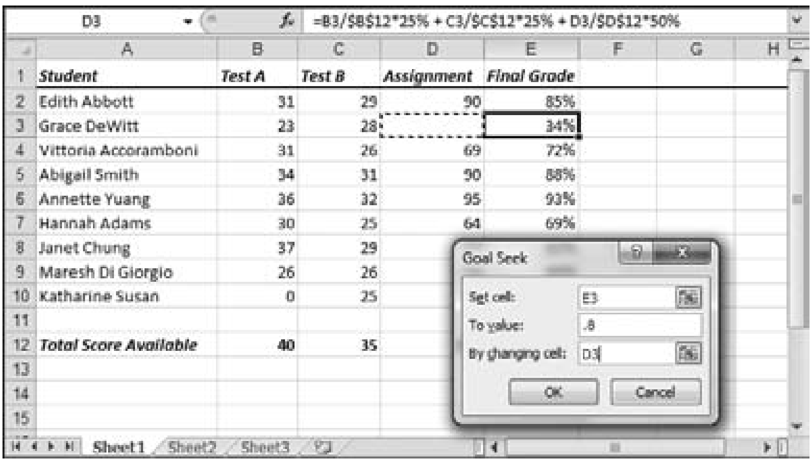

Figure 21-8: *This example shows the Goal Seek dialog box filled with the values needed to search for the minimum grade Grace needs to get on the assignment to increase her final grade to 80. Note that you have to enter the grade as a decimal value (0.80 instead of 80). That's because the cells in column E are formatted to use the percentage number style.*

2. **In the "Set cell" text box, enter the cell you want Excel to optimize.**

   In the student grade example, this cell is cell E3, which contains Grace's final grade.

3. **In the "To value" text box, enter your target.**

   Grace is seeking a final grade of 80 percent, so the target is *0.80*.

4. **In the "By changing cell" text box, enter the cell that Excel should modify to achieve the target.**

   Cell D3, which contains the assignment grade, is the value that Excel modifies to determine the grade needed on the assignment in order to boost her final grade to 80.

5. **Click OK to start goal seeking.**

   The Goal Seek Status dialog box appears, and Excel cycles through a series of trial values, incrementing the assignment grade and calculating the resulting value.

   If Excel can't find the answer after a series of attempts, the Goal Seek Status dialog box gives you the option to keep trying by clicking Step. If Excel stalls on a long analysis, you can stop it by clicking Pause, although this move's rarely necessary. In most cases, the goal-seeking process finishes before you can even take your hand off the mouse to reach for a nearby coffee mug.

   In a short amount of time, Excel discovers that an assignment grade of 91 (out of 100) raises the final grade to 80 percent. At this point, it stops the goal seeking process and shows you the answer.

6. **When an answer appears, click OK to accept the change and return to your worksheet, or click Cancel to return to the original version of your data.**

   Figure 21-9 shows the adjusted worksheet.

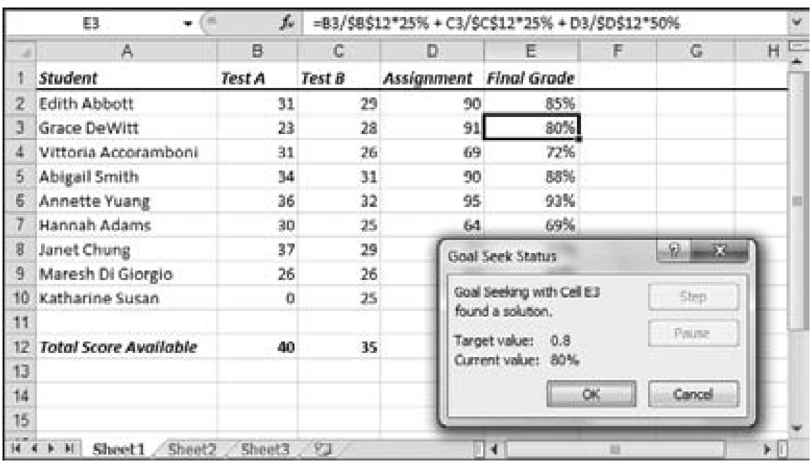

Figure 21-9: *Usually, goal seeking unfolds in a flash. In this example, Excel has met the target by setting the assignment grade to 91, as shown in cell D3. You can click OK to accept the change permanently, or Cancel to return the cell to its previous value.*

Goal Seeking with Complex Equations

In the student grade example, you don't necessarily need to use goal seeking to get the answer you're looking for. You could get the same information by rewriting the equation. To understand how this approach works, it helps to take a closer look at the formula that calculates the final grade:

```
=(B3/B12)*25% + (C3/C12)*25% + (D3/D12)*50%
```

In this formula, the two tests are each worth 25 percent of the final grade, and the assignment is worth the remaining 50 percent. However, using a dash of high-school math, you can rearrange this formula to find the grade you need on the assignment to get an 80 percent overall grade (assuming you already know your grades on the two tests). Here's the answer:

```
=(80% - (B3/B12)*25% - (C3/C12)*25%)/50%*D12
```

This formula looks a little intimidating at first glance, but it's really not too difficult to understand. First of all, you start with the 80-percent final grade you're trying to achieve. You subtract the two test scores to find the value you need from the assignment. Finally, you divide that value by 50 percent (the weighting of the assignment) and multiply it by D12 (the total score available in the assignment). The end result is 91.25.

Based on this example, you may assume that goal seeking is just a tool for quickly calculating a number that you could get on your own by rewriting one or more formulas. Although that's often the case, there are some types of formulas that you *can't* reorganize to provide an answer. These formulas often have exponents and use the value in the variable cell more than once. You can find many examples of this sort of equation in the scientific world.

One common example is the *quadratic function*, which shows up in lots of different scientific modeling and engineering formulas. A typical quadratic formula looks like this:

```
y = x2 + 5x + 10
```

In this example, it's easy to calculate the result (y) provided you know the input value (x). However, it's much more difficult to perform the reverse task—determining x if all you know is y. Try as you might, there's no way to rearrange the formula to solve the problem. In this case, goal seeking becomes very useful. It uses a trial-and-error approach that can usually get the answer you need. Excel tries different values of x in succession trying to get closer and closer to the target result.

To put the problem into Excel's formula format, assume that cell B1 contains the x value. Here's what the formula would look like:

```
=B1^2 + 5*B1 + 10
```

Figure 21-10 shows an example worksheet that puts this function to the test. It's easy to see that an input value of 20 produces a result of 510. But what input value do you need to get a result of 1,000? In this example, goal seeking is the only easy way to find the answer (which, incidentally, is 29.06343).

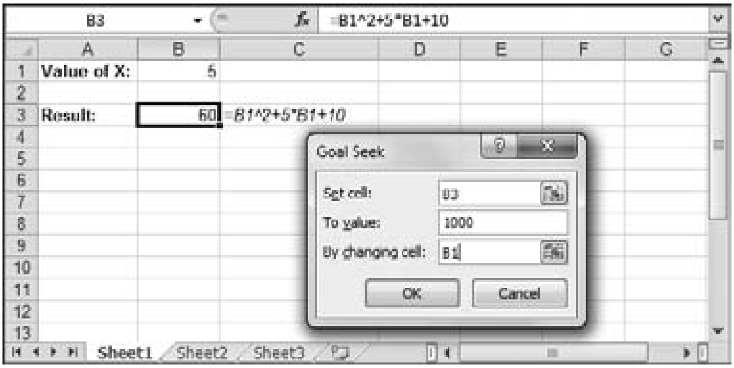

***Figure 21-10:** Figuring out what the value of B1 needs to be to generate a formula result of 1,000 is extremely difficult to resolve without using goal seeking. Figure 21-11 shows the answer.*

Note: It's often difficult to reverse-engineer a formula in financial calculations, like those that calculate how an investment appreciates in value. However, as you saw in Chapter 10, Excel includes financial functions that can calculate in both directions: the final value of an investment based on a given interest rate, or the required interest rate or investment length based on your desired goal. Of course, you can still use goal seeking with financial functions to find an answer quickly without writing a new formula.

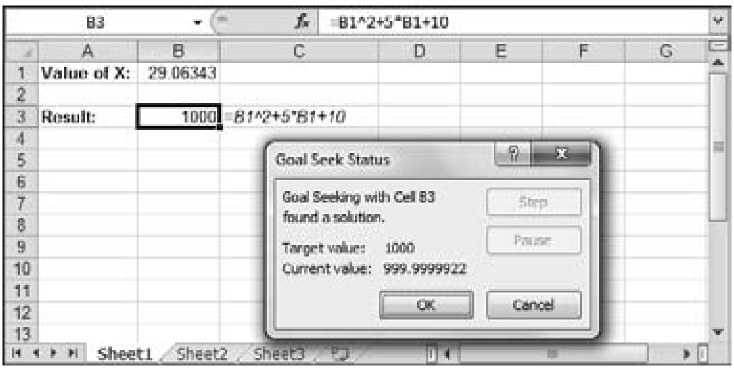

***Figure 21-11:** As you can see in the Goal Seek Status dialog box, the answer isn't perfect (it's 999.9999922 instead of 1000), but it's close enough for Excel to stop searching. When performing goal seeking, Excel doesn't need to match your target exactly.*

Solver

Goal seeking works well for simple problems. But it does have a few limitations:

- **Excel doesn't recognize what you may think of as common-sense limits in your data**. In the earlier student grade example, if you were looking to figure out the assignment grade required to get a final grade of 100, Excel would tell you you'd need a score of 113—even though a grade over 100 percent clearly isn't possible.

- **Excel adjusts only one cell**. There's no way to ask the goal-seeking tool, for instance, to predict the minimum grade combination you'd need on the two tests to get a particular final grade.

Solver is an Excel feature that goes several steps further than goal seeking. It uses the same basic trial-and-error approach (known to scientific types as an *iterative* approach), but it's dramatically more intelligent than goal seeking. In fact, it *needs* to be much more intelligent because it tackles much more complicated problems, including scenarios that have multiple changing cells, additional rules, and subtle relationships. Solver can (with a fair bit of work) suggest an optimal investment portfolio based on a desired rate of return and risk threshold. Incidentally, Solver is an add-in that you have to turn on the first time you use it; page 606 tells you how.

Note: Although it comes as part of the Excel package, the Solver add-in isn't a Microsoft product. Instead, it's provided by a company called Frontline Systems, which sells more enhanced versions of the Solver tool. To find out about these or to look through a number of interesting (and complicated) sample worksheets that show Solver in action, go to *www.solver.com*.

Understanding Solver

Every problem that Solver is capable of solving contains the same basic ingredients. Altogether, these ingredients make up a *Solver model*. They include:

- **Objective cell**. The objective cell (also known as the *target cell*) is the value you want to optimize. As with goal seeking, you work with only one objective cell at a time. However, you don't have to set a specific numeric goal for the objective cell. Instead, you can ask Solver to just make the target value as large or small as possible (without violating the other rules of the model).
- **Changing cells**. These cells contain the values (also known as *decision variables*) that Solver modifies in order to reach the target you want. Unlike goal seeking, you can designate multiple changing cells. When you do, Solver tries to adjust them all at once and gauge the most significant factors. With multiple changing cells, there's a definite possibility for multiple solutions (there's more than one combination of grades that can lead to a final grade of 80 percent). However, Solver stops when it finds the first matching solution.
- **Unchanging cells**. Unchanging cells can contain values that affect the calculation that's used in the objective cell. However, Solver never changes these cells.

Tip: Sometimes, it's useful to combine Solver with Excel scenarios. You may want to create multiple scenarios, each of which has different values for the unchanging cells.

- **Constraints**. Constraints are rules that you use to restrict possible solutions. Usually, you apply constraints to changing cell values. People often add constraints that set a maximum and minimum allowed value for each changing cell. You can also use constraints to restrict the target value. You could ask Solver to find the maximum value for the objective cell, but add a constraint indicating that this value can't fall in a certain range.

Note: Solver doesn't enforce constraints absolutely. Instead, it permits a certain *tolerance range*, depending on the configuration options that you've enacted. If you specify that a certain changing cell needs to be less than or equal to 100, Solver allows a value of 100.0000001, because it falls (just barely) in the standard tolerance range of 0.0000001. You'll learn how to adjust the tolerance setting later in this chapter.

You can probably already see that the service provided by Solver is similar to goal seeking, but with a few significant improvements. In the next section, you'll learn how to create a Solver model.

Defining a Problem in Solver

To learn how Solver works, take another look at the student grade example shown in Figure 21-7. This time, consider what happens when you add a new student named Katharine Susan. This student missed the first test and, as a result, received a score of 0 on it. She wants to know what grades in the second test and assignment she needs to guarantee her a final grade of 70. Excel's goal seeking tool can't tackle this problem. Not only is it unable to change two cells at once, but it doesn't heed any upper limits, so it cheerily recommends scores greater than 100 percent. Instead, you need Solver.

To create a Solver model for this problem, follow these steps:

1. **If this is the first time you're using Solver, you need to switch it on. Choose File→Options.**

   If this isn't the first time you're using Solver, skip to step 5.

2. **In the Excel Options window, pick the Add-Ins category.**

   This shows a list of all your currently active add-ins, and those that are installed but not active.

3. **In the Manage list at the bottom of the window, choose Excel Add-Ins, and then click Go.**

   The Add-Ins window appears, which lets you switch your add-ins on or off.

4. **Turn on the checkbox next to the Solver Add-in, and then click OK.**

   Once you have the add-in switched on, you'll find an additional entry in the Data tab of the ribbon.

5. **Launch Solver by choosing Data→Analysis→Solver.**

   The Solver Parameters dialog box appears.

6. **In the Set Objective text box, enter the cell you want Solver to optimize.**

   In the student grade example, this is cell E10, which contains Katharine Susan's final grade (see Figure 21-12).

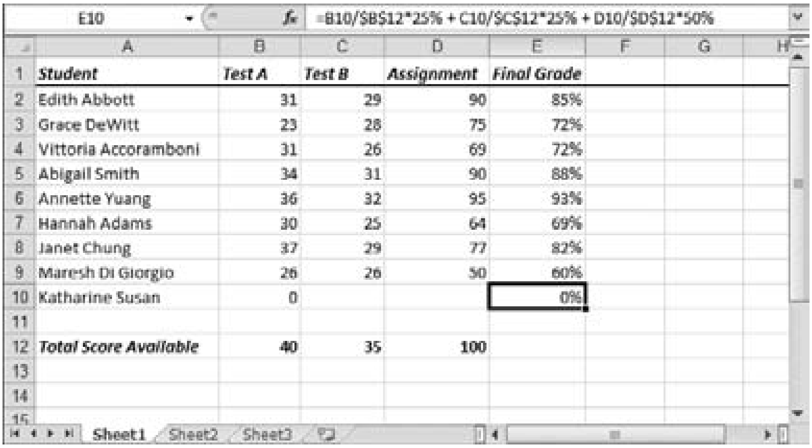

E10 =B10/$B$12*25% + C10/$C$12*25% + D10/$D$12*50%

	A	B	C	D	E
1	*Student*	*Test A*	*Test B*	*Assignment*	*Final Grade*
2	Edith Abbott	31	29	90	85%
3	Grace DeWitt	23	28	75	72%
4	Vittoria Accoramboni	31	26	69	72%
5	Abigail Smith	34	31	90	88%
6	Annette Yuang	36	32	95	93%
7	Hannah Adams	30	25	64	69%
8	Janet Chung	37	29	77	82%
9	Maresh Di Giorgio	26	26	50	60%
10	Katharine Susan	0			0%
11					
12	*Total Score Available*	40	35	100	

Sheet1 Sheet2 Sheet3

Figure 21-12: *The Solver tool helps calculate the values needed in C10 and D10 to boost E10 to 70 percent. A constraint ensures that the "Test B" grade isn't increased above the maximum grade available.*

7. **Choose the type of optimization you want by selecting one of the three To options: Max, Min, or Value Of. If you choose Value Of to match a set value, you have to also enter the value in the adjacent text box.**

   Max attempts to make the value as large as possible. Min attempts to make it as small as possible. Value Of attempts to get the cell to the value you indicate.

   In the student grade example, Katharine is trying to scrape past with a 70 percent grade. So to have Solver figure out the solution, you need to choose the Value Of option, and then set the target to 0.70.

8. **In the By Changing Variable Cells text box, enter the cell references that Excel should modify to achieve the target.**

   You can enter these cell references by hand (separating each reference with a comma), or you can click your worksheet to select the appropriate cells. In the student grade example, the changing cells are C10 and D10, which represent the two grades that are still up for grabs.

   Now, you need to indicate the restrictions that Solver follows when it calculates its solution.

9. **If you don't need any constraints, skip to step 12. Otherwise, click the Add button to add a constraint.**

   The Add Constraint dialog box appears (as shown in Figure 21-13). In the student grade example, you can use constraints to prevent scores that are above the maximum possible grade (or less than 0).

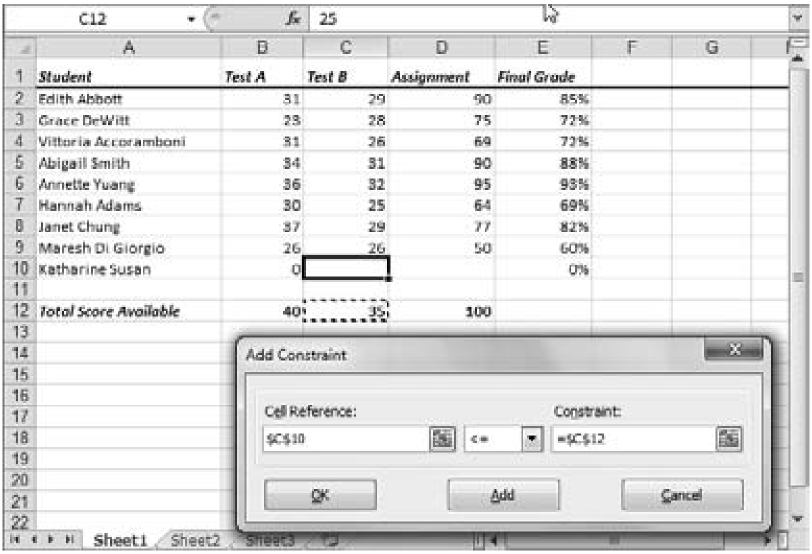

Figure 21-13: *This constraint lays down the law: the test score (which will be set in cell C10) can't be increased above the maximum grade available (which is found in cell C12).*

10. **Every constraint consists of a cell reference, an operator that's used to test the cell reference, and the constraint value that the cell needs to match. Enter this information, and then click OK to add the constraint.**

    A basic constraint compares a cell to a fixed value or another cell. You can perform an equal to comparison (=), greater than or equal to comparison (>=), or a less than or equal to comparison (<=). Figure 21-13 shows one of the constraints used in the student grade example. This constraint uses the less than or equal to (<=) operator to ensure that cell C10 (one of the changing values) doesn't exceed C12. It's not necessary to set a minimum value in this case because Solver can't solve the problem by lowering the grades.

    Also, three special types of constraint operators show up in the list as the words *int*, *bin*, and *dif*. If you apply an int constraint to a cell, Solver allows that cell to contain only whole integer values (with no digits on the right side of the decimal). If you apply a bin constraint, Solver permits only binary values. If you use dif, Solver makes sure that each cell in the group has a different integer value, so no two cells are alike. In all of these cases, all you need to specify is the cell reference and the operator. You don't need to fill in the third part (the Constraint box) because it doesn't apply.

11. **If you have another constraint to add, return to step 10. Otherwise, continue with the next step.**

    As you've probably already guessed, the Change and Delete buttons let you edit or remove the selected constraint. The Reset All button underneath is slightly less obvious. Click it to delete all your constraints and return all the Solver options to their original values.

    You might also be interested in the Make Unconstrained Variables Non-Negative checkbox, which sits just under the list of constraints. It acts as a sort of super-constraint. When it's turned on (which it is initially), every variable cell that doesn't already have a lower-limit constraint is limited to positive values. In other words, it's as if you added a "has to be greater than or equal to 0" constraint to each of these cells. However, if you specifically add a lower-limit constraint for a variable cell (for example, specifying that the variable has to be greater than –10), Solver assumes you're taking control and it *will* allow negative values in that cell, as long as they satisfy your constraint.

12. **Optionally, choose an algorithm from the Solving Method list.**

    The Solving Method gives you a peek into Solver's complex internals. It lets you tell Solver what sort of strategy it should use when it tackles your problem. Unless you're a Solver expert, you probably won't touch this setting. But if you want to get deeper into the details of how Solver works, jump ahead to page 617.

    Figure 21-14 shows the completed Solver model.

13. **Click Solve to put Solver to work.**

    When you click Solve, the Solver Parameters dialog box disappears, and the status bar indicates the number of trial solutions in progress. You can interrupt the process at any time by pressing the Esc key. If you do, the Show Trial Solution dialog box appears, which lets you abandon your attempt altogether by clicking Stop, or resume the trial process by clicking Continue.

    When Solver finishes its work, it shows the Solver Results dialog box, which indicates whether it found a solution, and what that solution is. Here are the most common messages:

    - **Solver found a solution** means Solver met your objective and satisfied your conditions. Hooray! Figure 21-15 shows the solution that Solver found for Katharine Susan.
    - **Solver could not find a feasible solution** means that no matter how Solver tweaked the variable cells, it couldn't get to the target value you want.

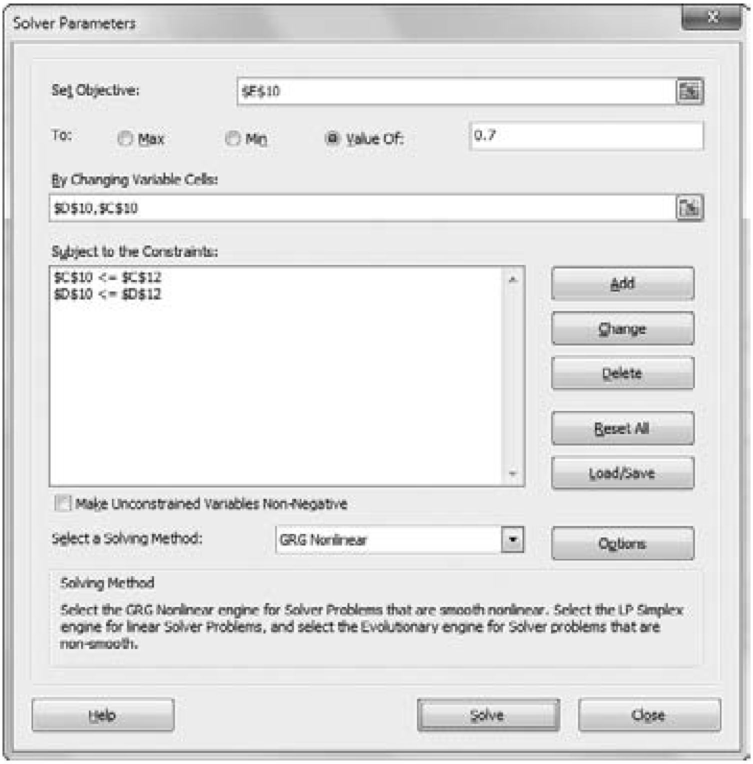

Figure 21-14: *This completed Solver Parameters dialog box is ready to find out how Katharine Susan can eke out a final grade of 70. It allows for two changing cells (the test and assignment grades in cells D10 and C10), and it sets two constraints to prevent solutions that don't correspond to real-world possibilities—in this case, test scores above 100%.*

- **The Objective cell values do not converge** happens only when you're using the Max or Min option (see step 7) to make the objective cell as big or small as possible. In this situation, Solver may find continuing increasing or decreasing your variable cells to ridiculous values, while continually getting a better and better solution, until it finally gives up. Usually, the problem is that you forget to set a constraint to limit one or more of your decision variables.
- **Solver encountered an error value** means Solver started changing your decision variables and an error value popped up in a variable cell or in the objective cell. (For example, maybe one of the formulas divides by zero.) Once again, you probably need to set a constraint to prohibit out-of-bounds values.

Tip: If you didn't get a successful solution, switch on the "Return to Solver Parameters dialog box", and then click OK to return to the familiar Solver window that you started at in step 5.

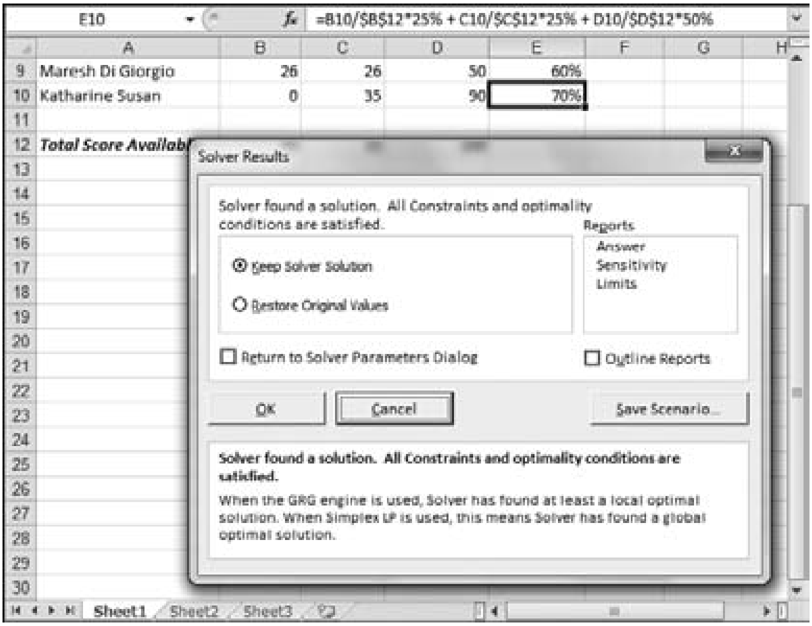

Figure 21-15: *There's hope for Katharine–provided she does well on the second test and the assignment. Solver identifies a solution by increasing the test score to its maximum (35 points, which translates to 100 percent), and then by incrementing the assignment grade to 90 (out of a possible 100).*

POWER USERS' CLINIC

Store the Solver Results in a Scenario

Instead of applying the values that Solver calculates, you can store them so you can refer to them later. Solver helps you out with a nifty feature that stores Solver results in a new scenario.

To use this feature, just click the Save Scenario button in the Solver Results dialog box after the Solver calculation is finished. Enter the name for the scenario you want to create, and then click OK. Excel creates the scenario and stores the value of each changing cell. To apply the scenario values later on, select Data→Data Tools→What-If Analysis→Scenario Manager, find the scenario in the list, and then click Show.

14. **Choose Keep Solver Solution if you want to retain the changed values.**

    Or, to undo Solver's work and go back to the previous version, choose Restore Original Values.

15. **Optionally, in the Reports list box, choose a Solver report.**

    Click a report to highlight it. You can choose just one report or as many as you want. These reports show various details about how Solver calculated its answer.

Solver's reports include Answer, Sensitivity, and Limits. Each report opens in a separate worksheet when you click OK. Most people won't find these reports very interesting at all because they contain a bare minimum of data, and what they do include are dry statistical calculations. The Solver reports aren't nearly as useful as the scenario summaries described earlier.

16. **Finally, click OK to return to your worksheet.**

More Advanced Solver Problems

The student grade example shows only a hint of Solver's true analytical muscle. To see Solver tackle a more interesting problem, you can add additional constraints. You can tell Solver to make sure that the assignment grade is always higher than the test score, or that both the assignment and the test end up having equal percentage values. Figure 21-16 shows the latter example.

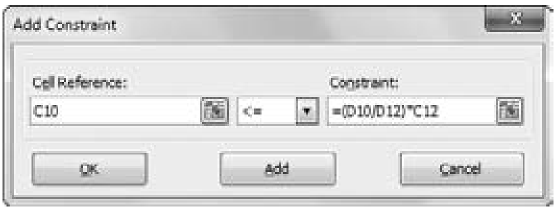

Figure 21-16:
This constraint calculates the percentage on the assignment (D10/D12), and multiplies it by the total available test score (C12). The constraint forces a solution in which Katharine has an equal percentage grade on the test and on the assignment. The solution is 33 out of 35 for the test and 93 out of 100 for the assignment, both of which equal 93 percent.

Note: Solver can't optimize more than one value at once. You can't tell Solver to calculate the lowest possible grade on test B that can, in combination with the assignment score, still result in a final grade of 80. However, you can approximate many cases like this with a crafty use of constraints.

Of course, not all problems have a solution. If you repeat the same problem but try to end up with a final grade of 90 for Katharine, you'll find that it's just not possible. Solver increases both changing cells to their maximum allowed values, and then shows the Solver Results dialog box with a message informing you that there's no feasible solution.

Solver is tailor-made for more complicated math problems, like the optimization problems that calculus teachers use to torture their students. Here's a common mind-bending example:

> You're standing at the edge of a slow-moving river which is one mile wide. You're hungry, and you need to return to your campground (which is on the opposite side of the river, and one mile upstream) to get a ham sandwich. To get it, you have to first swim across the river to any point on the opposite bank, and then walk the rest of the way. You can swim at two miles per hour and walk at three miles per hour, and you want to get your sandwich as quickly as possible. What route will take the least amount of time?

As Figure 21-17 shows, the answer isn't obvious.

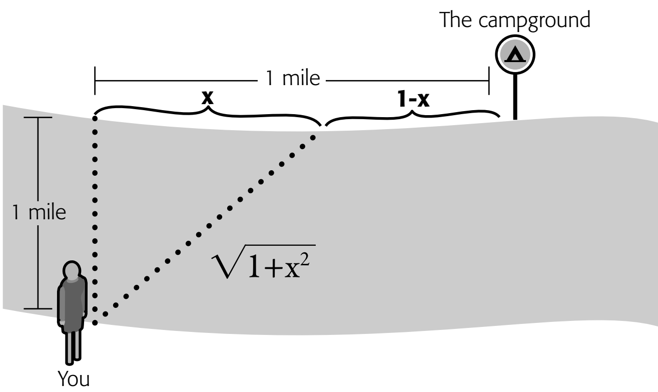

Figure 21-17: *If you swim straight across the river, you have a long one-mile walk to get the campground. If you swim diagonally to the campsite, you can avoid the walk, but you have to travel a much farther distance in the water. Assuming that* x *measures the distance upstream of the point to which you'll swim, some basic algebra gives you the total swimming distance and walking distance.*

To solve this problem, you need to find the right blend of swimming and walking. Math-heads can use calculus to get the answer. Excel lovers have another choice: they can plug all the equations into an Excel worksheet (Figure 21-18) and let Solver do the thinking (Figure 21-19).

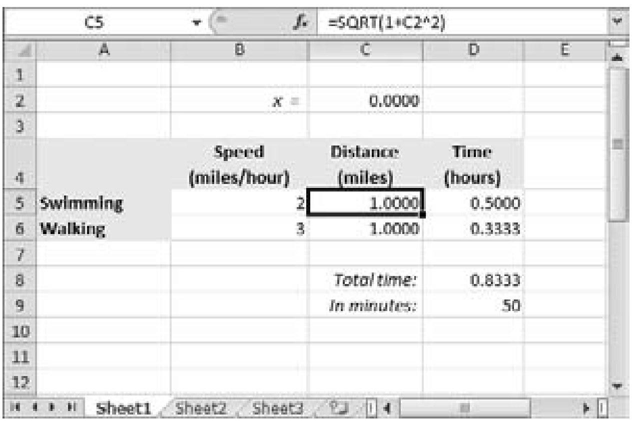

	A	B	C	D	E
1					
2		*x =*	0.0000		
3					
4		Speed (miles/hour)	Distance (miles)	Time (hours)	
5	Swimming	2	1.0000	0.5000	
6	Walking	3	1.0000	0.3333	
7					
8			*Total time:*	0.8333	
9			*In minutes:*	50	
10					
11					
12					

Figure 21-18: *The problem tells you the speed on land and in the water. The diagram measures the swimming and walking distances. With these values, you can create formulas that calculate the total time the trip takes. (To calculate the time, you divide the distance by the speed, using the formula* time = distance/speed.*) Here, you see the results of setting x to 0, swimming straight across the river, and then walking a mile to the campground. But you need Solver to find the best x value—in other words, how far upstream you should swim to get the fastest route.*

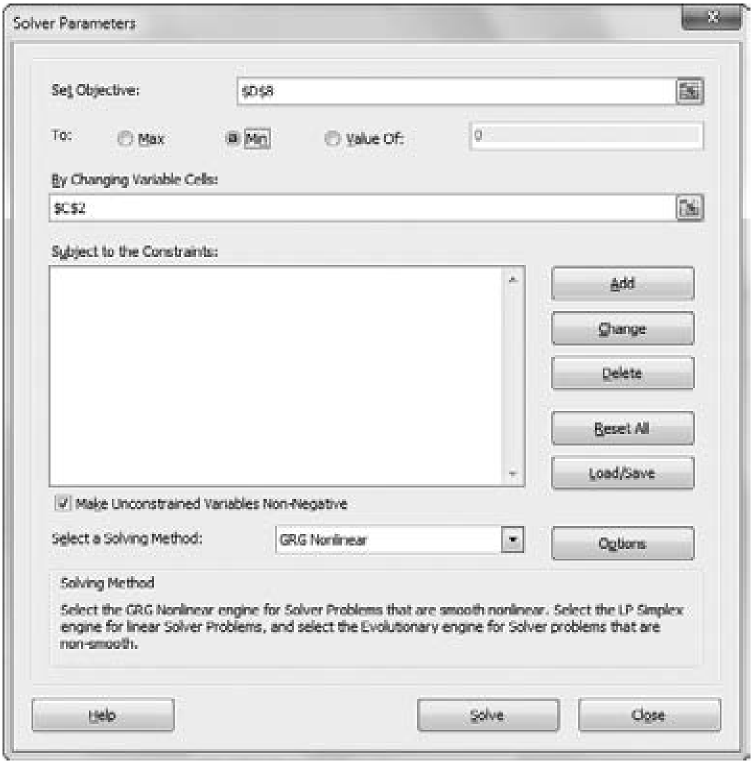

Figure 21-19:
Using the Min option tells Solver to make the cell with the total time as small as possible. The only cell that Solver can change is the one that contains the x value. You don't need constraints. Then click Solve; you'll see the result in Figure 21-20.

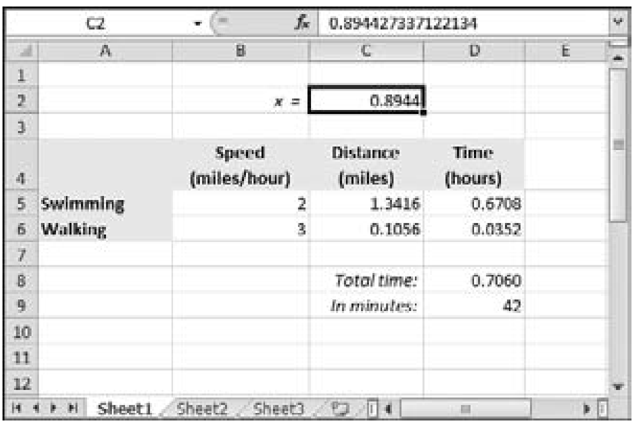

Figure 21-20:
Solver tells you that you need to swim to a point 0.89 miles upstream, which is most, but not all, of the way to the campground. Get out your water wings, because you'll be swimming 1.34 miles, walking a mere 0.1, and getting there in the best possible time of 42 minutes.

As complicated as this example may seem to non-math-geeks, it's actually a picnic for Solver. In fact, Solver's genius—which only math majors may appreciate—is its ability to find an answer in a dense thicket of complicated mathematical relationships. To work this sort of magic, it combines raw stubbornness with some higher-order mathematical reasoning.

POWER USERS' CLINIC

Solver Examples

With a little imagination, you can create Solver models for a variety of scenarios. People often use Solver to find the best ways to distribute work across multiple manufacturing centers (each of which has its own operating costs and maximum production level). Or you could use Solver to plan investments, using constraints to ensure that your assets are properly distributed across different industries or geographic locations. You can learn a lot about how Solver works by looking at some example worksheets that deal with these problems.

Excel provides a workbook called solvsamp.xls that includes several examples on separate worksheets. (You can find this file in a folder like *C:\Program Files\Microsoft Office\Office14\Samples*. If you have trouble locating it, try using the Windows Search feature by clicking the Start button, and then choosing Search.) Each worksheet has sample calculations and detailed information for using Solver, along with color-coded cells that show the objective cell, changing cells, and constraints.

You can find an even more extensive collection of examples at the Solver website (***www.solver.com/solutions.htm***). There, you'll find links for several different categories, each one with several downloadable workbooks that use the same format as solvsamp.xls.

Saving Solver Models

Each time you use Solver, Excel keeps track of the settings you've just used. And each time you save your workbook, it also saves the most recent Solver settings. However, in some cases you may want to keep track of more than one set of Solver settings. This situation could happen if you're trying to optimize the same data in different ways, or if you're using Solver to optimize data in different parts of the same worksheet or different worksheets in the same workbook. In either case, you need to explicitly save the Solver target value, changing cells, and constraints if you want to use them later.

Solver lets you save all of this information and load it later. The only catch is that you need to save it in a small block of cells in one of your worksheets. The actual number of cells you need depends on the Solver model you create. Excel uses one cell to store information about the target value and the objective cell, another cell to store the list of changing cells, and one additional cell for each constraint. The cells are stacked on top of each other.

To save a Solver model, follow these steps:

1. **If you haven't already started Solver, choose Data→Analysis→Solver.**

   The Solver Parameters dialog box appears.

2. **Click the Load/Save button.**

   The Load/Save dialog box appears.

3. **Click inside the text box. Then, click the worksheet where you want to place the first cell.**

   Make sure that there are enough empty cells underneath the cell you select so that the Solver information doesn't overwrite any worksheet data. The Load/Save dialog box tells you exactly how many cells the saved information requires.

4. **Click Save.**

   Solver writes the information from the current model to your worksheet and returns you to the Solver Parameters dialog box.

   Figure 21-21 shows what stored Solver data looks like. The first cell stores a formula that represents the problem you're trying to solve. In this example, it's *=MIN($D$8)*, because you're trying to minimize the value in cell D8. (In the student grade example, the formula is the somehwat nonsensical *=$E$10=0.7*. This formula doesn't do anything legitimate on its own—it's simply Solver's way of remembering that you want to set cell E10 to 70%.) The other cells track the variables in your Solver model and store miscellaneous Solver settings. They use array formulas to pack multiple pieces of information into a single cell.

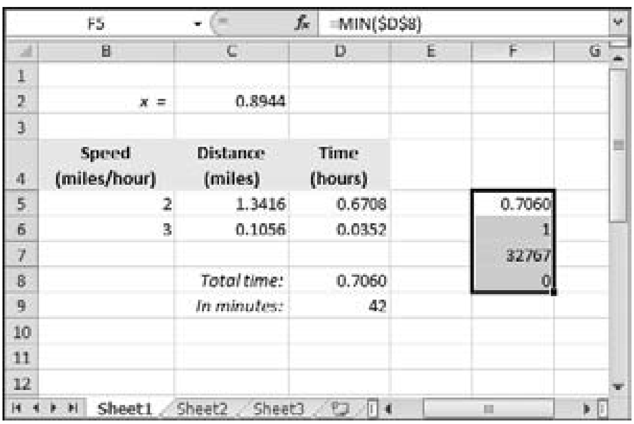

Figure 21-21: *Solver uses the cells F5 to F8 to store the current scenario information.*

To restore a Solver model that you saved earlier, follow these steps:

1. **If you haven't already started Solver, choose Data→Analysis→Solver.**

   The Solver Parameters dialog box appears.

2. **Click the Load/Save button.**

   The Load/Save dialog box appears.

3. **Click inside the text box. Click the worksheet, and then select the cells with the Solver data. Click Load.**

   Solver reads the information from your worksheet and configures the current Solver model accordingly. If you already have a Solver model in your current worksheet, Excel asks you if you want to replace the current Solver settings (click Replace), or combine them with the current settings where possible (click Merge), which is certainly more confusing. Finally, Excel returns you to the Solver Parameters dialog box.

Configuring Solver

The standard options that Solver uses are adequate for most situations. However, there's a wealth of options for more advanced tweaking.

To start, you can choose the solving strategy that Solver uses when it attacks your problem. If you have a good understanding of the mathematics that underpins your problem, this may help you get a better answer or a quicker one. You make your choice in the Select a Solving Method list in the Solver Parameters dialog box:

- **Simplex LP** works for simple, linear problems that have only a few variables linked in a straightforward relationship (like the student grade example). When used on the right problem, Simple LP is the fastest approach.

- **GRG Nonlinear** is the default, and it makes sense if you're not sure what mojo Solver needs to work on your numbers. It's intelligent enough to handle more complex relationships—what mathematicians call *smooth nonlinear functions*—like the campground problem described on page 612. If your objective cell has a formula that uses mathematical operations like log, sine, and cosine, or if takes the variables and multiples them, divides them, or raises them to a power, GRG Nonlinear is probably your best ticket.

- **Evolutionary** is an advanced approach that uses more complex *genetic algorithms*. This approach is generally slower, but it may let Solver arrive at an answer when other approaches fail. You may need this approach if your objective cell uses a formula that incorporates other Excel functions.

Choosing a solving method is just a starting point. If you want to tinker with some of Solver's advanced settings, you can configure them from Solver's Options dialog box (Figure 21-22). To get there, click the Options button from the main Solver Parameters dialog box.

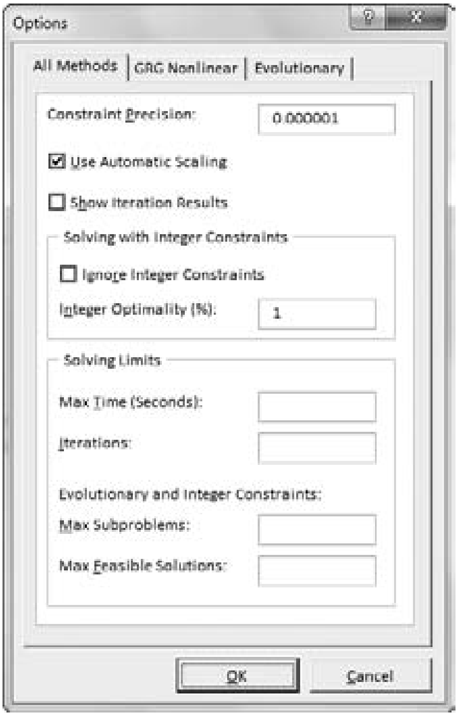

Figure 21-22:
The Options dialog box helps you tweak some options that control how Solver attacks a problem. Most of these enhancements are best left in the hands of experienced mathematicians (and even then, they may not improve Solver's rate of success or its performance). The first tab has settings that apply to all solving methods, while the GRG Nonlinear tab and Evolutionary tab have settings that apply to those specific solving methods only.

The Options dialog box shows a slew of Solver options, all of which you can change. The most useful ones include:

- **Constraint Precision**. This setting indicates how exact a constraint rule needs to be. The smaller this number is (the more zeros after the decimal place), the greater the precision, and the more exacting Solver is in trying to meet your constraints.
- **Show Iteration Results**. If you turn this checkbox on, Solver pauses to show the results of each guess it makes in its trial-and-error process. This choice slows down the process tremendously, but it also gives you some interesting insight into how Solver makes its guesses.
- **Integer Optimality**. This setting comes into play when you use *int* constraints to force one or more variable cells to be integers. For example, if the optimality is set to 1 percent, Solver considers any value that's within 1 percent of a whole number to be an integer.
- **Max Time**. You can use this setting to limit the time Solver takes to a specific maximum number of seconds. You can enter a value as high as 32,767, but Solver rarely reaches the standard limit of 100 seconds in typical small problems.

- **Iterations**. You can use this setting to limit how many trial-and-error calculations Solver makes. Once again, the number can be as high as 32,767. The only reason you'd change the Max Time or Iterations setting is if Solver can't find a valid solution (in which case you may want to increase these limits), or if Solver takes an exceedingly long time without any success (in which case you may want to decrease them).
- **Convergence**. This setting is in both the GRG Nonlinear and Evolutionary tabs, because it applies to nonlinear math problems. When solving this type of problem, Solver may have a hard time telling whether its trial-and-error guesses are approaching the best solution. Each time Solver perform a new attempt, it compares its new guess to the last guess. If five guesses pass and the change is always less than the convergence value, Excel decides that it has settled on the closest solution it can find, and returns an answer. If you make the convergence smaller, Solver takes more time trying to refine its solution, and it may deliver a slightly more precise answer.

CHAPTER

22

Pivot Tables

Creating neat, informative summaries out of huge lists of raw data is a common challenge. And while Excel gives you all the tools you need to create such summaries, the actual work of writing formulas, cutting and pasting information, and organizing your totals into a new table can be extremely tedious. Even worse, this approach isn't very flexible. Once you've created the perfect summary that compares, say, sales in different regions, if you want to compare sales across different product lines or different customers, you'll need to start from scratch and build a whole new report.

Fortunately, Excel has a feature called *pivot tables* that can solve all these problems. Pivot tables quickly summarize long lists of data. By using a pivot table, you can calculate summary information without writing a single formula or copying a single cell. But the most notable feature of pivot tables is that you can arrange them *dynamically*. Say you create a pivot table summary using raw census data. With the drag of a mouse, you can easily rearrange the pivot table so that it summarizes the data based on gender *or* age groupings *or* geographic location. This process of rearranging your table is known as *pivoting* your data: you're turning the same information around to examine it from different angles.

Pivot tables are a hidden gem in Excel. Many otherwise experienced spreadsheet fans avoid them because they seem too complicated at first glance. The real problem is that pivot tables are rarely explained properly. Most books and the online Excel help use no end of cryptic jargon like "cross-tabulated computations" and "*n*-dimensional analysis." But if you stick with this chapter, you'll discover that pivot tables are really just a convenient way to build intelligent, flexible summary tables—nothing more, and nothing less.

Summary Tables Revisited

When you analyze large amounts of data, you can look at the same information in many different ways. How you organize and group the data often determines whether you find or overlook important trends.

Note: In Chapter 15, you saw how to tame large tables of data with features like grouping and outlining. Pivot tables offer another approach—they're a little more work to set up, but they give you much more flexibility in rearranging your tables and calculating different summary information on the fly.

Consider the small table of information shown in Figure 22-1. This table lists all the customers of a small business, along with information about their gender, the city they live in, and their level of education. Looking at this table, an important question comes to mind: is there a relationship between these different pieces of information and the amount of money a customer spends?

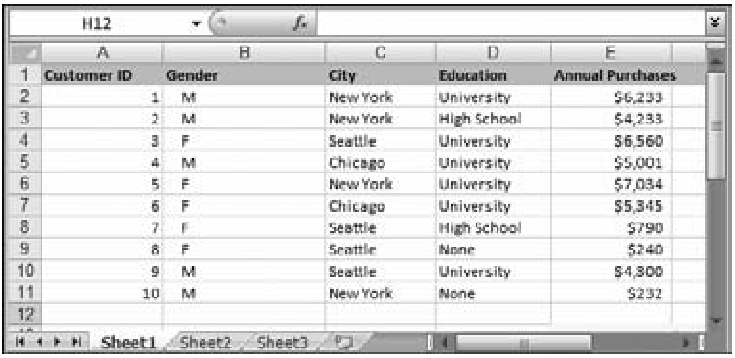

Customer ID	Gender	City	Education	Annual Purchases
1	M	New York	University	$6,233
2	M	New York	High School	$4,233
3	F	Seattle	University	$6,560
4	M	Chicago	University	$5,001
5	F	New York	University	$7,034
6	F	Chicago	University	$5,345
7	F	Seattle	High School	$790
8	F	Seattle	None	$240
9	M	Seattle	University	$4,300
10	M	New York	None	$232

Figure 22-1: *This example has only 10 records, so patterns aren't difficult to spot. However, if you extended this list to hundreds or thousands of rows, then you'd definitely need a summary table to spot any existing relationships.*

To look for trends and patterns in the customer list, it helps to build a *summary table*—a table that tallies key amounts, like the average amount spent for a customer in a specific city, education level, or gender. However, there are several potentially important relationships, and, therefore, several types of summary tables that you could create. Pivot tables are the perfect summary-building tool because they give you almost unlimited flexibility when you want to figure out different relationships. But before you learn about how to build pivot tables, it first helps to understand what life is like in Excel *without* them—because only then can you see why pivot tables make sense and decide whether you need them in one of your own workbooks.

Life Without Pivot Tables

The most basic way to calculate summary information is to use the SUMIF(), COUNTIF(), and AVERAGEIF() functions described in Chapter 13. To find the average annual purchases for an individual in New York using the worksheet shown in Figure 22-1, you use a formula like this:

```
=AVERAGEIF(C2:C11, "New York", E2:E11)
```

This formula scans the City column (C2 to C11) looking for the text. Every time it finds a match, it adds the corresponding purchase amount from the Annual Purchases column (E2 to E11). The final result is the average of all these amounts.

Next, you can make this formula more generic so that you don't need to create a new version for each city. Instead, change the formula so that it retrieves the text it should match (the city's name) from a cell just to the left of where you're going to place the formula. For example, this formula gets the city name from cell A15:

```
=AVERAGEIF($C$2:$C$11, A15, $E$2:$E$11)
```

Note that in this formula, all the search ranges are fixed as absolute references using the dollar sign ($). The search text isn't fixed—that way, when you copy the formula to a new row, the formula uses the search text on that row.

Figure 22-2 shows the result of these "city" formulas, which reside in cells B15, B16, and B17. It also includes some formulas that total up the numbers using different criteria (like the values in the Education and Gender columns).

Figure 22-2: *These three summary tables retrieve their data from the table shown in Figure 22-1 and calculate average purchases by grouping the customers into different categories. The purchase-by-city summary shows that there isn't a significant amount of difference based on location (although Seattle customers tend to spend a little less). The purchase-by-education summary shows a dramatic difference, with less educated individuals making smaller purchases. Finally, the purchase-by-gender comparison turns up no variance at all.*

In this example, building each table isn't terribly difficult, but the situation clearly becomes more tedious the more ways you want to compare the same data. It's not difficult to imagine a more realistic scenario where you may want to look at customer purchases based on age bracket, income level, and the answers to different customer survey questions. To get the full picture with this information, you need to build each table from scratch. In a worksheet that could have thousands of rows of data, this is a tall task.

In fact, in some of these more complex scenarios, you may need to group and then *subgroup* your information. Figure 22-3 shows a more advanced example of a table that calculates the variance in average annual purchases by city, and then shows the subdivided totals in each city by gender.

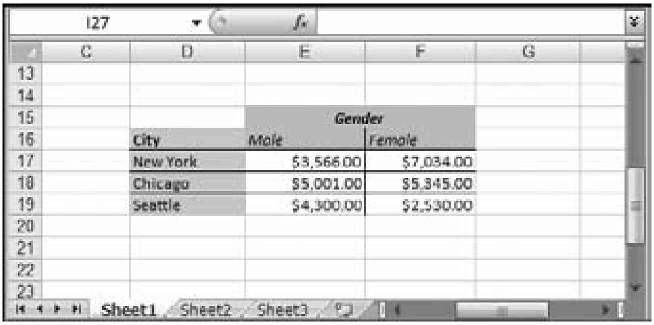

Figure 22-3:
This summary table is a little more advanced than the one shown earlier in the chapter. It groups and then subgroups data, which lets you find hidden trends. In this case it identifies New York females as the best purchasers. Of course, this example has far too few rows for the results to be truly persuasive. In a table with thousands of rows, however, a grouped and subgrouped breakdown becomes much more meaningful.

The summary in Figure 22-3 performs a *two-dimensional* comparison. In other words, it compares two different groupings—one by city in different rows, and one by gender in different columns. This is a step up from the one-dimensional summary tables you saw earlier, but it's also more difficult to correctly calculate. You could use the AVERAGEIFS() functions (page 369), or create the table in two steps (first grouping the records by city, and then totaling them by gender). In either case, life gets much more complicated, and that's when pivot tables really shine.

Life with Pivot Tables

With a pivot table, you can create summary tables like the ones shown so far just by choosing the columns you want to compare. And once you've built your summary, you don't need to stick with it—instead, you can transform a purchase-by-education table into a purchase-by-city table just by dragging and dropping. You can even drill down into the details, apply filter criteria, and apply advanced subtotaling calculations. The following sections describe all of these techniques.

Building Pivot Tables

Now that you've learned the role that pivot tables play in summarizing data, it's time to create your own. Before you begin, you need to have a long list of raw data that you want to summarize. You could use the customer list from the previous example, but it's too small to really show the benefits of pivot tables. A better example is something like the list of order information shown in Figure 22-4.

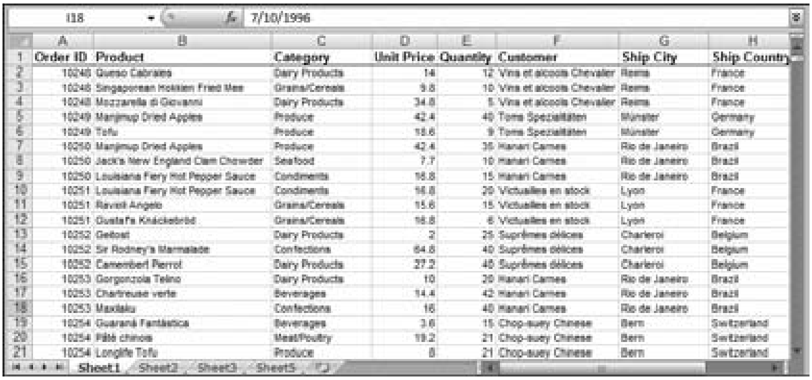

	Order ID	Product	Category	Unit Price	Quantity	Customer	Ship City	Ship Country
2	10248	Queso Cabrales	Dairy Products	14	12	Vins et alcools Chevalier	Reims	France
3	10248	Singaporean Hokkien Fried Mee	Grains/Cereals	9.8	10	Vins et alcools Chevalier	Reims	France
4	10248	Mozzarella di Giovanni	Dairy Products	34.8	5	Vins et alcools Chevalier	Reims	France
5	10249	Manjimup Dried Apples	Produce	42.4	40	Toms Spezialitäten	Münster	Germany
6	10249	Tofu	Produce	18.6	9	Toms Spezialitäten	Münster	Germany
7	10250	Manjimup Dried Apples	Produce	42.4	35	Hanari Carnes	Rio de Janeiro	Brazil
8	10250	Jack's New England Clam Chowder	Seafood	7.7	10	Hanari Carnes	Rio de Janeiro	Brazil
9	10250	Louisiana Fiery Hot Pepper Sauce	Condiments	16.8	15	Hanari Carnes	Rio de Janeiro	Brazil
10	10251	Louisiana Fiery Hot Pepper Sauce	Condiments	16.8	20	Victuailles en stock	Lyon	France
11	10251	Ravioli Angelo	Grains/Cereals	15.6	15	Victuailles en stock	Lyon	France
12	10251	Gustaf's Knäckebröd	Grains/Cereals	16.8	6	Victuailles en stock	Lyon	France
13	10252	Geitost	Dairy Products	2	25	Suprêmes délices	Charleroi	Belgium
14	10252	Sir Rodney's Marmalade	Confections	64.8	40	Suprêmes délices	Charleroi	Belgium
15	10252	Camembert Pierrot	Dairy Products	27.2	40	Suprêmes délices	Charleroi	Belgium
16	10253	Gorgonzola Telino	Dairy Products	10	20	Hanari Carnes	Rio de Janeiro	Brazil
17	10253	Chartreuse verte	Beverages	14.4	42	Hanari Carnes	Rio de Janeiro	Brazil
18	10253	Maxilaku	Confections	16	40	Hanari Carnes	Rio de Janeiro	Brazil
19	10254	Guaraná Fantástica	Beverages	3.6	15	Chop-suey Chinese	Bern	Switzerland
20	10254	Pâté chinois	Meat/Poultry	19.2	21	Chop-suey Chinese	Bern	Switzerland
21	10254	Longlife Tofu	Produce	8	21	Chop-suey Chinese	Bern	Switzerland

Figure 22-4: *This worksheet shows some entries from an order list of 2,155 grocery store items. Lists like this make great pivot table candidates.*

UP TO SPEED

Pivot Tables vs. Grouping and Subtotals

When should you use pivot tables instead of other Excel features? In Chapter 15, you learned about two techniques for grouping and organizing data:

- **Grouping**. Excel's grouping tools let you collapse detailed information in a table so that only the summary information remains. This technique is the perfect solution when you have a worksheet that contains both detail and summary information and you want to make it more manageable. However, you still need to write all the summary formulas and decide how you want to organize the information.
- **Automatic subtotaling**. Automatic subtotaling generates subtotals for grouping data according to a column that you choose. Automatic subtotaling is a good technique for building a quick-and-dirty summary report, but it doesn't provide a lot of flexibility. It also breaks up the structure of your list so that you can't effectively edit it any longer.

Compared to either of these features, pivot tables are more flexible and more dynamic. A single pivot table can group and subtotal the same list in different ways—all you need to do is drag and drop your columns. Thanks to this flexibility, pivot tables are a much more flexible tool for analyzing your information and discovering hidden relationships.

Not all data is suited for a pivot table. To work well, your data needs to meet a few criteria:

- **It has to include at least one column that has duplicate values**. In the order table shown in Figure 22-4, there are multiple records with the same value in the Customer column. Accordingly, you can create a separate group of ordered items for each customer.
- **It has to include some numeric information**. You'll use this information to create subtotals. Often, you'll be interested in generating a simple count, total, or average, although you can also find maximum, minimum, and standard deviation, and you can use your own formulas.

Note: It's technically possible to create subtotals without using a numeric column. In this case, the subtotals just count the number of values in the group. This approach is occasionally useful, but it's not as powerful as other types of subtotals.

The order information table is perfect for a pivot table, because it has several columns you can use to group the order rows. These include:

- **Product and Category**. Find out how well specific products are selling or what the hottest product categories are.
- **Customer**. Find out who's making the most purchases.
- **Ship City and Ship Country**. Find out where the majority of your customers reside.

A pivot table can handle all of these comparisons. You don't need to choose one column or another before you start building the pivot table.

Tip: The best way to learn about pivot tables is to perform the steps in this chapter, and then start experimenting. If you don't happen to have a table with hundreds of records on hand, you can download the workbook shown in Figure 22-4 from the Missing CD page at *www.missingmanuals.com/cds*. It gives you 2,155 rows to summarize and pivot to your heart's content.

Preparing a Pivot Table

Creating a new pivot table is a two-step process. First you need to run the PivotTable and PivotChart wizard, which asks you to identify the data you want to summarize and select the location where you want to place the pivot table. The next step is to actually define the structure of the pivot table and try out different ways of organizing and grouping your data.

The following steps lead you through the first step in creating a new pivot table:

1. **Select the range of cells that you want to use for your pivot table.**

   If you plan to add more rows later, consider using a table (Chapter 14) instead of selecting a range of cells. That way, when you add rows to the data list and refresh the pivot table, Excel automatically adds your new rows to the pivot table.

2. **Select Insert→Tables→PivotTable→PivotTable.**

   If you're creating a pivot table for a table that you've defined with the Insert→Tables→Table feature (page 396), there's a shortcut. Just move anywhere inside your table, and then choose Table Tools | Design→Tools→Summarize with Pivot.

   Either way, the Create PivotTable dialog box appears (Figure 22-5). Excel automatically chooses the "Select a table or range," with the table name or cell range you selected.

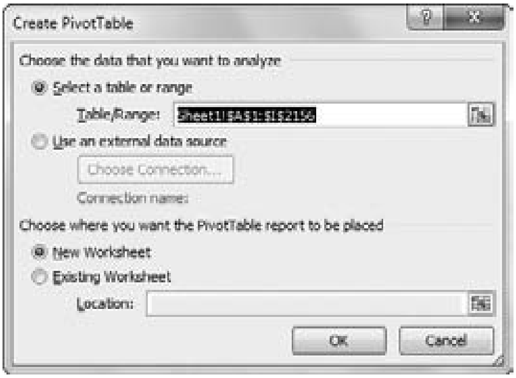

Figure 22-5:
The Create PivotTable dialog box asks where your raw data is (in a table, cell range, or external data source) and where you want to put the pivot table that summarizes it.

Note: You can also build a pivot table based on records that you select from an external database. But before you take this step, you must configure your database as an Excel data source. To learn how to do so, see page 658.

3. **Select "New worksheet" to create a new worksheet for your pivot table, which is usually the easiest option.**

   Pivot tables are fairly complex creations, so the easiest way to manage them (and keep them separate from the rest of your data) is to pop them into a new worksheet. Alternatively, you can choose "Existing worksheet" to insert your pivot table on a worksheet that's already in your workbook. In this case, you need to specify the cell reference for the top-left corner of the pivot table. If there's data under this cell or to the right of it, then Excel may overwrite it as it generates the pivot table. Usually, the best approach is to place a new pivot table on a separate worksheet.

4. **Click OK.**

   Excel inserts the new pivot table. Because you haven't defined the columns it should use for grouping and analyzing rows, this pivot table appears as an empty placeholder, as shown in Figure 22-6.

Note: If you choose to create a new worksheet, Excel gives the worksheet an unhelpful name (like Sheet5), and then places it before the worksheet that contains your data. You can rename the new worksheet, and then drag it into a better spot using the techniques described in Chapter 4.

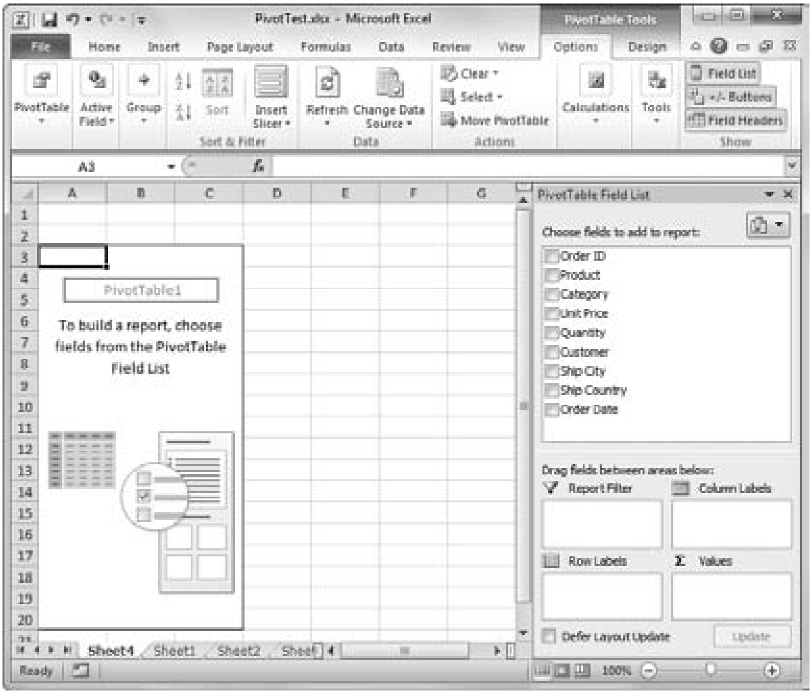

Figure 22-6:
This worksheet shows a newly created pivot table that doesn't yet contain any information. When you move to a cell inside the pivot table, Excel displays the PivotTable Field List pane on the right, which contains a list of all the columns in the data that you're summarizing.

Note: When you create an empty pivot table in Excel 2003 compatibility mode (page 43), it looks a little different. Instead of the graphical box shown in Figure 22-6, you get a series of linked white boxes, with messages inviting you to drop fields in each one. Microsoft designed this strange appearance to mimic the way that blank pivot tables appeared in Excel 2003. But don't worry—the instructions in this chapter will still work just as well.

Pivot Table Regions

To build a pivot table, simply drag columns from the PivotTable Field List pane on the right side of the Excel window, and then drop them into one of the four boxes underneath. As you work, Excel generates the pivot table, updating it dynamically as you add, rearrange, or remove columns.

Note: Excel refers to all your source data's columns as *fields*.

To understand how to fill up a pivot table with data, you need to know how each region works. Altogether, a pivot table includes four regions:

- **Values**. These are the fields that you want to subtotal and display in your pivot table. Usually, this is the numeric information you use to build averages and totals.
- **Row Labels**. These fields let you break down the data into categories. The pivot table uses the Row Labels fields to create separate groups, and puts the information for each group into a single row. Using the product list in Figure 22-4, you could group products into categories by dropping the Category field here.
- **Column Labels**. Often, you want to subdivide your data in more than one way at once. Like the Row Labels fields, the Column Labels fields are used to create groups. The difference is that these groups are displayed in separate columns. If you use Category for a row field and Ship Country for a column field, you wind up with a table that divides sales figures into rows by product category, and then divides each category row into columns by country.
- **Report Filter**. These fields use filtering to limit which data the pivot table displays. If you want to look at a breakdown of sales by category, but you want to consider only U.S. sales, you could add Ship Country as a page field and configure it accordingly.

It really doesn't matter whether you use a field for row grouping or column grouping. The pivot table shows the same data either way, but one approach may be more readable than another. If you have a field with extremely long names, it probably works better as a row field than as a column field (where it would stretch out the width of the whole column).

Also, consider how many different groups you want to create. If you want your pivot table to compare sales by category and country, and your list features five categories and 20 countries, you'll probably be best off if you use the country field as a row field and the category field as a column field. That keeps the table long and narrow, which is easier to read and print.

Laying Out a Pivot Table

To get a better understanding of how to create a pivot table, it helps to follow along with a basic example. These steps lead you through the process of creating a summary that compares the products and shipping locations shown in Figure 22-4.

1. **In the PivotTable Field List pane, drag the Product field into the Row Labels box underneath.**

   When you drop the field, Excel fills in the names of all the products from the list from top to bottom, in alphabetical order (see Figure 22-7). When you finish this step, Excel creates one row containing the subtotals for each product.

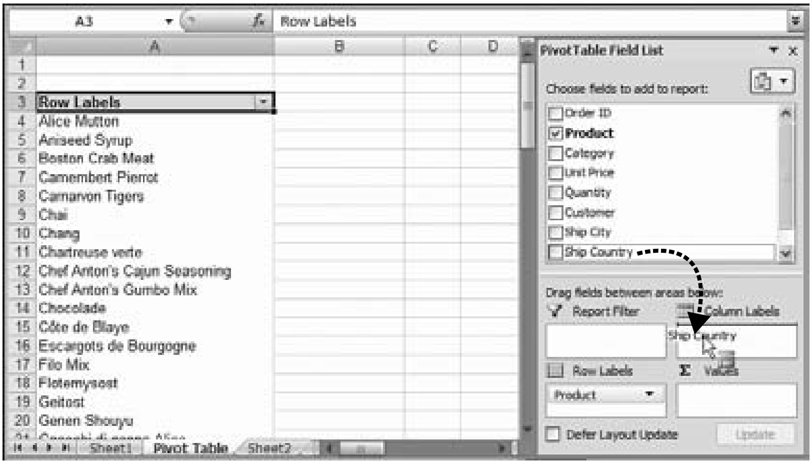

Figure 22-7: *In this example, the list of products has already been added to the row area (and it appears in column A), and the second grouping criteria (the list of countries) is being dragged into the Column Labels box. Notice that once a field is linked, its name is listed in boldface in the PivotTable Field List.*

2. **From the PivotTable Field List, drag the Ship Country field to the Column Labels box.**

   When you drop the field, Excel fills in the names of all the countries from the list from left to right, in alphabetical order. In other words, each country is listed in its own column.

3. **Now you need to choose what data you want to examine. Drag the Quantity field to the Values box.**

   This step is designed to actually fill the table with data—specifically, the numbers of products that were ordered by customers in various countries. It's helpful to remember what's in the table just prior to this step: a list of products in column A and, from column B on, each column is labeled with a different country.

   Once you complete this step, Excel generates the pivot table using its standard calculation option, which is to perform sums on all the fields in the Values box. In this example, it adds all the values in the Quantity field for a given group, which tells you the total number of units of a specific product shipped to a specific country. Figure 22-8 shows the result.

Note: Pivot tables also calculate row and column subtotals. If you want to find the total number of units shipped for a given product across all countries, scroll to the far right end of the chart. If you want to find the total number of units sold in a given country, scroll to the totals at the bottom of the chart.

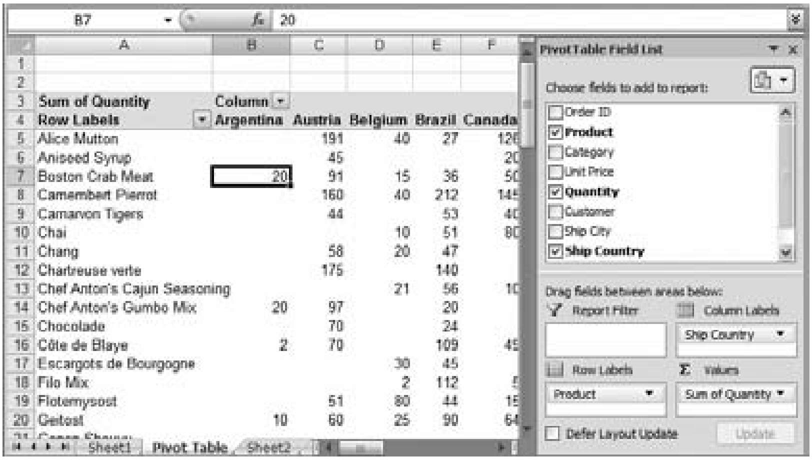

Figure 22-8: *This pivot table totals the quantity of product units in every order. For example, it tells you that 20 units of Boston Crab Meat were shipped to Argentina.*

This example built a fairly sophisticated two-dimensional pivot table, which means that it compares two different groupings (one represented with rows, and the other represented with columns). Most of the pivot tables you'll see in real life are two-dimensional, but, there's no reason you can't create simpler one-dimensional pivot tables. All you need to do is leave out either the Column Labels or Row Labels. Figure 22-9 shows a pivot table that simply totals the number of units sold for particular products.

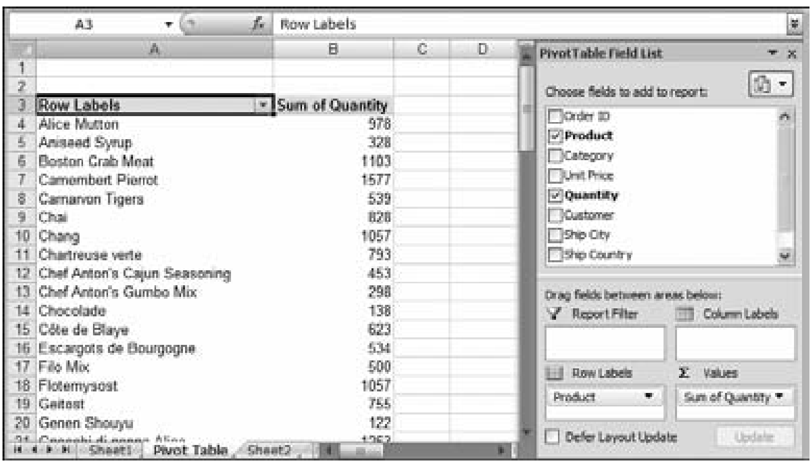

Figure 22-9: *This pivot table is one-dimensional. It groups orders by product, but it doesn't subgroup them by country. If you want to create a pivot table like this, don't drop anything in the Column Labels box.*

Unlike most other elements in Excel, pivot tables don't refresh themselves automatically. That means that if you change the source data, the pivot table may show out-of-date totals. To correct this problem, you can refresh the pivot table by moving to one of the cells in the pivot table, and then selecting PivotTable Tools | Options→Data→Refresh (or the keyboard shortcut Alt+F5). This action tells Excel to scan the source data and regenerate the pivot table.

Formatting a Pivot Table

As you've probably already noticed, when you move into one of the cells of a pivot table, two new tabs spring up in the ribbon under the PivotTable Tools heading. These are similar to the tabs that appear when you select a chart, picture, or table.

You can use the PivotTable Tools | Options tab to access a few advanced features that you've yet to consider, like grouping, pivot table formulas, and pivot charts. You'll consider these features in the rest of this chapter. The PivotTable Tools | Design tab is more modest; you use it to format and otherwise fine-tune the appearance of your pivot table.

The PivotTable Tools | Design tab is carved into three sections. At the far right, you'll find the PivotTables Styles section, with a familiar style gallery. If you choose one of the entries here, Excel adjusts your PivotTable automatically, giving it new colors and shading that can range from subtle to dramatic.

Tip: The colors that Excel uses in the pivot table styles actually come from your workbook theme (page 174). So if you want to get a completely different set of accent colors, just choose a new theme from the Page Layout→Themes→Themes gallery.

Some pivot table layouts use *banding*—a pattern of alternating shading that separates one row (or one column) from the next. Banding works well in some pivot tables, but you may find that the effect is overwhelming in others. To give you more control over your pivot table's formatting, Excel lets you apply just a portion of the full pivot table style. You can do so by using the checkboxes in the PivotTable Tools | Design→PivotTable Style Options section of the ribbon. If you want to use a specific style but you don't want to apply its banding effects, clear the Banded Rows and Banded Columns checkboxes. Similarly, if you don't want to apply the style formatting to headers, clear the Column Headers or Row Headers checkboxes.

Finally, the PivotTable Tools | Design→Layout section lets you choose from various preset options that control spacing and subtotals. This section contains four submenus:

- **Grand Totals**. Pick an option from this list to show or hide the totals at the end of each row and column.
- **Subtotals**. Pick an option from this list to show or hide subtotals at the end of every group. This setting has an effect only if your groups are subdivided—in other words, if you have more than one field in the Row Labels section or more

than one field in the Column Labels section. Otherwise, the group "subtotals" are actually the grand totals.

- **Report Layout**. Pick an option from this list to choose how tightly packed your pivot table is. In Excel 2010 and Excel 2007, pivot tables are shown in "compact" form, which keeps columns as narrow as possible. (Typically, the column is just wide enough to fit the data inside and the column header.) Older versions of Excel made each column as wide as the widest column, which takes considerably more room. If you still want this space-chewing display choice, then choose PivotTable Tools | Design→Layout→Report Layout→Show in Outline Form.
- **Blank Rows**. Pick an option from this list to add a blank line between groups. This option has an effect only if there's more than one field in the Row Labels section.

Rearranging a Pivot Table

So far, you've seen how to use a pivot table to quickly build a summary table. However, pivot tables have another key benefit: flexibility. There's no limit to how many times you can move fields or recalculate your summary so that it performs different calculations.

To change a pivot table, you can use the following techniques in the PivotTable Field List pane:

- **To remove** a field from a pivot table, click the field header (in the appropriate box), and then drag it out of the PivotTable Field List pane. The mouse pointer changes to an X symbol to indicate you're removing the field.
- **To move** a field from one position to another, just drag the field header from one box to another. You can reverse the example shown earlier by dragging the column field (Country) to the Row Labels region, and dragging the row field (Product) to the Column Labels region.

Figure 22-10 shows one way you could rearrange the pivot table shown in Figure 22-8.

Getting to the Source

As you review the summarized information in a pivot table, you just might discover trends and relationships you didn't expect. For example, the pivot table in Figure 22-10 exposes a fact you probably wouldn't have realized by just looking at the raw data—namely, Austrians love their dairy products.

When faced with a fact like this, it's often worthwhile to dig a little deeper. For example, you might wonder exactly how the high value in cell C8 was calculated. Did customers in Austria place a huge number of orders, or were there just a few orders that were unexpectedly large? Now that the pivot table has done its number crunching job by highlighting this sort of information, you need to return to the source data to further your investigation.

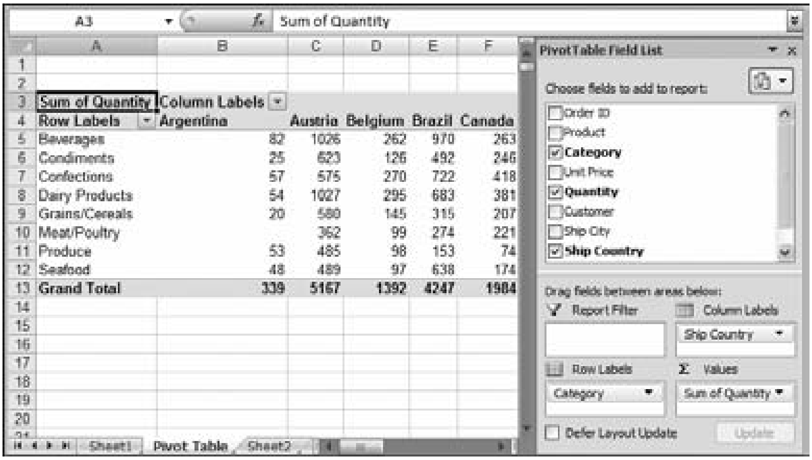

Figure 22-10: *Another way to organize the data shown in Figure 22-8 is to see how the product categories perform in each of the countries, as shown here. This summary table makes it easy to spot categories that do unusually well in certain countries.*

Although you could flip back to the worksheet that has your source data and search for the corresponding records, Excel has an easier solution. If you click any value in a pivot table, Excel adds a new worksheet. In this worksheet, it copies the records that were used to calculate the pivot table cell in question, and nothing else. For example, if you double-click cell C8 in Figure 22-10, Excel adds a worksheet that shows all the Austrian dairy sales (Figure 22-11).

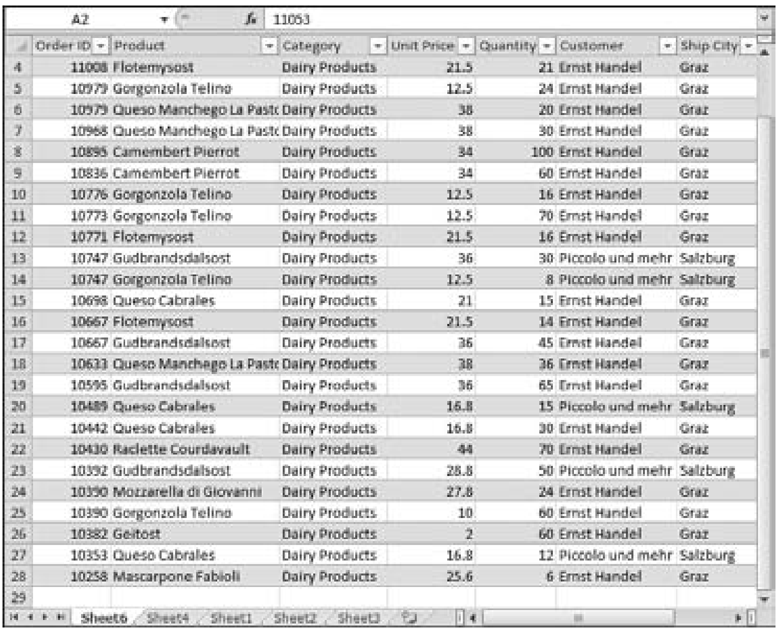

Figure 22-11: *When you double-click the calculated total for Austrian dairy sales, Excel creates this worksheet, which shows the 27 records that were added together to produce the pivot table result.*

Once you've finished reviewing the data, you can delete the worksheet Excel created. After all, the full information is in your original worksheet. To remove the worksheet, right-click the corresponding worksheet tab at the bottom of the window, and then choose Delete.

Multi-Layered Pivot Tables

So far, you've seen examples of how to create one- and two-dimensional pivot tables. But the fun doesn't need to end there. In fact, there's no limit to the number of groupings you can add to a pivot table. To add additional levels of grouping, simply drag the appropriate fields from the PivotTable Field List onto the row or column area of the pivot table. Each time you add a new grouping, Excel *subdivides* your current groups.

If you add Product to the row area of your pivot table, Excel groups all your records into rows so that each row totals the information for a separate product. Next, say you add *another* field to the row area—this time, the Order Date field. Excel responds by dividing each product row into *multiple* rows. Each individual row shows the total units for a given product sold on a given day.

The problem with subdividing is that it can needlessly enlarge the size of your summary table. If you're not careful, your summary table may not be a summary at all! Consider the table in Figure 22-12, which shows the Category and Order Date fields in the row area, and the Ship Country field in the column area. The problem here is that there aren't many orders that fall on the same date. Even when they do, they're often for products in different categories. As a result, many of the rows aren't true totals—instead, they display only the results for a single order.

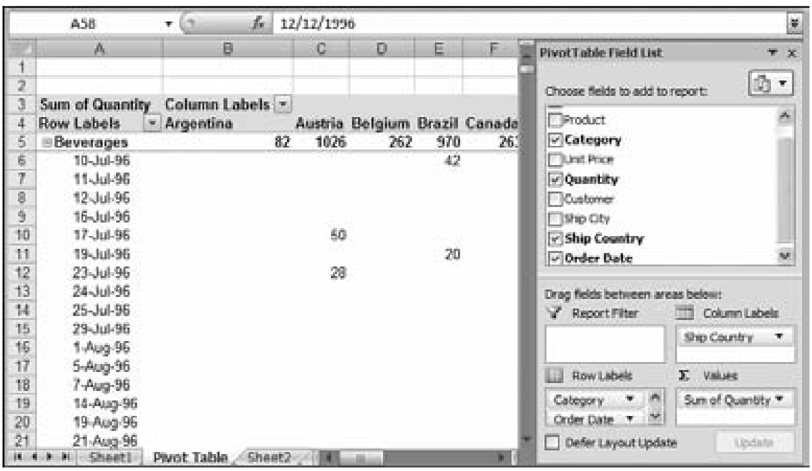

Figure 22-12: *In this pivot table, rows are grouped by category and subdivided by order date. That means each row shows information for the orders made for a given product category on a given day. Each row is further broken up into columns by country. At the bottom of each group (not shown), Excel subtotals all the rows for that product across all days.*

One hint that the table isn't performing a good summary is the number of rows in the pivot table. At 1,621 rows, it's not much smaller than the total 2,155 rows of source data in the original table. You'll also notice that the table is quite sparse, because each row is further broken down into columns (by country). The end result is a lot of blank, wasted space.

This example may lead you to conclude that pivot tables with more than two groupings aren't much use—unless you have a staggering amount of data (thousands or even tens of thousands of records). However, subgrouping does come in handy if you have *related* fields. Figure 22-12 shows two good examples of related fields in the orders table:

- **Categories and Products**. Each product is a part of one category. Thus, you can group first by category, and then subdivide each category by product.
- **Ship Country and Ship City**. Each city is located in one country. You can group first by country, and then subdivide each country into cities.

In both of these examples, it's important to make sure you add the fields in the correct order. You don't want to group by product and then group by category, because each product is only a part of one category. Instead, you want to group by category, and *then* list products within these categories. Figure 22-13 shows the difference.

Row Labels	Sum of Quantity
⊟Alice Mutton	**978**
Meat/Poultry	978
⊟Aniseed Syrup	**328**
Condiments	328
⊟Boston Crab Meat	**1103**
Seafood	1103
⊟Camembert Pierrot	**1577**
Dairy Products	1577
⊟Carnarvon Tigers	**539**
Beverages	539
⊟Chai	**828**
Meat/Poultry	828
⊟Chang	**1057**
Beverages	1057
⊟Chartreuse verte	**793**

Row Labels	Sum of Quantity
⊟Beverages	**9243**
Carnarvon Tigers	539
Chang	1057
Chartreuse verte	793
Côte de Blaye	623
Guaraná Fantástica	1125
Ipoh Coffee	580
Lakkalikööri	981
Laughing Lumberjack Lager	184
Outback Lager	817
Rhönbräu Klosterbier	1155
Sasquatch Ale	506
Steeleye Stout	883
⊟Condiments	**4547**

Figure 22-13:
Top: When subgrouping, make sure you place your fields in the correct order. Here, the pivot table groups the records by product, and then subdivides the products by category, which really doesn't make sense. The result is a table where each group contains a single subgroup.

Bottom: This pivot table's rows are grouped by category, and then subdivided by product, which makes much more sense.

To make sure you have the right grouping, look in the Row Labels box (if you're grouping rows) or the Column Labels box (if you're grouping columns) in the PivotTable Field List pane. In the example shown in Figure 22-13, bottom, the Category field should appear above the Product field in the list. If it isn't, you can drag the Category field into place, or you can click it, and then choose Move Up from the pop-up menu.

Hiding and Showing Details

Subgrouping gives you another interesting ability—you can hide or show individual groups. This feature lets you show detailed information for just the part of the table that you're interested in, while hiding the rest. In fact, this feature works just like the collapsible outlines you learned about in Chapter 15.

Imagine you create a pivot table that uses the Category and Product fields to group rows. When you create this table, Excel shows you every product in every category. But what if you want to show only the products in a specific category? In this case, the trick is to hide every category you *don't* want to see.

- **To hide** (or *collapse*) the products in a specific category, click the plus/minus icon next to the category name. You can repeat this process to collapse as many categories as you want. All you'll see is the row with the category totals.
- **To show** a collapsed category, click the plus/minus icon again.

GEM IN THE ROUGH

Avoiding Slow Refreshes

If your pivot table involves a huge amount of information, and if you plan to use multiple levels of grouping, you may not like Excel's standard behavior. Every time you drag a field to a box in the PivotTable Field List pane, Excel regenerates the pivot table with your new settings. This refresh takes a bit of time, and it can make designing a complex pivot table seem just a little sluggish.

If you're tired of waiting, Excel has a solution. Before you make any changes to your pivot table, switch on the Defer Layout Update setting, which appears at the bottom of the PivotTable Field List pane. While the Defer Layout Update is on, Excel doesn't refresh your pivot table as you build or change it. If you're a fast mover and you're dealing with a massively large pivot table, this lets you work much faster.

When you're finished making your changes (or you just want to see what the pivot table looks like so far), click the Update button to perform a one-time refresh. Or, turn off the Defer Layout Setting to go back to normal, and let Excel refresh the pivot table automatically.

Of course, there's one disadvantage to using Defer Layout Update. You don't see the effects of your changes until you're finished. As a result, the Defer Layout Update setting is a great tool for pivot table pros, but you'll probably want to stay away from it while you're learning.

Tip: For an even quicker shortcut to hide or show a category, just double-click the cell with the category name (like *Beverages* in A6). When you double-click an expanded category, Excel collapses it. When you double-click a collapsed category, Excel expands it. This feature helps you quickly drill down to the most interesting parts of your summary.

- **To collapse** all the categories in your pivot table, move to any category, and then choose PivotTable Tools | Options→Active Field→Collapse Entire Field.
- **To expand** all your categories, move to any category, and then choose PivotTable Tools | Options→Active Field→Expand Entire Field.

Figure 22-14 shows a pivot table that takes full advantage of Excel's ability to hide and show details.

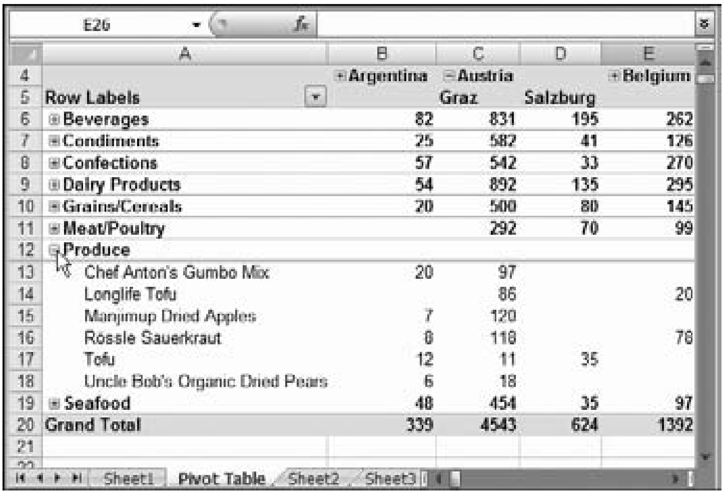

Figure 22-14: *In this pivot table, rows are grouped by Category and subgrouped by Product. Columns are grouped by Country and subgrouped by City. All the category groups are collapsed except for Produce, and all the country groups are collapsed except for Austria. This way, the pivot chart highlights produce sales in Graz and Salzburg, two picturesque Austrian cities.*

There's no limit to how many levels of grouping you can add. If you use the Show Detail command to try to expand the last level of your pivot table, Excel prompts you with a Show Detail dialog box that lists all the fields you aren't using currently. If you choose one of these fields, and then click OK, Excel adds another layer of grouping to the pivot table, as shown in Figure 22-15.

Note: There's a screencast—an animated online tutorial—that shows how to dig through the levels of a multi-layered pivot table, online at *www.missingmanuals.com/cds*. It was created using Excel 2007, but the process is the same. Surf to the Missing CD page for *Excel 2007: The Missing Manual* to check it out.

Fine-Tuning Pivot Table Calculations

As you saw earlier, when you add a field to the Values box, Excel guesses what calculation you want to perform. In most cases, it assumes you want to perform a sum operation that totals up all the values in the field. However, this calculation isn't always the right one. Consider the sales summary pivot table that you've been exploring in this chapter. Although it makes sense to examine the total units sold, you may be just as interested in the maximum, minimum, or average order size, or the order *count* (the number of times the product was ordered, without considering the number of units in each order).

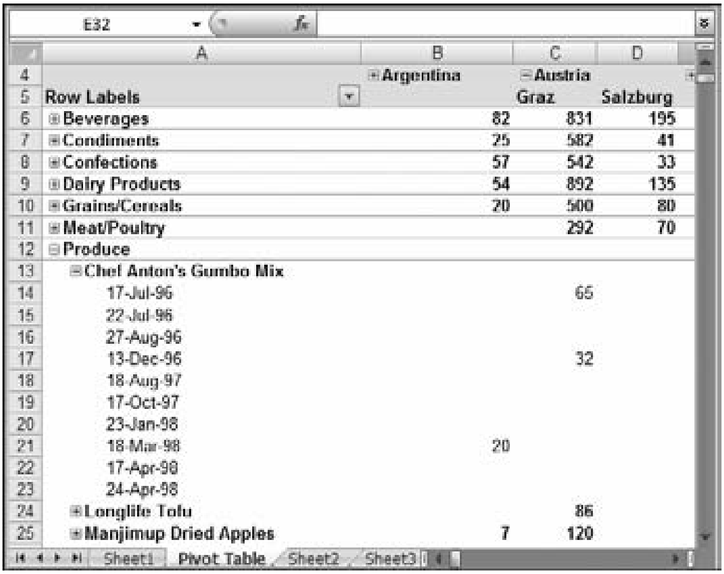

***Figure 22-15:** This pivot table drills down through three levels of row groupings. It shows a detailed breakdown that indicates when Chef Anton's Gumbo Mix was ordered, and exactly where the shipments were headed.*

Fortunately, Excel makes it easy to change the type of calculation you're using. And as you'll see in the following sections, you can even perform more than one calculation in the same pivot table, and throw your custom formulas into the mix.

Changing the Type of Calculation

To modify the calculation that the pivot table performs, follow these steps:

1. **Find the appropriate field in the Values box of the PivotTable Field List pane. Click the drop-down arrow, and then choose Value Field Settings.**

   If you want to change the current operation, which sums together the Quantity value for each row in a group, then click the "Sum of Quantity" item in the Values box.

   When you do, the Value Field Settings dialog box appears (Figure 22-16).

2. **In the "Summarize by" tab, choose a different option from the list.**

   Choose Count to add up the number of different orders for a product, or choose Average to calculate the average order size.

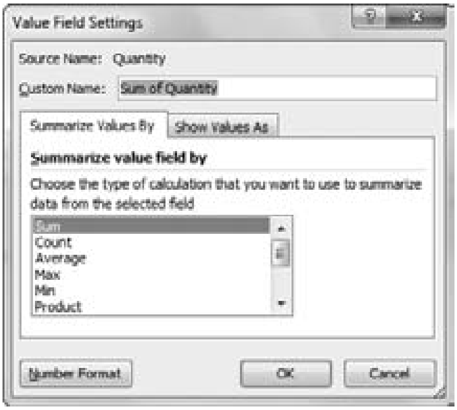

Figure 22-16: *Using the "Summarize by" tab, you can choose the subtotaling calculation you want to perform. You can use counts, averages, sums, or find the maximum value, minimum value, or standard deviation. The "Show Values as" tab lets you configure more complex calculations that compare fields (like differences, percentages, and so on) or calculate running totals.*

3. **If you want to define a different number format to display the summary information, click the Number Format button, choose a new format, and then click OK.**

   When you click Number Format, Excel shows a slimmed-down version of the Format Cells dialog box, which only includes the Number tab. You can use this tab to change the number of decimal places, get a currency symbol, and so on.

4. **Click OK to close the Value Field Settings dialog box.**

   Excel refreshes the pivot table with the new information.

Tip: These column headings that are used for data values can use up some valuable space. To get narrower columns, you can apply a shorter custom column name. To do so, click the field in the Values box, and then choose Value Field Settings. In the Value Field Settings dialog box, enter a new field name in the Custom Name box.

Adding a Calculated Field

Using the pivot tables you've learned about so far, you can perform a variety of operations on any field, including averages, sums, and other preset calculations. But what if you want to branch out a little further and create a *custom calculation?* You could alter the source data, add a new column containing your calculation, and then recreate the pivot table. But Excel provides a much easier option by letting you define a calculated field.

POWER USERS' CLINIC

Adding Multiple Values

As you've seen, you can add more than one field to the Row Labels and Column Labels boxes. You can use the same technique to add multiple fields to the Values box. When you do, each field is calculated and shown on the pivot table, in a separate column (as shown in Figure 22-17).

You could decide to show the total for the Quantity field *and* the average of the Unit Price field. To do so, drag both fields into the Values box. Then, follow the steps on this page to configure the type of calculation that's performed for each field.

You can also use this technique to perform multiple calculations with the same field. If you want to average and total the Quantity field, drag the Quantity field into the Values box twice. You'll end up with two separate items, which you can configure separately.

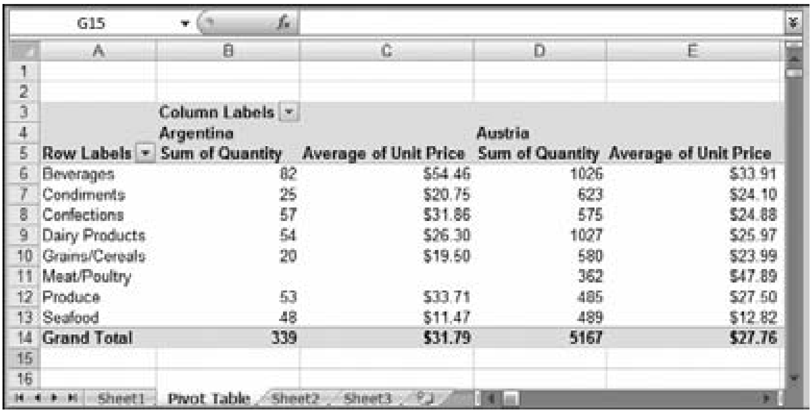

Figure 22-17: *When your pivot table contains multiple values, they appear in different columns with headings like "Sum of Quantity" and "Average of Unit Price."*

The following steps show how to create a calculated field within a pivot table. Using these steps, you can modify a pivot table that shows a list of orders so that it totals the amount of revenue generated instead of just adding up the number of units shipped.

1. **Move anywhere inside the pivot table.**

2. **Choose PivotTable Tools | Options→Calculations→Fields, Items, & Sets→ Calculated Field.**

   The Insert Calculated Field dialog box appears (as shown in Figure 22-18).

3. **Enter a name for the new field in the Name text box.**

   In this case, the new name is Revenue.

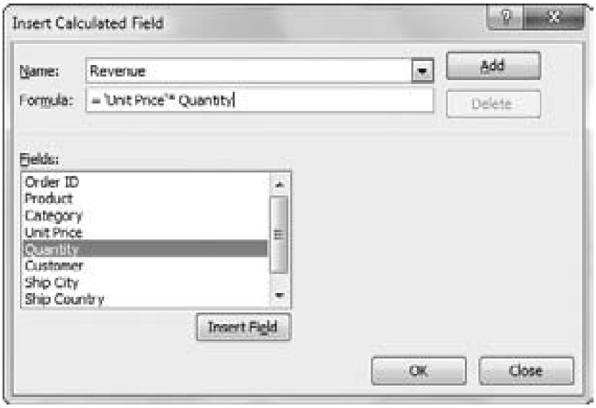

Figure 22-18:
Here, Excel creates a new field named Revenue to calculate how much money is being made on any given order. To arrive at the total revenue, the Unit Price field is multiplied by the Quantity field.

4. **Enter the formula this field uses in the Formula text box.**

   (This is a standard Excel formula that modifies one or more of the other fields provided in the Fields list.)

   Your formula can use Excel's built-in functions, or you can enter a formula that alters or combines one or more of the fields that are in the Fields list. The new Revenue field calculates the total revenue generated by multiplying the price of the item by the quantity purchased. It uses the following formula:

   ```
   ='Unit Price'* Quantity.
   ```

   Note that field names that include spaces or special characters, like Unit Price, need to be enclosed in apostrophes. If you're in any doubt, double-click the field name in the list to have Excel insert it into the Formula text box with apostrophes if they're necessary.

5. **Click OK.**

   Excel automatically adds the calculated field to the PivotTable Field List pane, and then inserts it in the Values box, so that it appears in the pivot table (Figure 22-19).

   As with any other numeric field, Excel assumes you want to perform a sum calculation that totals up your formula for every row. If you want to perform a different calculation, you can customize your field as described in the previous section.

Note: Even if you drag your custom field out of the Values box to remove it from the pivot table, it remains in the field list. If you want, you can drag it back to the Values box later on.

	A	B	C	D	E	F
1						
2						
3	Sum of Revenue	Column Labels				
4	Row Labels	Argentina	Austria	Belgium	Brazil	Canada
5	Beverages	$31,262.50	$834,958.80	$58,557.00	$1,713,456.50	$97,165.35
6	Condiments	$1,037.50	$210,231.35	$12,480.30	$225,213.00	$45,965.10
7	Confections	$14,526.45	$200,249.50	$70,883.10	$436,672.82	$143,562.10
8	Dairy Products	$8,521.20	$720,029.70	$76,316.50	$511,225.50	$147,904.20
9	Grains/Cereals	$390.00	$208,684.00	$22,127.00	$80,923.50	$47,910.15
10	Meat/Poultry	$0.00	$173,358.18	$10,360.35	$166,260.46	$34,608.60
11	Produce	$8,933.15	$146,688.25	$12,583.20	$44,377.65	$6,734.00
12	Seafood	$2,752.80	$62,689.80	$10,917.35	$378,582.82	$24,447.00
13	Grand Total	$366,391.20	$17,929,231.65	$1,868,036.16	$22,613,746.08	$3,784,480.00
14						
15						

Figure 22-19: *The worksheet now shows the total revenue by category and customer country.*

Filtering a Pivot Table

As you've seen, pivot tables are a miraculously powerful tool for creating detailed summary tables. The only problem is that sometimes these reports are too detailed—leaving you with summaries that are nearly as detailed as the original table.

To simplify your pivot tables, you may want to restrict them so that they show only a *portion* of the total data. This process is known as *filtering*, and you can do it in three ways: report filtering, slicers, and group filtering.

Note: As you already know, you can create a pivot table that draws its data from a table (the kind you learned to create in Chapter 14). Excel tables support filtering, but their filtering settings have no effect on linked pivot tables. In other words, even if you apply a filter to hide some of the raw data in your source table, all the data still appears in the pivot table. If you want to filter information from the pivot table, you need to use one of the three methods of pivot table filtering.

Report Filtering

Report filtering lets you filter out your data so that your pivot table uses only the rows that really interest you. For example, maybe you want to narrow down your sales summary so it focuses on the sales made in a specific country. To do so, you need to drag the appropriate field (in this case, Ship Country), to the Report Filter box in the PivotTable Field List pane. You'll see it appear on your worksheet, as shown in Figure 22-20.

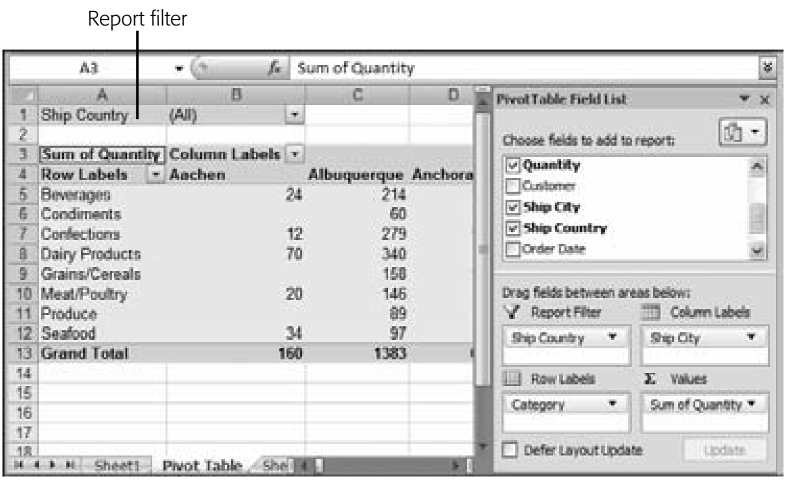

Figure 22-20: *The report filter fields appear just above the pivot table. If you have more than one report filter, Excel gives each one a separate row.*

Note: Before you go any further, it's important to understand one basic fact about report filtering. Any fields you use to perform report filtering can't be in used your pivot table for grouping. So if you're filtering by Ship Country, you can't also group by Ship Country. (If you try it, Excel will remove your grouping when you add the report filter.) This limitation doesn't apply to the slicers or group filters that you'll learn about in the following section.

Once you've added a field to the Report Filter box, it's available for filtering. To actually *set* the filter, click the drop-down arrow in the field box. This action pops up a list with all the values that are currently in that field. You can then choose the one value you want to display. Figure 22-21 shows an example.

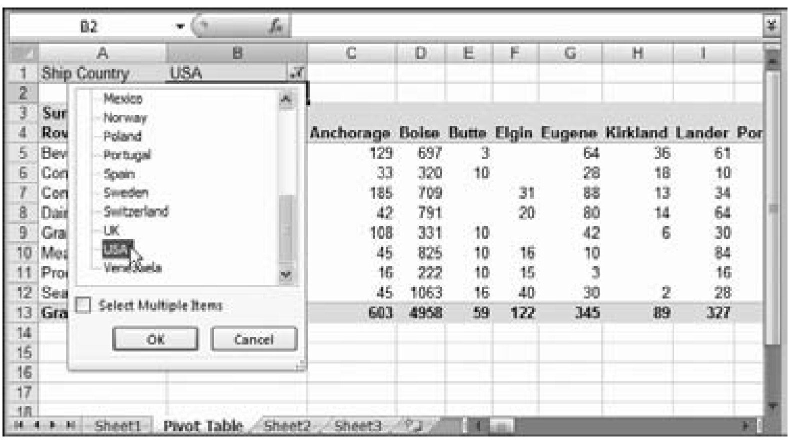

Figure 22-21: *In this example, report filtering is used to hide every country except the U.S. by using the Ship Country field. (Ship Country isn't used as a grouping field, although Ship City is.) The resulting pivot table shows only U.S. cities. You could also add the Customer field to the Report Filter box to show an order summary for specific customers.*

If your report filter offers a long list of potential filtering values, there's no need to spend all day scrolling. Instead, type the first few characters of the value you want in the search box that sites at the very top of the report filter window (Figure 22-21). As you type, Excel will reduce the options in the list. For example, type "Po" into the Ship Country filter and you'll see a list that includes just two choices: Poland and Portugal.

If you want, you can filter for several values at once. To do so, switch on the Select Multiple Fields checkbox at the bottom of the list. When you do, a checkbox appears next to every entry in the list, with a checkmark. Excel uses checked values to create the pivot table, and ignores unchecked entries. Figure 22-22 shows an example.

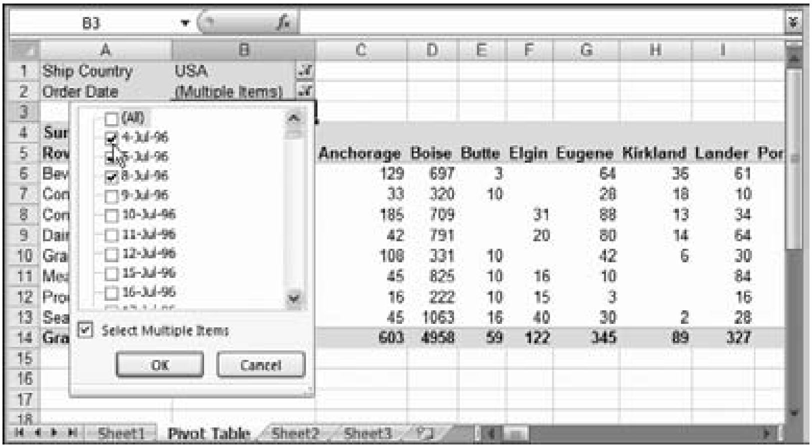

Figure 22-22: *Clear the checkbox next to items you don't want to see in your pivot table. In this example, the pivot table is using the sales records from three specific days.*

To quickly remove a report filter, choose the first item in the report filter list: "(All)."

Slicers

In Excel 2010, Microsoft has added a new way to filter with pivot table *slicers*.

The first thing you need to understand about slicers is that they don't do anything you can't already accomplish with report filters. (And they're a fair bit *less* powerful than the advanced filters you'll learn about in the next section.) However, slicers show filter settings in a unique way—right in the Excel window, so you can change them with a single click (Figure 22-23).

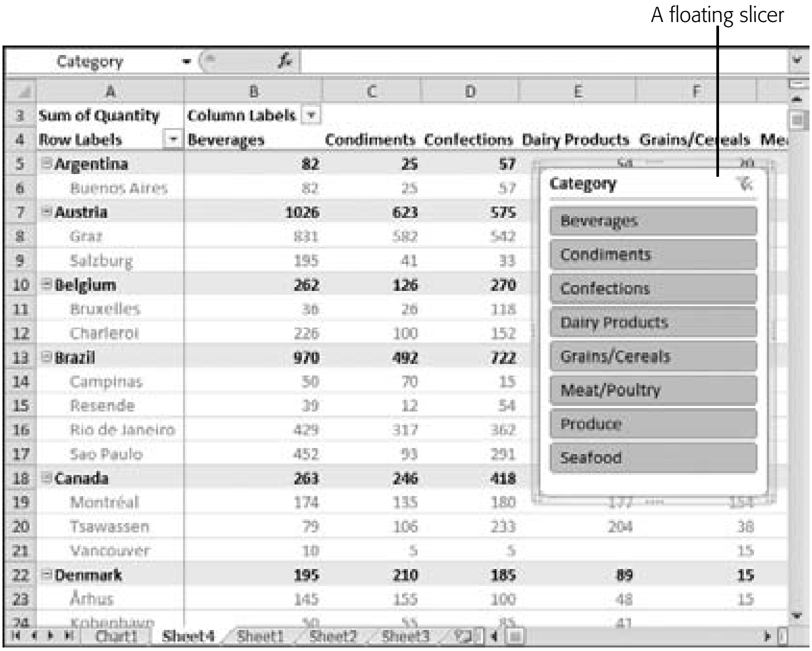

Figure 22-23: *Each slicer gets its own floating window, which you can format, and then drag around the main Excel window. This makes sure your filtering options remain at your fingertips all the time, but it can also clutter your display.*

FREQUENTLY ASKED QUESTIONS

Choosing the Right Filtering

Should I use report filters or slicers?

Slicers create a sort of "data dashboard" effect. They're perfect if you need a way to quickly switch between different filtering choices. For example, if you frequently compare sales in different product categories, slicers make it easy and intuitive. And unlike report filters, you can create slicers for the same fields that you're using to group your pivot table.

However, slicers aren't a clear-cut replacement for report filters. In fact, they have two potential drawbacks. First, slicers take a lot of space in your Excel window, and they can easily obscure parts of your pivot table data, making it difficult to read. This is particularly true if you want to perform filtering on multiple fields—in this case, the report filtering lists are far more compact than the floating slicer windows. Second, slicers don't work well where single fields have a lot of possible values, because all the possibilities won't fit into the floating slicer window at once. You can scroll your way through the list, but there's no handy search box, as there is with report filters.

Finally, if you want more power from your filters—for example, the ability to intelligently pick out ranges of data—neither report filters nor slicers will cut it. Instead, you need the more advanced group filtering feature described on page 650.

To get a closer look at a slicer, you need to create one of your own. Here's how:

1. **Click somewhere inside your pivot table.**
2. **Choose PivotTable Tools | Options→Sort & Filter→Insert Slicer.**

   The Insert Slicers dialog box appears. It lists all the fields in your pivot table (not including any calculated fields you've created), with a checkbox next to each one.

3. **Add a checkmark next to the field you want to use for filtering.**

   In order for a field to make a good slicer, it should have a relatively small number of different values. For example, the Category field makes a great pivot table slicer, because there are only a handful of different category values. The Product field works too, but not as well, because the list of products is much longer. On the other hand, Unit Price and Order Date are poor choices, because there are far too many different values. They just won't fit well in the floating slicer window.

   If you want to add several slicers, you can add a checkmark next to several fields.

4. **Click OK.**

   Excel adds a separate floating window for each slicer.

5. **Drag the floating window wherever you like.**

   Ideally, you'll position each slicer in a place where it won't obscure important pivot table data.

   To move a slicer, move the mouse pointer over the border, so it turns into a four-way arrow. Then, drag as you would to move a floating chart or picture.

   To resize a slicer, move the mouse pointer to a corner or the middle of an edge, so it turns into a two-way resize arrow. Then, drag as you would to resize a floating chart or picture.

   You'll get more help blending slicers into your worksheet in step 7, when you apply formatting.

6. **Use the slicer to apply your filtering.**

   The floating slicer window shows a list of all the unique values in the field you've chosen. In the case of the Category field, for example, you'll see a list of all the categories. Each category appears as a separate button, and each currently visible category is shaded. When you first create a slicer, there's no filtering in place, so every value is shaded. Here's how you can filter it:

   - To filter the pivot table to a single value, click the corresponding button. For example, if you click Beverages, Excel shows only the sales records for beverages.

- To filter the pivot table to show multiple values, hold down the Ctrl key while you click each value (see Figure 22-24).
- To clear the filtering (and show everything), click the tiny funnel-with-an-X icon in the top-right corner of the slicer window.

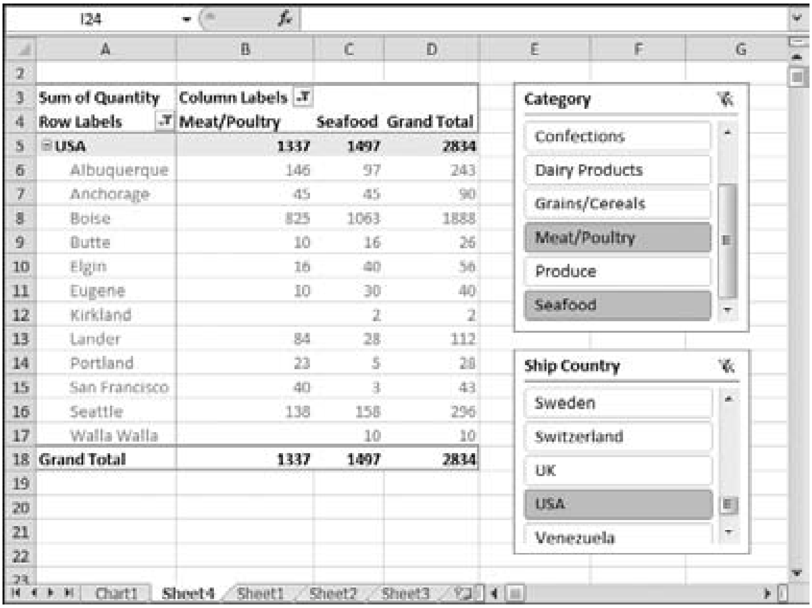

Figure 22-24: *Here, the Category slicer limits the pivot table summary to the Meat/Poultry and Seafood categories. In addition, the Ship Country slicer limits the display to U.S. sales. Changing either of these filter options requires nothing more than a quick click, making slicers a great tool for exploring your data.*

7. **To get that finishing touch, format your slicer.**

   If you plan to keep your slicer around as a permanent part of your pivot table, it's worth the effort to make it fit in. Here are some ways you can refine its appearance:

   - To change the button colors in your slicer, pick a new style from the Slicer Tools | Options→Slicer Styles gallery.
   - To change the title that appears at the top of your slicer, go to the Slicer Tools | Options→Slicer section, and then change the text in the Slicer Caption box. To hide the title altogether, choose Slicer Tools | Options→Slicer→Slicer Settings, and then remove the checkmark from the "Display header" setting.
   - To make a slicer more compact, you can shrink the size of the buttons or tile them into multiple columns (Figure 22-25). To do so, go to the Slicer Tools | Options→Buttons section of the ribbon, and then change the numbers in the Columns, Height, and Width boxes.

- To sort the values in reverse order, right click the slicer, and then choose a different sorting option (like Sort Z to A for alphabetical values or Sort Largest to Smallest for numeric values).

At any point, you can repeat the first four steps to add new slicers. To remove a slicer, click it, and then press Delete.

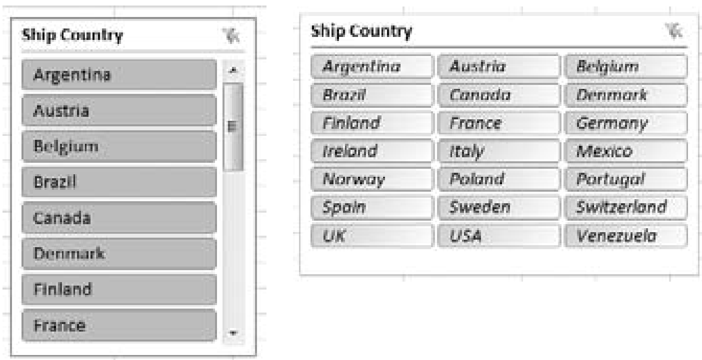

Figure 22-25: *On the left is the Ship Country slicer as it first appears. On the right is the revamped version that uses three columns, smaller buttons, and slightly different styling.*

POWER USERS' CLINIC

Custom Slicer Styles

To get even more control over your slicer formatting, you can create a *custom slicer style*. Best of all, you only need to create the style once, and then you can reuse it with as many slicers as you want. Here's how:

1. **Click the drop-down arrow in the Slicer Tools | Options→Slicer Styles section of the ribbon, and then choose New Slicer Style.**

   The New Slicer Quick Style dialog box appears.

   In the Name box, give your slicer style a name (like "High Contrast Slicer Style").

2. **Pick an entry from the Slicer Element list.**

   Each item lets you configure how a specific part of the slicer looks in a specific state. The most useful options are at the top of the list: Whole Slicer (which applies formatting to the whole shebang) and Header (which formats the title at the top of the slicer). Underneath are more specialized choices that let you set how an item looks when it's selected, when it's currently being hovered over, and when there's no corresponding data for that value in the pivot table.

3. **Click the Format button to format your selected element.**

   The Format Slicer Element dialog box appears. It looks a lot like the Format Cells dialog box but it only contains three tabs, which let you change the font, border, and fill.

4. **Once you've finished choosing your formatting settings, click OK.**

   If you want to format another element in the slicer, return to step 3.

5. **Click OK to finish creating the style.**

   Your custom style appears in the Slicer Styles gallery. You can now apply it to any slicer. You can also right-click it, and then click Modify (to edit your formatting settings), Delete (to remove it), or Duplicate (to create another style with the same settings, which you can then fine-tune).

Group Filtering

Excel also lets you apply filtering to the fields you're using to group your pivot table. (We call it *group filtering*, but Excel doesn't have an official name for this feature.)

Group filtering gives you two abilities. First, it lets you hide and show specific items, just like a report filter. This is a handy feature, because you can't create a report filters for a grouping field.

Second, group filtering lets you create much more powerful filter conditions. For example, you can pick out dates from a set time period, numbers that fall in a specific range, and so on. Using this sort of filter condition, you can home in on the exact subset of data that interests you.

Note: Pivot table filtering is *additive*, which is a fancy way of saying that Excel applies every filter you create at the same time. For example, imagine you're creating a pivot table that analyzes a decade's worth of chocolate sales. Using a report filter, you can focus on sales at a specific store. Using an advanced grouping filter, you can ignore sales of less than $5.

To apply group filtering, just click the drop-down arrow at the right of the Column Labels or Row Labels cell. Figure 22-26 shows exactly where to go.

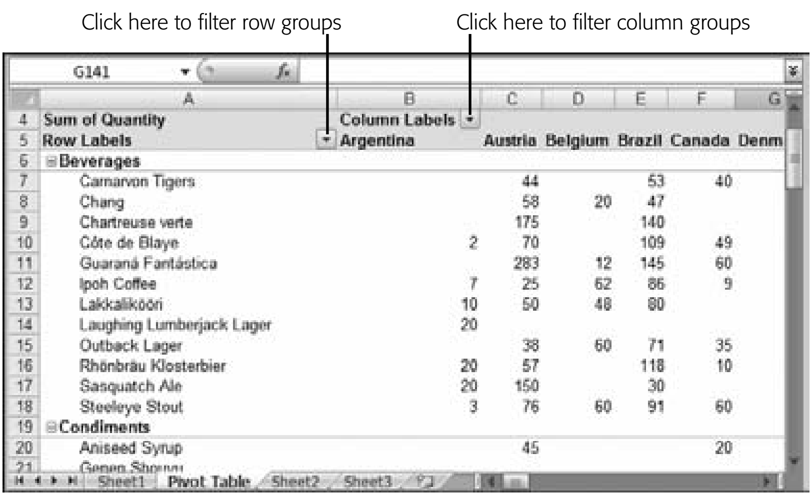

Figure 22-26: *To filter by Category, you need to click the drop-down arrow next to the Row Labels cell. Figure 22-27 shows the drop-down list that appears.*

You can add filter settings to more than one field. You can hide certain category and certain products. However, you need to configure them separately. Here's how:

1. **Click the drop-down arrow to show the filter list (see Figure 22-27).**
2. **In the "Select field" box at the top, choose Category.**

3. **Now uncheck the categories you don't want to see, and then click OK to make it official.**

4. **Click the drop-down arrow to show the filter list again.**

5. **Now, choose Product in the "Select field" box.**

6. **Hide the products you don't want to see (by clearing the checkboxes), and then click OK.**

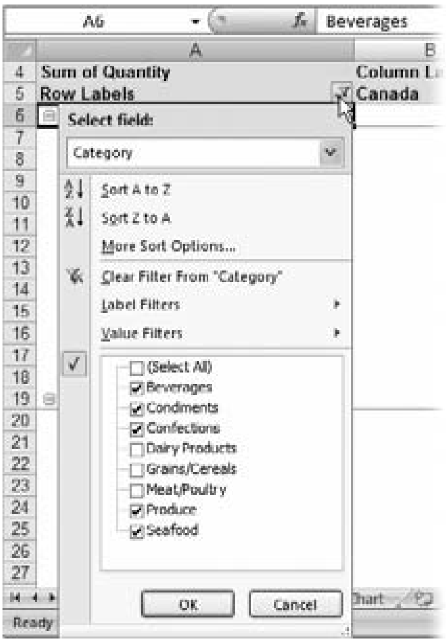

Figure 22-27:
Because there are two row fields (Category and Product), you need to start by choosing the one you want to use for your filter. Then, hide the items you don't want to see by clearing the checkbox next to them in the list. In this example, only five categories appear in the pivot table. Excel also adjusts the subtotals and grand totals accordingly.

The drop-down filter menu has additional options that let you filter out groups according to more sophisticated criteria. You could use these options to show or hide categories (or products) that contain specific text, start or end with a certain letter, and so on. To set these options, click the drop-down arrow next to the Column Labels or Row Labels cell, pick your field, and then choose an option from the Label Filters submenu. For example, to show only categories that begin with a "C", choose Label Filters→Begins With, and then type *C* into the dialog box that appears.

Note: The filtering options in the Label Filters submenu work the same way as the filtering options for structured tables (tables you create with the Insert→Tables→Table command). To learn more about your filtering options, refer to page 411.

You'll also spot a similar submenu named Value Filters. The Value Filters menu lets you filter out information based on the *calculated data*. If, say, you're calculating the sum of all your product orders, you can filter out the products that sold less than 500 units. Just choose Value Filters→Less Than, and then fill in the number 500 when asked. Once again, you've seen this sort of filter before—when you created tables in Chapter 14.

Filtering can be a bit tricky at first, particularly if you have several row fields or several column fields to play with. Before you set a label or value filter, make sure you pick the right field. The Value Filters→Less Than command has a different effect depending on whether you apply it to the Product field or the Category field.

If you apply it to the Product field, you'll see slow-selling products that sell less than 500 units. If you apply it to the Category field, you'll see only categories that have less than 500 units of sales across *all* their products. In the current example, this action would hide every category, because every category has sold significantly more than 500 units.

Note: If you apply a filter on a row field (like Products), your column fields have no effect. Thus, in this example, it doesn't matter that you've subdivided your product sales by country. When evaluating the filter condition, Excel considers the *total* sales for that product.

You can remove your filters using the Clear Filter command in the filter menu. However, if you've added filters to more than one field, you need to remove each filter separately. A faster alternative is to choose PivotTable Tools | Options→Actions→Clear→Clear Filters. Doing so removes every filter and returns your pivot table back to normal, so it shows all the data.

Pivot Charts

Excel lets you create charts based on the data in a pivot table. These charts work more or less the same as ordinary Excel charts (Chapter 17). However, pivot charts are typically very dense, so they warrant a couple of extra considerations:

- Because pivot charts are so dense, many of the specialized chart types don't work well. Instead, stick with simple chart types like column charts and pie charts. One popular choice is the stacked column chart, which helps you see the breakdown of your various groups.
- Before you create a pivot chart, it's often useful to limit the amount of information in your pivot table. Too much information can lead to a chart that's hard to read. Avoid using too many levels of grouping, and use filtering (as described in the previous section) to cut down on the total amount of information that's shown in the pivot table.

Figure 22-28 shows a sample pivot chart.

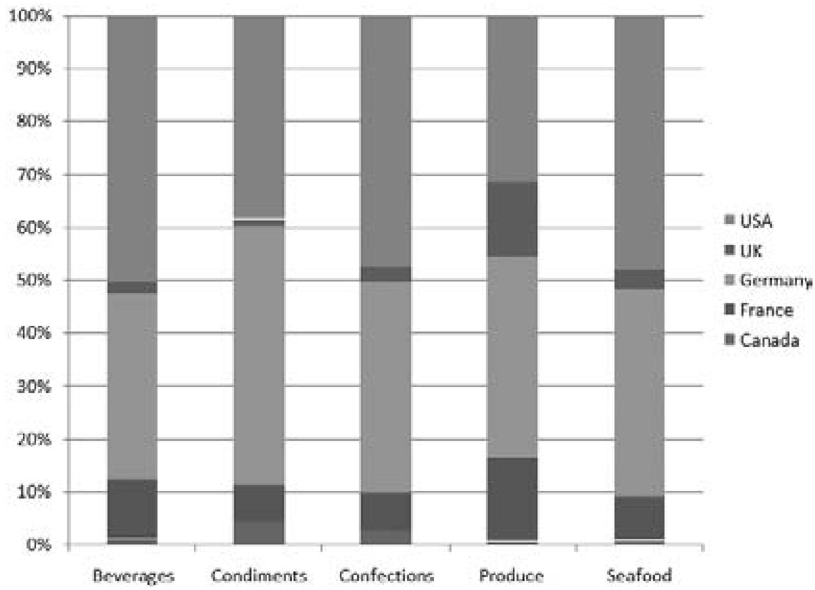

Figure 22-28: *In this example, a stacked column chart shows the breakdown of orders by product category and country. Each bar represents a different category and is subdivided into color-coded sections that represent a country (as detailed in the legend). A pivot chart doesn't necessarily show all the data, because you can use filtering to show just a subset of the total information.*

Creating a Pivot Chart

To create a pivot chart, follow these steps:

1. **Choose PivotTable Tools | Options→Tools→PivotChart.**

   The Insert Chart dialog box appears.

Note: You may have noticed the other options in the Tools section of the PivotTable Tools | Options toolbar tab (OLAP Tools and What-If Analysis). Both of them are turned off, because they apply only when you pull your information out of a high-powered database, as you'll learn in Chapter 23.

2. **Choose the type of chart you want, and then click OK to generate it.**

   Your pivot chart appears on the worksheet.

   When you create a new pivot chart, Excel places it in a floating box. If you want to move it to a separate worksheet, select it, and then choose PivotChart Tools | Design→Location→Move Chart. When the Move Chart dialog box appears, choose "New sheet," pick a sheet, and then click OK.

Manipulating a Pivot Chart

When you select a pivot chart, you'll notice several changes:

- Excel adds several new tabs to the ribbon under a PivotChart Tools heading. These tabs are very similar to the charting tabs you learned about in Chapter 18. They let you change the formatting and layout of the chart, and configure chart elements like titles, axes, and gridlines.
- Excel changes the names of the sections in the PivotTable Field List to help you understand how the different parts of your pivot table are used to create the chart. The Row Labels section becomes Axis Fields. Excel uses these fields to create the different chart categories. The Column Labels section becomes Legend Fields, and Excel uses them to create different series for the chart.
- Excel shows a floating box named the PivotChart Filter Pane. You can use this box to apply filter options directly from your chart.

Like pivot tables, pivot charts are *interactive*. That means you can rearrange your pivot chart on the fly, simply by dragging the fields in the PivotTable Field List from one section to another. As you make your changes, Excel regenerates the pivot table and your pivot chart.

You can also change your filter settings (or apply new ones) right on your pivot chart. Just click one of the filter buttons on the chart to show the familiar drop-down list of filtering options (Figure 22-29). It makes no difference whether you adjust filter settings on the chart using the PivotChart Filter pane or on the original pivot table in the worksheet using the PivotTable Field List pane.

Tip: If you open an Excel 2007 workbook that contains a pivot chart, you won't see any of the filtering buttons, because previous versions of Excel didn't use them. To add filter buttons to an old pivot chart, use the PivotChart Tools | Analyze→Show/Hide→Field Buttons list.

Overall, pivot charts are a specialized tool. You're likely to use pivot tables much more frequently than pivot charts. However, if you need a way to quickly visualize the numbers and get a feel for the most important groups in a summary, a pivot chart can help you out.

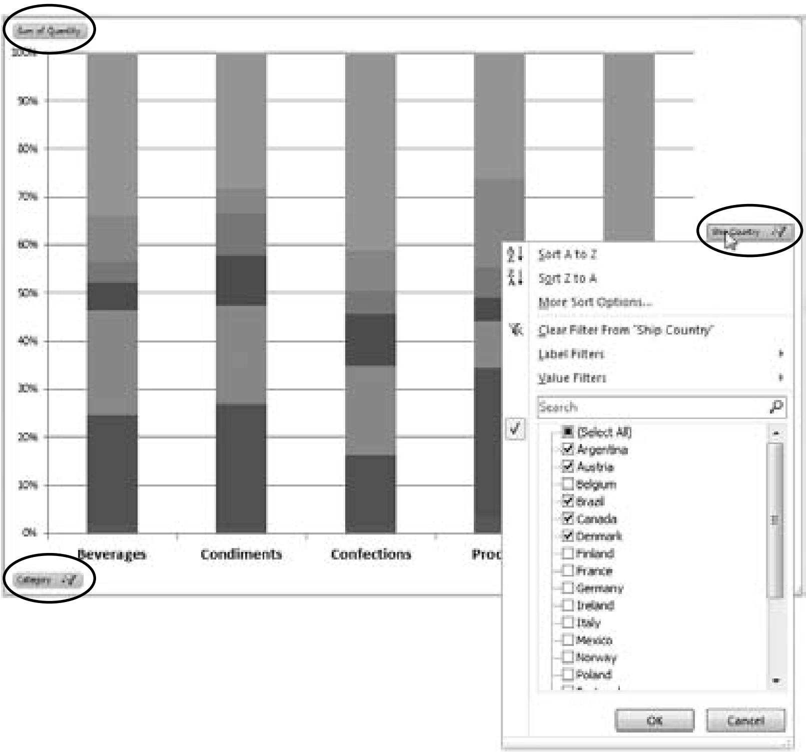

Figure 22-29:
A standard pivot chart includes several filter buttons, arranged in various places. If you find these buttons are cluttering your chart, you can hide some or all of them by choosing the corresponding options from the PivotChart Tools | Analyze→Show/Hide→Field Buttons list.

CHAPTER

23

Analyzing Databases, XML, and Web Pages

An ordinary Excel workbook stores a bunch of information, and uses Excel tools (like formulas and charts) to analyze it. However, a great deal of the world's business information isn't stored in Excel files. Instead, it sits inside *databases*—organized catalogs of information managed by high-powered computer software.

Typical mid-size companies (or bigger ones) store product catalogs, customer lists, and purchase orders in databases. When they need to change this information (or just take a look at it), they use another program that talks to the database. Databases are extremely fast, secure, and durable. They're the best place to put mission-critical information.

At first glance, this approach (storing big volumes of information in a database) seems to cut Excel out of the loop. However, Excel's greatest strength isn't storing information, but *analyzing* it. In fact, Excel is the perfect tool for analyzing information, no matter where it comes from. Your company might track orders using a custom application that dumps its results into a database. If you want to find out what all the numbers amount to, and make a prediction about your company's future, you can pull that info out of a database (using a process called *querying*), and then start charting away. In other words, your Excel workbook gets a copy of the information that's stored in the database so that you can put it under the microscope. Best of all, you can refresh your workbook at any time by getting the latest information from the same database, so your workbook never goes out of date.

In this chapter, you'll learn how to pull information out of a database so you can analyze it in a workbook. Then, you'll consider how to pull information out of an *XML* (Extensible Markup Language) file, which gives you even more options for plugging business data into Excel. Finally, you'll look at a messier but occasionally useful technique called web queries, which lets you extract information from a web page and transfer it to your worksheet. All of these approaches have one thing in common: they use Excel to analyze information that's stored in another place.

Excel and Databases

Databases provide the best way to store massive amounts of data, like a product catalog for an online store or a customer list for a mail order company. Databases break this data into separate tables and let you search it with unrivaled flexibility. Databases are particularly adept at answering questions like "What products did this customer order last week?" and "How many employees have completed sensitivity training?"

The world has two basic types of databases:

- **End-user databases** are designed for small-scale or single-person use. These databases store data and provide people-friendly tools for searching and managing the information they contain. Microsoft Access, which comes with some versions of Office, is the best-known end-user database.

Tip: For the full scoop on Microsoft Access, check out *Access 2010: The Missing Manual*.

- **Server-side databases** work on a larger scale by storing data and feeding it to other programs and websites. Ordinary mortals don't use server-side databases directly. Instead, they use applications that, in turn, rely on databases. For example, when you browse Amazon.com, the Amazon web page code is hard at work getting product information and customer reviews from some of the most high-powered databases on the planet. Microsoft SQL Server is one well-known server-side database.

In Excel, you can just as easily extract information from a small-scale database like Access as you can get data from a full-fledged database server running SQL Server. You'll find a few minor differences—for example, if you want to use Access data, you need to supply only the file name, whereas you need to know the server name, user ID, and password to get into a typical SQL Server database. However, the process of searching and importing the information is the same in both cases.

Note: Data querying is a one-way street. You can extract the data and analyze it, save a copy on your PC, and even modify the information in your worksheet. However, Excel doesn't provide any tools to let you route your changes back to the original source.

Connecting to an Access Database

Before you can get information out of a database, you need to create what's known as a *data connection*. The connection tells Excel how to connect to your database server (or database file) so that it can extract the information it needs. You need to perform this step only once for each data source you want to use. (If you don't have a database handy, use the Northwind.accdb Access database file, which you can find on the Missing CD page at *www.missingmanuals.com/cds*.)

The following steps walk you through the process of creating a new connection for an Access database:

1. **Choose Data→Get External Data→From Access.**

   The Select Data Source dialog box appears, which lets you browse through the files on your computer.

2. **Find the Access file you want to use, select it, and then click Open.**

   Access database files have the file extension .mdb or .accdb.

Note: Keep in mind that this connection remains valid as long as you don't move the file. If you move the file, you'll need to modify the connection, as explained on page 668.

Once you pick a database file, the Select Table dialog box appears (Figure 23-1).

Figure 23-1:
A typical database includes several tables. Each table contains a set of records that are structured in a particular way. For example, the Northwind database (which you can find on the Missing CD page at www.missingmanuals.com/cds*) includes a table named Order Details, which is selected here.*

Note: Along with the list of tables, you may also see Access *queries*. Queries present a customized view of a table, which can apply sorting, hide columns, or filter out unimportant records. Queries are a great shortcut for getting combined results from separate tables and for calculating and summarizing database date before it reaches Excel. For example, the worksheet shown later (in Figure 23-4) uses a query to get a list of sales summarized by product category.

3. **Choose a table from the list, and then click OK.**

   If you're using the Northwind.accdb database from the downloadable samples, you can use any of the tables, including Customers (for a list of customers with contact information) and Order Details (for a list of purchased items from every order made by every customer). Once you choose a table, the Import Data dialog box appears (Figure 23-2).

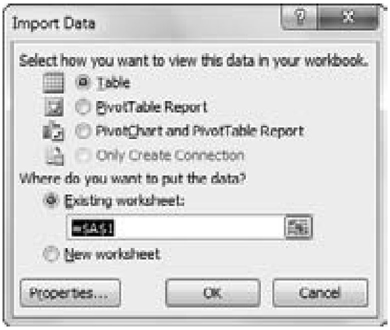

Figure 23-2:
After you've picked a database and a table, use this dialog box to tell Excel where in your workbook it should put the data.

4. **Choose whether you want to see the data as an ordinary table or a pivot table.**

   Usually, you'll choose to insert the data as a table. That choice gives you all the features you learned about in Chapter 14 for working with the data, including sorting, filtering, and more.

   If you're dealing with extremely large amounts of information that you want to summarize, a pivot table is more practical. That way, your worksheet doesn't actually store the raw data (which would bloat your workbook file, and slow down Excel). Instead, it uses the information to create the pivot table. This method's only disadvantage is that if you decide to rearrange your pivot table later, Excel needs to contact the database again and fetch all the data, which can be a bit slow. To learn more about pivot tables, see Chapter 22.

Tip: If you decide to use a pivot table with a slow database, you should consider using the Defer Layout Update option, which is discussed on page 637.

5. **Choose whether you want to place the data in a new worksheet.**

   You can choose "Existing worksheet" to place your table somewhere in the existing workbook. Then, switch to the worksheet you want to use (using the worksheet tabs), and then click the top-left corner where your data should appear. Be sure to leave plenty of space for all your data.

You'll find it easier to choose "New worksheet", which keeps your imported data separate from the data that's already in your workbook. If you opt for this choice, you'll probably want to move and rename the worksheet after Excel creates it, as described in Chapter 4.

6. **Click OK.**

   Excel pulls the information out of the database and inserts it into your workbook, as shown in Figure 23-3. Next you can apply formulas, create charts, and start analyzing. You can also use the Table Tools | Design tab to format the table by applying a different table style.

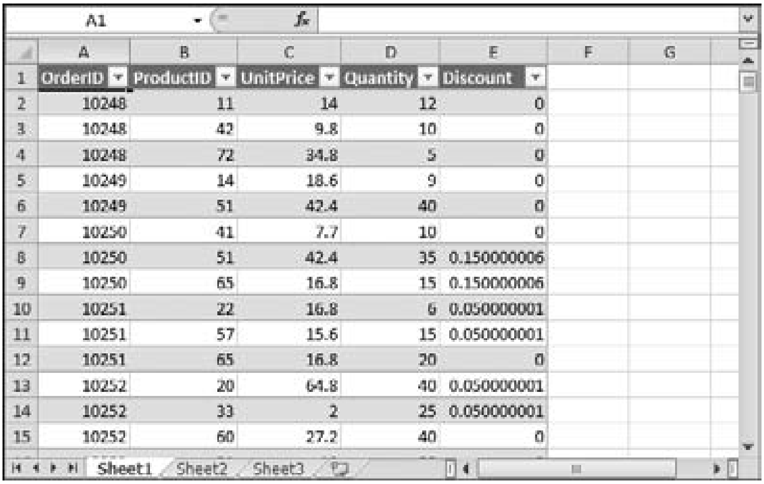

	A	B	C	D	E	F	G
1	OrderID	ProductID	UnitPrice	Quantity	Discount		
2	10248	11	14	12	0		
3	10248	42	9.8	10	0		
4	10248	72	34.8	5	0		
5	10249	14	18.6	9	0		
6	10249	51	42.4	40	0		
7	10250	41	7.7	10	0		
8	10250	51	42.4	35	0.150000006		
9	10250	65	16.8	15	0.150000006		
10	10251	22	16.8	6	0.050000001		
11	10251	57	15.6	15	0.050000001		
12	10251	65	16.8	20	0		
13	10252	20	64.8	40	0.050000001		
14	10252	33	2	25	0.050000001		
15	10252	60	27.2	40	0		

Figure 23-3: This worksheet shows a portion of the results extracted from the Order Details table of the Northwind database. As you can see, Excel isn't smart enough to determine the number formats for your data, so it's up to you to apply the Currency format to the UnitPrice column and the Percentage format to the Discount column. Furthermore, a table might not have all the information you want–the Order Details table, for example, doesn't tell you the shipping address of the person who made the order. To get that sort of information, you need to create a join query in your Access database.

Refreshing Data

Once you've pulled some information out of a database, you can work with it in the same way you work with ordinary Excel data. You can edit cells, apply conditional formatting, cut and paste chunks of data, and so on. Assuming you've placed your data into an Excel table, you can also use sorting, filtering, and all the other handy table features described in Chapter 14.

One of Excel's database support's nicest aspects is that it can refresh the data at any time. If you suspect that the data's changed in the database, just move somewhere in the list, and then choose Table Tools | Design→External Table Data→Refresh All.

Excel wipes out your current table and replaces it with the up-to-the-minute information taken from the database, without disturbing any other content that might be in your worksheet. You can use this technique to build a "window" that can analyze the most recent information in your database, as shown in Figure 23-4.

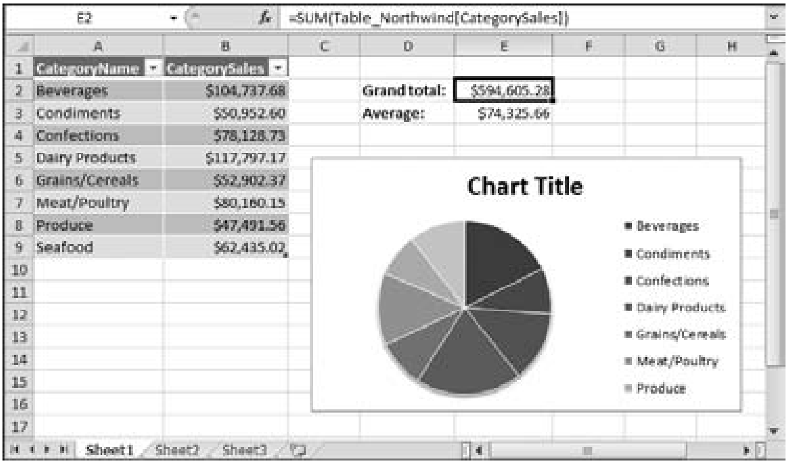

Figure 23-4: *In this example, the data you've pulled from the database is used to calculate a commonplace sum and average, and to create a chart. Here's the neat part: If you refresh your data (Table Tools | Design→External Table Data→Refresh), the data in the list changes, and the formulas and chart update themselves automatically to take the new data into account. To get this information, the worksheet uses the "Category Sales for 1995" query. Another option is to get the raw sales information from the OrderDetails table and use Excel to subgroup it, possibly using a pivot table.*

Tip: A workbook can have as many connections as you want. To refresh them all at once, choose Table Tools | Design→External Table Data→Refresh→Refresh All.

If you're using a database that you can access quickly (like an Access file on your computer), the update process happens almost instantaneously, and you might not even notice the refresh. If you're grabbing a large amount of information from a heavily taxed server on your network, this process takes more time.

The Excel database features' only limitation is that you're always working with a *copy* of the data that's stored in the database. Even if you change the information in your worksheet, you can't change anything in the database. As a result, the next time you refresh your data, Excel wipes out all your changes and replaces them with whatever content is currently in the database. Interestingly enough, this limitation doesn't apply to formatting, which Excel preserves even after you update. In other words, if you embolden a few rows, and then perform a refresh, the same rows are bolded even after the new data appears.

Note: Remember, Excel's database features are designed to let you analyze information, not change it. If you need to modify a database, you need a different tool. If you want to modify the information in an Access database, you need to use the Access software.

If you ever get tired of your database and decide to break free, just click somewhere in your table and choose Table Tools | Design→External Table Data→Unlink command. This action changes your database-linked table into an ordinary table. In other words, you can't use the Refresh command any longer. Before you take this step, be certain that you don't want to connect to the database any longer. Once you use the Unlink command, all the database information disappears from your workbook, and there's no going back. (If you do want to get the most recent data from the database, you have to create the connection all over again.)

POWER USERS' CLINIC

Controlling the Refresh Process

You can change a few details about the refresh process using the External Data Properties dialog box (Figure 23-5). To get there, move into your table, and then choose Table Tools | Design→External Table Data→Properties. You can use this dialog box to control whether or not Excel keeps your formatting, sorting, filtering, and row width settings when it gets new data. (Ordinarily, it does.)

You can also use this dialog box to tell Excel what to do if it finds that records have been removed or added when it performs a refresh. Ordinarily, it inserts the new data and removes anything that's been deleted (the first option in the list), which is by far the simplest choice. Alternatively, you can choose one of two much more complex options that insert the new information in new rows, so you can take stock of what's changed.

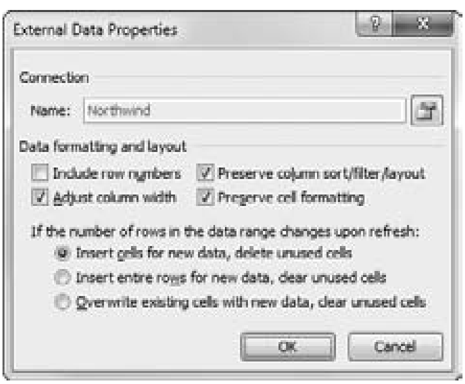

Figure 23-5:
Use "Include row numbers" to add an extra column that numbers each row, which helps you keep track of them when you perform refreshes. Clear the other checkboxes if you don't want to keep your custom formatting when you refresh your data. Finally, choose one of the three options at the bottom to control what happens if Excel finds out that records have been deleted or added when it performs a refresh.

Data Source Security

In order to provide the database update feature, Excel stores the information about your data connection in your Excel file. You can save your workbook, open it later, and perform a refresh to get the most recent information.

Unfortunately, Excel is also relentlessly paranoid. When you open a workbook that has a data connection in it, Excel keeps your data connection closed and shows you a message (Figure 23-6). Excel takes this step because it knows there's an incredibly slim chance that you're opening a workbook created by a devious computer hacker. This workbook could contain a database connection that doesn't just read data, but somehow executes a command that changes or even deletes it. To guard against this possibility, Excel stops you, temporarily, from using the connection.

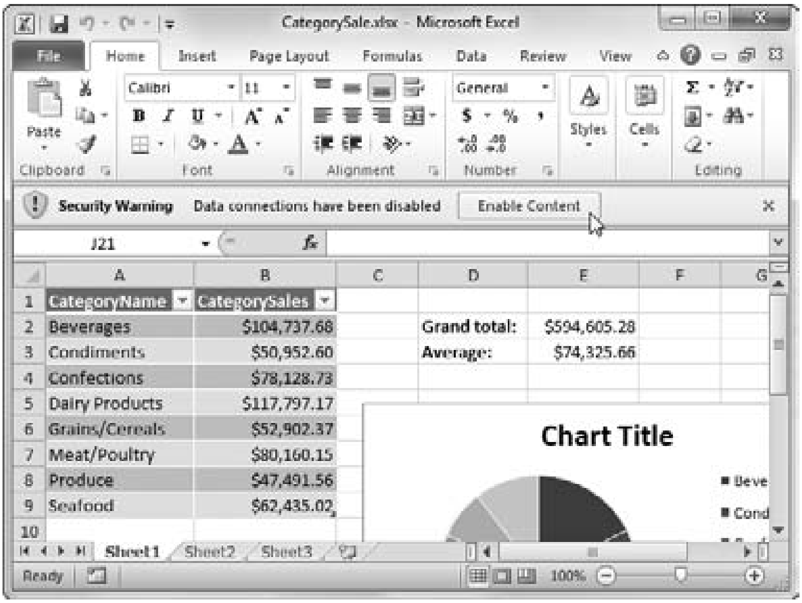

Figure 23-6: *This warning message tells you your workbook contains one or more data connections, and that Excel closed them just to be safe. You can switch them back on by clicking the Enable Content button.*

Of course, if you've created the workbook yourself (and you know that you aren't a computer hacker), there's no way your workbook could contain any devious trickery. Therefore, your connections are completely safe, and there's no reason you shouldn't turn them on. To do so, click the Enable Content button. This tells Excel to treat your workbook as a trusted document, so it won't hassle you the next time you open it. (Page 804 has all the details about Excel's trusted document system.)

Connecting to a SQL Server Database

The process for connecting to a full-fledged SQL Server database is similar to the process for connecting to an Access database—you just need to supply more information. To get the process going, follow these steps:

1. **Choose Data→Get External Data→From Other Sources→From SQL Server.**

   The Data Connection Wizard starts up (Figure 23-7).

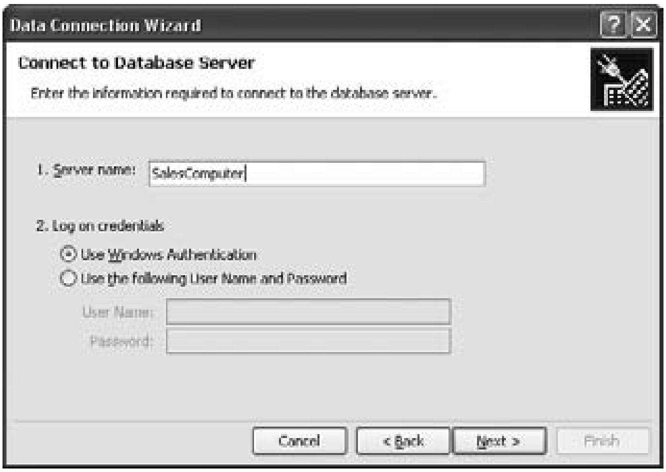

Figure 23-7:
In the first step of the Data Connection Wizard, you need to choose your database server and supply login information.

Note: SQL Server isn't the only server-side database that you can connect to using Excel. In fact, you can pull information out of just about any database that's been created in the last 10 years. Get started by choosing Data→Get External Data→From Other Sources→From Data Connection Wizard. When the Data Connection Wizard appears, choose Other/Advanced, and then click Next to see a long list of supported databases. To get the exact settings you need to use, check with your organization's resident network administrator.

---

2. **Enter the database server's name.**

   The server is the computer that's running the database software, which has to be accessible on your network. (Check with the nearest network administrator for help.) If the SQL Server database software happens to be on your computer, just enter the word *localhost*, which is code for "this computer right here." If you've installed the slightly scaled-down but absolutely free version of SQL Server that's known as SQL Server Express (and it's on your computer), you should use the server name *.\SQLEXPRESS*. Reminder: Don't forget the initial period.

3. **Supply your login information.**

   Usually, you'll choose Use Windows Authentication, which tells Excel to log onto the database using your current Windows account. Depending on the setup of your network and database, that option may not work, and you may need to choose "Use the following Name and Password" instead. If so, supply the correct user name and password in the appropriate boxes.

4. **Click Next to move to the next step.**

   Now you need to pick your data (Figure 23-8).

5. **Choose the database you want from the "Select the database" list.**

   A SQL Server computer can hold hundreds (or more) databases. Every database is a collection of tables. You need to pick the database that has the table you need.

6. **Make sure the checkbox next to the "Connect to a specific table" setting is *not* turned on.**

   When you choose a database, Excel shows you a list of all the tables that are in the database (see Figure 23-8). However, don't pick the table you want just yet. You may want to reuse the same connection to connect to different tables in the same database. In order to do this, you need to wait just a few more seconds before you pick your table.

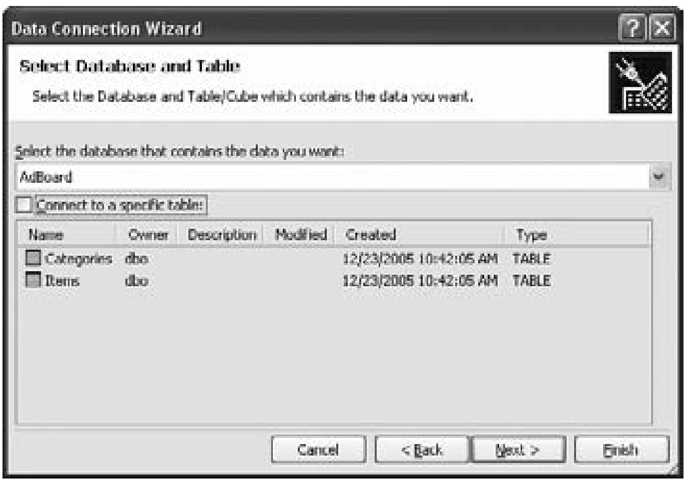

Figure 23-8:
A database is chock-full of information. You need to pick the database you want to use (using the drop-down list at the top). In this example, a connection is being created to the AdBoard database, which contains two tables (Categories and Items).

7. **Click Next.**

   This move gets you to the last step (Figure 23-9).

Figure 23-9:
In the final step, you set some descriptive information about your connection.

8. **Fill in the Description and Friendly Name information.**

   When you create a new connection, Excel saves it in a special location on your computer called My Data Sources. That way, if you want to use the same connection again, you don't need to go through the Data Connection Wizard. (See page 665 for the full details.)

   To make the connection easier to identify, you should fill in the Description and Friendly Name boxes. (The "friendly name" is a name that makes sense to ordinary humans. The actual file name that's used to store the data connection on your computer may be considerably more cryptic, although you're free to change that, too.)

9. **Click Finish to end the wizard.**

   Now, Excel shows you all the tables in your database and asks which one you want to use (Figure 23-10).

10. **Pick a table, and then click OK.**

    Excel shows the Import Data dialog box that you saw earlier (Figure 23-2), which lets you choose where you want the new data. Once you make your selection, Excel contacts the database server, fetches the contents of the table you selected in the previous step, and then inserts it into your workbook. From this point on, you can refresh the data and modify the connection using the techniques you learned earlier in this chapter.

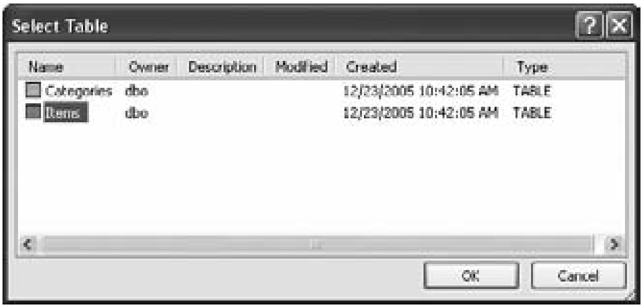

Figure 23-10:
The AdBoard database has two tables. In this example, Excel will insert the contents of the Items table into the worksheet.

POWER USERS' CLINIC

Creating a Data Source That Updates Automatically

Excel leaves a few of its advanced database connection options off the Data Connection Wizard, including a nifty timesaving tool: the ability to have Excel automatically suck in a fresh load of data at regular intervals. To make Excel do this updating, follow these steps:

1. Choose Data→Connections→Connections. The Workbook Connections dialog box appears, with a list of all of the connections in your workbook.
2. Select the connection you want to modify.
3. Click the Properties button to open the Connection Properties dialog box, which has every setting you can change (Figure 23-11).

If you want to tell Excel to refresh the data in your workbook periodically, add a checkmark in the "Refresh every" box, and then enter your preferred number of minutes. Excel performs these updates automatically, without bugging you, so you can keep working on other tasks.

Of course, if you've changed imported data in your spreadsheet, your changes will be lost when Excel performs its refresh. So use the "Refresh every" section with caution. Don't use it with an extremely slow database, or the performance of your workbook may suffer. And don't make the refresh interval too small, or you'll create an unnecessary amount of extra work for your database server. (Remember, there's always the Data→Connections→Refresh All button to trigger an immediate refresh.)

Reusing Your Database Connection

Once you've created a connection, Excel stores the connection in your workbook so that you can reuse it—a real timesaver if you want to grab *another* table from the same database. To do so, just choose Data→Get External Data→Existing Connections. The Existing Connections dialog box appears (Figure 23-12). You'll find your existing connection(s) in the "Connections in this Workbook" section at the top of the list. Click Open to use it again, without being forced to re-enter all the same settings.

Figure 23-11:
The Connection Properties window has several advanced database connection options. You can choose to have Excel retrieve the latest information from the database whenever you open the workbook using the "Refresh data when opening the file" setting. You can also tell Excel to automatically refresh your data every few minutes using the "Refresh every" setting.

Figure 23-12:
The Existing Connections dialog box gives you a head start for connecting to the AdBoard database that you created previously. If you go ahead and click Open, Excel asks you which table you want to use. If you want, now's the time to import a different table in a different place in your workbook.

This technique works for *all* connections, whether they point to Access databases or more powerful server-side databases like SQL Server. However, if you're using a server-side database, there's one more trick you can use. You can tell Excel to make your connection information available to *other* workbooks on your computer. So you have yet another way to cut down on the amount of time you spend creating (or recreating) database connections. Here's what you need to do:

1. **Choose Data→Connections→Connections.**

   The Workbook Connections dialog box appears, with a list of all of the connections in your workbook.

2. **Select the connection you want to reuse, and then click the Properties button.**

   This action shows the Connection Properties dialog box (Figure 23-11), where you can tweak all the connection details you supplied in the Data Connection Wizard (plus a few more).

3. **Choose the Definition tab, and then place a checkmark in the "Always use connection file" checkbox.**

4. **Click OK to save your changes.**

   Excel warns you that you'll now actually have two connections (one in your workbook and one stored on your computer). There's a possibility for confusion here because if you modify the connection in your workbook, it doesn't affect the copy that's stored on your computer. Click Yes to go ahead.

5. **Click Close to return to Excel.**

Now, you can use the Data→Get External Data→Existing Connections command from any workbook. In the Existing Connections dialog box, you'll see your connection in the "Connection files on this computer" section.

Note: Behind the scenes, Excel stores your database connections in a special My Data Sources folder deep within the Documents folder on your computer. For example, if you're running Windows Vista or Windows 7 and your user name is MattG, you might find the database connection that was created in the previous example in the file *C:\Users\MattG\Documents\My Data Sources\SalesComputer_AdBoard.odc.*

To delete a connection, you need to remove the connection file from the My Data Source folder. The easiest way to do this is to open the Existing Connections dialog box (choose Data→Get External Data→Existing Connections). Then, click the "Browse for More" button to open a dialog box that shows all the files in the My Data Sources folder. When you spot the file for your data connection, right-click it, and then choose Delete.

Understanding XML

Databases aren't the only place you can find massive amounts of information. All kinds of specialized software programs store vast quantities of data in their own way. More often than not, when you need to get information out of (or into) one of these programs, you'll use *XML*.

XML is an all-purpose system for structuring and organizing data in a file. XML lets you exchange information with just about anyone, and send your spreadsheet data to other businesses that don't use Excel, or analyze raw information created with other programs.

For example, instead of saving data in Word documents, Excel spreadsheets, or ordinary text files, you can save data in an XML file. XML alone *sounds* pretty modest, but this simplicity is deceiving. Two factors make XML really special:

- **XML is flexible**. You can tailor XML to store pretty much any type of information: pictures, product catalogs, invoice data, receipts, catalog listings, the maintenance specs for every Dodge Minivan ever built, and on and on.
- **XML is widespread**. Computer programs written in different programming languages (Java, Visual Basic, C++, and so on), or running on different operating systems and computer hardware (Windows, Mac, Linux, and so on), can all use XML in exactly the same way. XML is a perfect solution for exchanging information between people, companies, and even computers programmed to send data to one another automatically. (Features like this last one cause supply-chain management types to start drooling when they talk about XML.)

Note: In a confusing move, Microsoft's decided to limit XML support to certain editions of Office 2010. If you don't have Office's full Professional or Enterprise Editions, you may want to stop reading now. All other Office versions (including Standard, Small Business, Student and Teacher) omit the XML goodies that this chapter explores.

What Is XML, Really?

Contrary to what many people believe, XML is *not* a data format (as is, for example HTML, the format used to create web pages). If XML were an ordinary data format, it wouldn't be nearly as useful, because no matter how good a format is, it can't suit everyone. Even though almost every company needs to create invoices, most companies wouldn't be happy with a generic format for storing invoice information. One company might need to track customer names, while another might track customer ID numbers. The bottom line is that most companies need to store slightly different data in slightly different ways. A one-size-fits-all solution is pretty much always doomed to failure.

So if XML isn't a data format, what is it? Technically, XML is a *meta-language*, which is a fancy way of saying that XML is a language for creating *other* languages. XML does this by setting out a few simple rules that let you build your *own* data format that's just right for *your* data.

For example, Acme Company can build an XML format for invoices and call it AcmeInvoice. Meanwhile, Budget Company can build *its* own XML invoice format and call it BudgetInvoice. Even though both these formats are designed to store invoice information, they can contain completely different kinds of data. XML's strength is its flexibility.

At the same time, XML's flexibility can create problems. Imagine that a bank named Worldwide Green sets up a system to automatically process XML invoices in a specific format. The system works smoothly until Acme Corporation sends along its own homegrown invoice. Even though Acme's invoice uses XML, it doesn't conform to the XML that the bank expects, and so it gums up the bank's automated invoice-processing application. Suddenly, XML doesn't look so useful.

The bottom line: XML holds the *promise* of universal data sharing—but if you don't create some rules and follow them, you're left with a bunch of incompatible formats.

Note: XML is really quite simple. However, you'll find a slew of other standards with names like XML Schema and XSLT that work in conjunction with XML, and provide solutions for validating XML, searching XML, transforming XML, and so on. These other standards are quite complex and aren't discussed in this book. For more information, pick up a book like *Learning XML* by Erik Ray (O'Reilly), or see the website *www.w3schools.com/xml*.

Three Rules of XML

To better understand how to configure Excel to handle XML, look at a simple example. Technically, you don't need to know what XML looks like in order to use Excel's XML features, but the more you understand, the less confusing life will be. In this section, you'll learn the three most important rules that shape all XML documents. If you already know a little about XML, feel free to skip ahead.

By the way, good news before you even start: XML is written in a text-based, human-readable format. That means you can use a program like Notepad to crack open an existing XML file and get a basic idea of its format and structure. You can even write an XML file from scratch using Notepad. You can't do the same with the average Excel spreadsheet file, because they're stored in a binary format that you can read only when you're looking at the file in Excel. (If you open an ordinary Excel file in Notepad, you see a jumble of indecipherable symbols.)

The prolog

All respectable XML documents start with something called a *document prolog*. This bit simply announces that what you're looking at is an XML document. It can also indicate the *encoding* of the document, which sometimes specifies that the document uses a special character set (for example, a non-English alphabet).

Here's a typical document prolog, indicating that this document uses Version 1.0 of the XML standard (the most prevalent version):

```
<?xml version="1.0" ?>
```

If you're creating an XML document by hand, you should make sure to place the document prolog as the very first line of the file.

Elements

The *element* is the basic building block of any XML document. Elements are information containers. If you wanted to store a person's name, you might create an element called Name. (For more on the infinite variety of elements anyone can create, see the box on page 674.)

A typical element is composed of a start tag and an end tag. The actual information goes between these two tags. You can easily recognize start tags and end tags because they use angle brackets <>. For example, here's one possible start tag:

```
<Name>
```

This tag marks the start of the Name element. The end tag looks almost identical, except it begins with the characters </ instead of just <. Here's what you need to end the Name element:

```
</Name>
```

To actually store some information in an XML document, you just insert the content between the start and end tag of an element. Here's how you might store someone's name in an XML document:

```
<Name>Patrick</Name>
```

You could create a list of names by putting one <Name> element after the other, or you could add other elements that store different types of information, like address, title, employer, and so on. You put all these tags together in a file to make an XML document.

Nesting

So far, you've seen examples of XML elements that contain text. You can also create an element that contains one or more additional elements. This is a basic principle for organizing information in XML.

UP TO SPEED

A Closer Look at Tags

Tags follow fairly strict naming rules. Tags can be of any length, are case-sensitive, include any alphanumeric character and hyphens (-), underscores (_), and periods (.). You can't use other special characters, including spaces, and the tag name *must* start with an underscore or letter. XML documents also support characters from non-English alphabets.

The most important thing you should understand about tags is that it's up to you to create them. If you decide that you need to store a list of names, you might create an XML format that uses a <Name> tag. Meanwhile, someone else might decide to track name information by creating another XML format that uses elements like <firstName> and <lastName>. These two elements might store the same type of information as your <Name> element, but they're different, and a document written with the <firstName> and <lastName> tags isn't compatible with your documents.

Because so many possible XML formats exist, a lot of intelligent people have invested a lot of time and energy in trying to create ways to define and manage different XML formats. Also, companies and organizations have come together to define specific XML standards for different industries. (If you search on the Internet, you'll find predefined XML formats for law, science, real estate, and much more.) When you use an XML document in Excel, you probably won't be responsible for creating it. Instead, you'll be given a document that's already in a specific XML format, and you'll just be interested in retrieving the information that it contains.

Imagine you want to keep track of several people's names and ages. The following format isn't especially clear, because it's hard to tell which person connects to which age:

```
<Name>Lisa Chen</Name>
<Age>19</Age>
<Name>Bill Harrison</Name>
<Age>48</Age>
```

Better to group the <Name> and <Age> elements together for each person, and put them inside *another* element. Here's an example:

```
<Person>
    <Name>Lisa Chen</Name>
    <Age>19</Age>
</Person>

<Person>
    <Name>Bill Harrison</Name>
    <Age>48</Age>
</Person>
```

Here, the two <Person> elements each represent a distinct individual. Information about each person is stored in <Name> and <Age> elements that are *nested* inside the appropriate <Person> element.

There's no limit to how many layers deep you can nest information, making this method of organizing information extremely flexible. In fact, it's part of the reason that XML can work with so many different types of data.

XML imposes one more rule. Every document must start with a single element that you place right after the document prolog. You place all the other content inside this element, which is called the *root* or *document element*. So far, the examples you've seen are only excerpts of XML. The following listing shows a complete, valid XML document—a list with information about two people—that starts off with the document element <PeopleList>:

```
<?xml version="1.0" ?>
<PeopleList>
    <Person>
        <Name>Lisa Chen</Name>
     <Age>19</Age>
    </Person>

    <Person>
        <Name>Bill Harrison</Name>
        <Age>48</Age>
    </Person>
</PeopleList>
```

You could enhance this document by adding more <Person> elements, or by adding different elements to track additional information about each person.

You've probably noticed that these XML examples indent each level of elements. That indentation makes the overall structure easier to read, but it's not required. In fact, applications that read XML (including Excel) ignore all the white space between elements, so it doesn't matter if you add spaces, tabs, and blank lines. In fact, as far as computers are concerned, the previous document is exactly the same as the following, much less human-friendly version:

```
<?xml version="1.0" ?>
<PeopleList><Person><Name>Lisa Chen</Name><Age>19</
Age></Person><Person><Name>Bill Harrison</Name><Age>48<
/Age></Person></PeopleList>
```

XML Files and Schemas

As you've already learned, a file is one place you can store XML documents. But you can just as easily place XML documents in databases or other storage locations. In fact, sometimes XML data isn't stored anywhere—instead, you just use it to send information between applications over the Internet. However, when you use XML with Excel, you're always using XML files (unless your company's created a custom solution using Excel's heavy-duty programming features). Most XML files have the extension .xml. It makes perfect sense to take the person list document shown earlier and place it in a text file named PersonList.xml.

Tip: The Excel XML features make the most sense when you have information that's already in an XML format. If you *don't* have XML information that you need to work with, there's really no reason to play around with Excel's XML features.

There's another type of XML document that's extremely important: XML *schemas*. Schemas are designed to solve a common problem—namely, defining the rules for a specific XML-based format. A schema may indicate the element names you can use, how you can arrange the elements, and the type of information each element can contain. An XML-friendly application like Excel can use the schema to verify that an XML document uses the right structure and contains the appropriate content. In an ideal world, every time a company created an XML format, they'd write an XML schema that defines it. (You probably won't be surprised to learn that this doesn't always happen.) Excel doesn't include schemas of its own, so it's up to you to supply the schema based on the XML you want to use.

This book doesn't look at XML schemas in detail, because they're a more complicated standard than XML. However, even if you don't know how to build your own schemas, you can still use other people's schemas with Excel. In fact, you'll find that doing so actually simplifies your life. In order to use a schema, you simply need to have a copy of it in a file (schemas themselves are complex and ugly and beyond the scope of what a typical office needs—or wants—to learn). Usually, schema files have the extension .xsd.

Tip: For a more comprehensive beginner's introduction to XML and XML schemas, check out the excellent online tutorial provided by W3 Schools at *www.w3schools.com/xml*.

Excel and XML

XML is a great way to exchange data between different computer programs. But what does that have to do with Excel, which already has its own perfectly good file format? Here's the deal: More and more companies today use XML to pass data back and forth. For example, when companies exchange business orders, news organizations post stories, or real estate firms list properties for sale, chances are they're using an XML-based format. If you want to crack open these documents and analyze this data using all of Excel's features, including formulas and charts, then you'll need to use Excel's XML tools.

There's another side to this story. Instead of trying to get XML information *into* Excel, you may need a way to get your worksheet data *out of* Excel. You might want to take an expense worksheet, export it to XML, and then feed that XML into an automated expense-processing program. That program could then track your expenses, submit them to your supervisor for authorization, and notify the payroll department when a payment is required. In a small company, it could be just as easy to print out the expense report and deliver it by hand (or email it). But in a large company, an automated application can help the whole process flow seamlessly, without forcing anyone to sort through stacks of paper or dozens of email messages. In these situations, XML really shines.

Note: Experts estimate that Excel spreadsheets contain more data than all the world's relational databases combined. Excel's XML features can help you extract information that's trapped in your spreadsheet files and use it in other automated applications.

Because XML is so flexible, there's no single-step solution to importing and exporting XML. You can't just perform an *Open XML* command because Excel doesn't know *which* XML format you're using. Instead, you first need to give Excel some information about the specific format you're using, and tell Excel how to extract the data you need.

Excel makes this possible through a set of features called *XML mapping*. XML mapping lets you link a specific XML format to a specific spreadsheet. Once you set up this link, you can use it in two ways: to export data *from* your worksheet into an XML document, or to import the contents of an XML document *into* your worksheet.

Some of the options for Excel's XML features are tucked away on a special tab that doesn't ordinarily appear. Before you do anything with XML and Excel, you need to display this tab. To do so, choose File→Options, and then choose the Customize Ribbon section. On the right of the window, a large box lists all the tabs that are currently shown in the ribbon. Near the bottom, you'll see an unchecked item named Developer. To show the Developer tab, check this box, and then click OK.

Mapping a Simple Document

The simplest way to map an XML document is to link each element in the XML document to a corresponding cell in a worksheet. Then, when you import the document, the data flows out of the elements and into the linked cells.

Note: You can find all the XML files used in this chapter on the Missing CD page at *www.missingmanuals.com/cds*. You can use these files to map your own worksheets.

To try this out, you can use the simple Student.xml document shown here, which stores the test and assignment scores for a single student:

```
<?xml version="1.0" ?>
<Student>
    <Name>Lisa Chen</Name>
    <StudentID>45349920</StudentID>
    <Test1_Score>75</Test1_Score>
    <Test2_Score>63.23</Test2_Score>
    <Assignment1_Score>94</Assignment1_Score>
    <Assignment2_Score>90</Assignment2_Score>
</Student>
```

Keep in mind that in real life, you wouldn't create this document by hand. Instead, it might be extracted from a database or, even more likely, generated by some sort of automated student grading program.

Before you can map an XML document to an Excel workbook, you need to prepare the workbook. Simply follow these steps:

1. **Create a new blank Excel document.**

   You could perform mapping with an existing workbook, but, in this case, it's easier to start from scratch.

Note: Excel worksheets can contain a mix of linked cells (those that'll be receiving the XML file's content) and nonlinked cells containing other information (like your own descriptive labels or formulas that use the imported data). Excel doesn't include these nonlinked cells in any import or export operations.

2. **Choose File→Open.**

   The Open File dialog box appears.

3. **Browse to the Student.xml file, and then open it.**

   Excel shows an Open XML dialog box with three options.

4. **Select "Use the XML Source task pane," and then click OK.**

   You have three choices when it comes to opening a basic XML document. You can import the data into a basic table, dump all its content into the worksheet in read-only mode, or take full control of the linking process. The last option is the best for this example, because it lets you choose where you want to display the data.

5. **Excel warns you that you're mapping a document without a schema. Click OK to continue.**

   A schema defines the structure of an XML document, as explained on page 675. Using a schema is the best way to link XML documents to worksheets because its strict rules prevent errors. You can, however, get away without using a schema, and, because you don't have a schema file for the Student.xml file, you can proceed without it.

   When you click OK, Excel doesn't yet import the data into your document. Instead, the XML Source pane appears, showing a "tree" that includes all the elements Excel found in your XML document (see Figure 23-13). At this point, you're ready to start the mapping process.

To map the Student.xml elements to your spreadsheet, follow these steps:

1. **In the XML Source pane, select the element you want to link.**

2. **Drag the element from the XML Source pane to the appropriate place on your spreadsheet.**

   You may want to place the student's name in cell B1 (leaving room for a label in cell A1). Simply click Name in the XML Source pane, drag it over cell A1, and then release it, as shown in Figure 23-14. Excel outlines the cell in blue to indicate that it's linked.

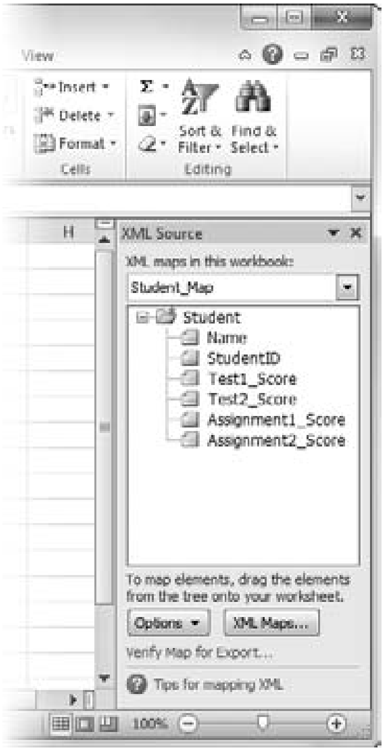

Figure 23-13:
Nothing's happened yet, but the XML Source pane shows a tree diagram with the structure of the Student.xml file. It's now up to you to start the linking process.

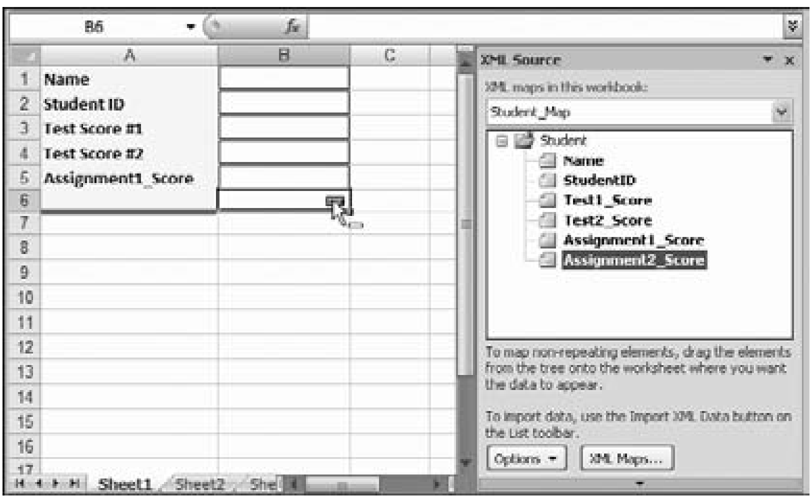

Figure 23-14:
To map the <Assignment2_Score> element to cell B6, begin by dragging the element to cell B5, creating the link.

3. **Optionally, add a caption to your element by clicking the smart tag icon that appears next to the newly mapped cell. Choose "Place XML Heading to the Left" or "Place XML Heading Above," as shown in Figure 23-15.**

   When you choose one of these options, Excel inserts a piece of static text with the name of the linked element in bold formatting. You can use this technique to insert the caption "Name" in cell A1 after you link cell B1 to the <Name> element.

Note: Excel doesn't immediately insert any information in a linked cell. You won't see the student's name appear when you drag the <Name> element to cell B1. Instead, you'll need to import the XML data once you've finished the mapping process. You'll learn how to do that in a moment.

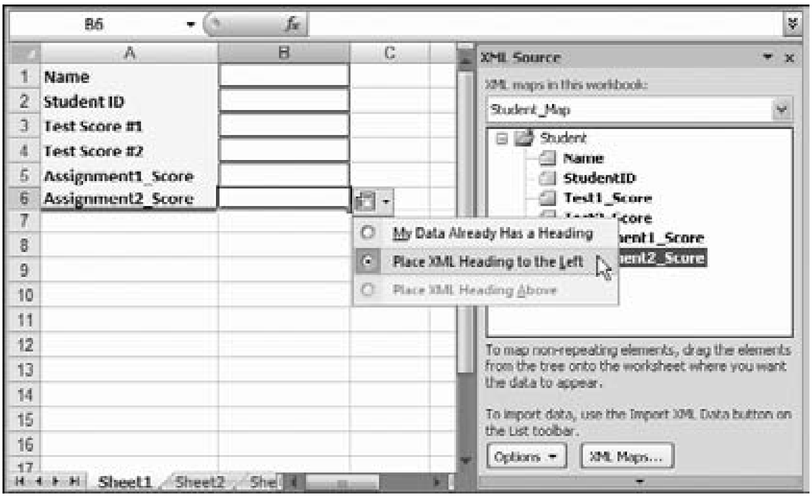

Figure 23-15: *Once you link an element, a smart tag appears. You can use this tag to quickly insert a caption for the linked data. This step is optional, because you can add your own labels. However, using the smart tag helps you quickly create the basic structure of your mapped worksheet. In this case, you still need to do some cleanup to make the XML element names more attractive; you might want to replace Assignment1_Score with Assignment Score #1.*

4. **Return to step 1, and then repeat the process until you've mapped all the elements.**

   You don't need to map every element. It's quite possible that you want to work with just a portion of your XML file's data. In that case, you'd want to map only the elements you need. Of course, if you don't map an element, you can't extract its data and put it into your worksheet. Similarly, if you eventually export the Excel file as an XML document, that element won't be included.

   If you change your mind while mapping, you can remove any element by right-clicking the corresponding element in the XML Source pane, and then choosing Remove.

Once you've finished mapping the document and adding your headings, it's time to import the XML. Choose Date→Connections→Refresh. Figure 23-16 shows the worksheet with the XML data imported.

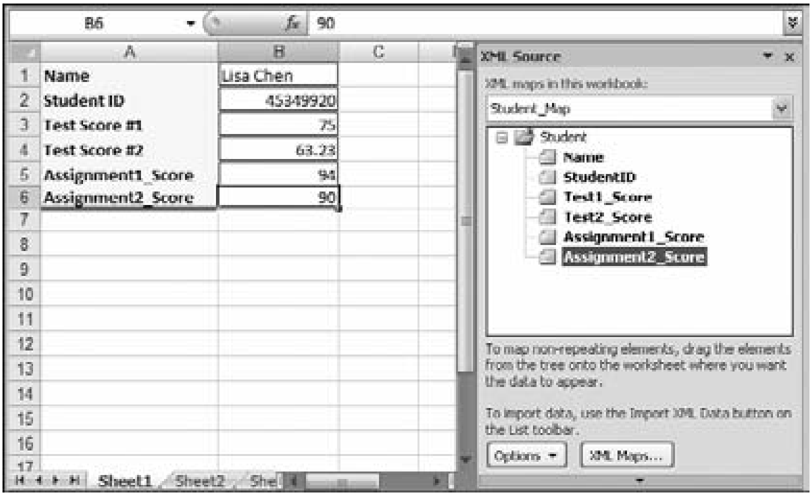

Figure 23-16: *Once you've mapped the XML document, actually importing your data's a snap. In this example, the data's been imported into the worksheet. The cell coloring and borders have been added manually. XML doesn't support any formatting features–it's just a container for raw data.*

Importing and Exporting XML

Now that you've mapped the Student.xml file, you've not only extracted some information from the XML document, you've also given yourself a range of options for managing the link between your spreadsheet data and the world of XML. The real magic of XML lies here.

Once you've defined a link between your worksheet and a specific XML format, you can perform three tasks:

- **Refresh**. If the XML source changes, you can easily refresh your worksheet with the new content. Just click any linked cell, and then select Data→Connections→Refresh or Developer→XML→Refresh Data. (Both buttons have the same effect.)
- **Export**. You can export the data to a new XML document. Click any linked cell, select Developer→XML→Export, and then choose a new XML file. Optionally, you can choose to replace the original XML source with the new XML file. (You might want to use this option if you've edited the content inside your Excel spreadsheet.)
- **Import**. You can import XML data from another file that has the same structure. In this case, click any linked cell, select Developer→XML→Import, and then choose a new XML file. The new XML data flows into the existing worksheet, replacing the original data.

This ability to import means that in the Student.xml example, you could import a student's information, modify it, and then save it back to a *different* XML document. In fact, you could use your spreadsheet as a way to create dozens of different student documents in separate XML files, just by exporting different data to different files. A crafty developer could create a custom application that scans student XML files, and then automatically generates and mails report cards.

All you'd need to do is export the XML. In fact, the whole process could get even easier if someone develops a little piece of Excel macro code to handle the XML export. In that case, you'd just click a button on the worksheet, and Excel would export the student data to XML and submit it to the custom application. Developing this type of workflow takes a fair bit of work, and the first step's learning Excel's VBA macro language, which is introduced in Chapter 29.

Note: Remember, if you save a mapped workbook, you're saving only the Excel spreadsheet file, with whatever data it currently contains. If you want to save the content to an XML file, you need to Export the XML data. Choose Developer→XML→Export.

POWER USERS' CLINIC

XML Mapping with a Schema

Serious XML gurus don't map a worksheet using an XML document. There are too many possible problems, including optional elements that Excel might ignore, data type rules that it doesn't enforce, and certain types of structures that Excel might misunderstand.

A better way to map a worksheet to XML is to use an XML schema. A schema defines the structure of the XML document you want to import. If you're creating a worksheet to analyze lists of student grades, you could use a file called StudentList.xsd that defines the elements you'll use and the document structure.

XML mapping works more or less the same with a schema as it does with an actual document. The key difference is that the schema removes the possibility for error. To map a schema, follow these steps:

1. Open the workbook in which you'd like to add the XML schema, or create a new workbook.
2. Select Developer→XML→XML Source to show the XML Source pane. In the XML Source pane, click the XML Maps button.
3. In the XML Maps dialog box, click the Add button.
4. Browse to your XML schema file, and then click OK.
5. Click OK to close the XML Maps dialog box. You now see the schema-defined elements in the XML Source pane.
6. Drag the elements from the XML Source pane to the worksheet.
7. When you're finished, you need to import an XML document that has the same structure as the schema. To do so, select Developer→XML→Import, and then browse to the XML file.

When using an XML schema, Excel can perform validation with the schema data types to prevent invalid input. If you want to use this feature, you have to turn it on for the document. Select Data→XML→Map Properties. In the list of options, turn on the first checkbox, "Validate data against schema for import and export." Now, every time you import or export XML data, Excel checks it against the rules defined in the schema.

Mapping Lists

Excel's XML features really get interesting when you need to map XML documents that contain *lists* of information, like product catalogs, order tables, and—as in the following example—a class report. As you've no doubt noticed, most Excel documents use lists of some kind. Whether you're tracking student grades, monthly expenses, or employee contact information, you rarely have just one piece of information.

XML documents are often designed to hold repeating elements. The Student.xml file, in contrast, only held the information for a single student. But you can readily create a document that holds a list of students, each one in separate <Student> element containers.

Here's an example (available in the StudentList.xml file, which you can find on the Missing CD page). Only two students are shown here, but the actual StudentList.xml file contains many more students:

```
<?xml version="1.0" ?>
<Students>
    <Student>
        <Name>Lisa Chen</Name>
        <StudentID>45349920</StudentID>
        <Test1_Score>75</Test1_Score>
        <Test2_Score>63.23</Test2_Score>
        <Assignment1_Score>94</Assignment1_Score>
        <Assignment2_Score>90</Assignment2_Score>
    </Student>
    <Student>
        <Name>Edwin Albott</Name>
        <StudentID>45349921</StudentID>
        <Test1_Score>85</Test1_Score>
        <Test2_Score>73.23</Test2_Score>
        <Assignment1_Score>94</Assignment1_Score>
        <Assignment2_Score>95.6</Assignment2_Score>
    </Student>
    ...
</Students>
```

When you try to map this document, Excel quickly notices that the <Student> element repeats. Instead of mapping the <Student> element to a single cell (as it did in the previous section), Excel creates a mapped table that you can use to manage the list of students with filtering, sorting, and searching.

Note: This example, in fact, creates the same type of table you studied in Chapter 14. So why bother? It gives you all the fancy table tools, like sorting, filtering, alternating row formatting, and so on.

To map the StudentList.xml file, follow these steps:

1. **Choose File→Open.**

   The Open File dialog box appears.

2. **Browse to the StudentList.xml file, and then open it.**

   If you haven't already gotten this file, download it from the Missing CD page at *www.missingmanuals.com/cds*. Excel shows an Open XML dialog box with three options for the file.

3. **Select "Use the XML Source task pane," and then click OK.**

   Excel warns you that you're mapping a document without a schema.

4. **Click OK to continue.**

   The XML Source pane appears, with the structure of the StudentList.xml file. It looks similar to the previous example, but there's a difference now. When you drag an element onto the worksheet, Excel creates a table column complete with a header that has a drop-down list. Why? Because Excel recognizes that the StudentList.xml file contains multiple students. And Excel can't store all these students in a single cell on your worksheet.

5. **Click the Student element, and then drag it to cell A1.**

   When multiple columns belong to the same XML list, it's often easiest to drag these elements onto your spreadsheet in one operation. You can select multiple elements by holding down the Ctrl key while you select items in the XML Source pane or by selecting the *parent* element (the element that contains all the elements you want to insert).

   In the student list example, if you select the Student entry in the XML Source pane, you also select *all* the elements that contain student information. You can then drag them all at once. It's entirely up to you whether you create your table as a series of contiguous columns (the easiest approach) or as separate columns spread out over your spreadsheet. Either way, the data is equivalent.

6. **Choose Data→Connections→Refresh All.**

   Now the student information flows into the table, filling it up automatically, as shown in Figure 23-17.

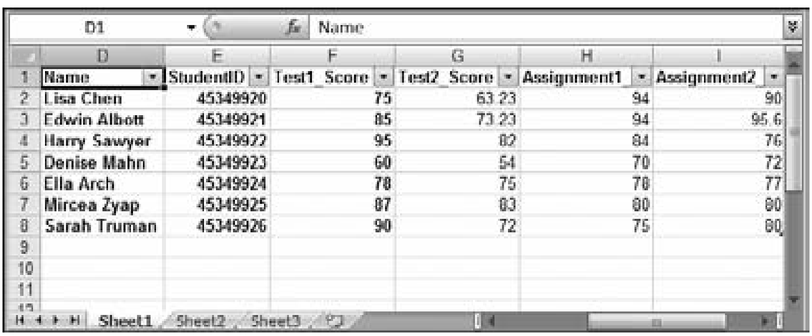

Figure 23-17: *This example shows a mapped worksheet with the table of student information. You can insert new students, remove existing students, and edit student data before you export the information back to XML.*

Tip: In this example, all the XML data is contained in a repeating list of students. However, XML documents often use hybrid structures where they include some repeatable information (like the list of students), and some information that appears only once (like the name of the class and the instructor who's teaching it). In this case, you'd probably link individual cells in the top portion of your worksheet, and then add the table a little lower down.

Gaining the Benefits of XML Mapping

Regardless of whether you want to import or export XML, you should always save a copy of your mapped spreadsheet file. Excel stores the mapping information in that file. Essentially, you should think of this spreadsheet as a window that lets you analyze any XML file, as long as it has the same structure as the XML document you mapped.

When a semester of classes finishes up, an automated student-grading application might generate a new XML document with the most up-to-date information. Fortunately, you don't need to map this document—as long as it matches the structure of the grading document you used initially, you can import the new XML document with a couple of mouse clicks. The new information then flows seamlessly into your existing workbook.

In a very real sense, you can reuse a mapped workbook to examine different XML files in the same way you can use a database query to get and analyze the most up-to-date information from a table in a database (as shown earlier).

For example, consider the StudentList.xml workbook that you created earlier. To make it a better tool for analyzing student grades, you'd probably add a few extra ingredients. Here are some possibilities:

- A calculated column in the table that determines each student's overall grade.
- A calculated field outside the table that determines the average or median grade.
- A chart that shows the distribution of grades in the class.

The beauty of XML mapping is that once you've added these extra touches, you can *reuse* them with the data in other XML documents, provided these documents have the same structure as the document you used to map your worksheet. For example, you might receive a new file, called StudentList_Geography2010.xml, with a whole series of grades for another class. This document uses the same elements, so you don't need to go through the whole mapping process again.

Instead, you can just import this new information into your existing worksheet. All you need to do then is move to a mapped cell, select Developer→XML→Import, and then choose the StudentList_Geography2010.xml file. The student information flows into the linked cells on your worksheet, and Excel immediately updates the other information on your worksheet, including the total grade calculation, average grade calculation, and the chart that shows grade distribution (see Figure 23-18).

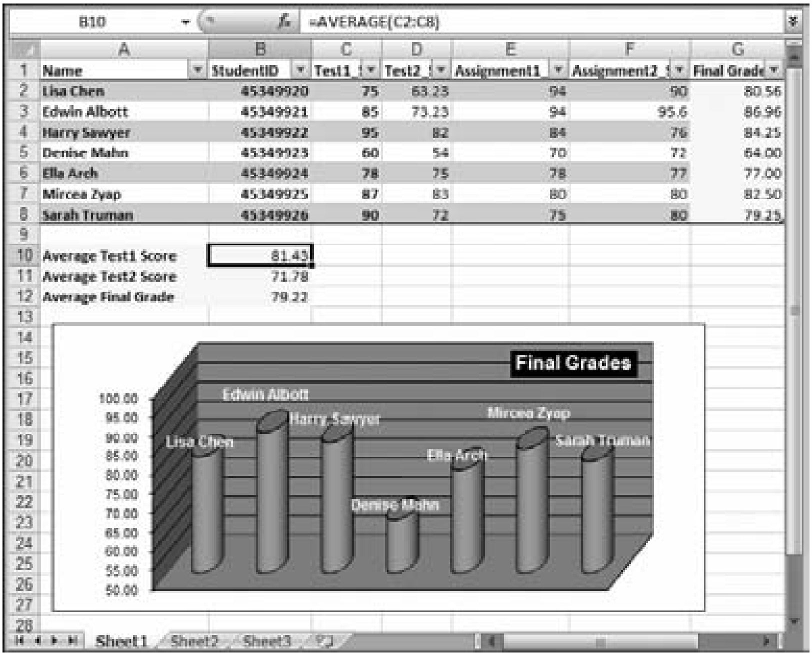

Figure 23-18: *This worksheet enhances the linked table with a new column that calculates a student's final grade, performs a few summary calculations, and adds an attractive chart. Best of all, you can import another XML document with the exact same structure into your worksheet, potentially saving hours of work.*

Web Queries

The Internet's greatest attraction (other than its ability to help Uncle Charlie auction off his *The Y2K Bug Will End Life As We Know It* T-shirts) is that it provides a huge mass of continuously updated information. If you know where to look, you can find up-to-the-minute product prices, stock quotes, and sports scores.

And if you're like most web surfers, you can see this information only by using a web browser. That approach works well most of the time, but it limits what you can actually *do* with the information you find. Because one of Excel's strengths is helping you analyze data—whether it's in your worksheet, a relational database, or an XML document—you'll be glad to know that Excel can also help extract important information right off a live web page. This feature is called a *web query*.

When you perform a web query in Excel, you suck the data out of one or more tables on a web page and insert it into the cells in a worksheet, where you can work on it just like any other data. Excel also stores the web address you used and some information about the table you queried. That way, when you refresh a web query, Excel returns to the web page, grabs the new information from the same location, and replaces it in your worksheet. That definitely beats copying and pasting!

The Limitations of Web Queries

Excel Web queries seem downright miraculous. After all, who could pass up the ability to insert and update information from across the globe? But beware: Web queries are notoriously fragile. Due to the always changing nature of the Internet, a web query that works one day may not work properly the next. Here are some of the pitfalls:

- **Web queries are tied to specific web addresses**. If the website owner moves the information somewhere else, your web query can't find it.
- **Web queries find information based on the structure of a web page**. Minor changes in that web page (like adding a new table or even rearranging the order of elements) can throw Excel off, so it may download the wrong information or fail altogether.
- **Web queries rarely work on websites that require authentication**. That's because you need to log in before you can retrieve the data. Excel can't tell that you haven't logged in, and its attempt to get the information fails.

Note: Most of these shortcomings are factors that limit the *repeat use* of Web queries. In other words, web queries usually work great the first time you use them. Just don't assume that you'll be able to easily update the information in the future, unless you're using a website that never changes.

So, what types of web pages are good candidates for Web queries? Ideally, you'll use a page that has a simple, rigid structure, and puts the information you need in a table. Unfortunately, you often can't tell how complex a page is (or whether it uses tables) just by looking at it, so you may want to try a web query to assess whether it'll work. Also look for web pages where the structure and formatting never change, unless you're willing to create a new web query each time the web page you're drawing from changes. Some examples you may want to try include getting a temperature from a city-specific weather web page, reading a stock quote from a financial site, or getting the list price of a book on Amazon.com.

Creating a Web Query

Now that you're forewarned, you're ready to create a web query of your own. Just follow these steps:

1. **Decide where you want to go to get the data you want to import into Excel.**

   Although you can actually browse different websites from within Excel (more on that in step 3), it's usually easiest to find the correct address using your web browser *before* you create a web query. That's because Excel's New Web Query window doesn't let you access your favorite Internet shortcuts.

Excel finds a page based on its web address. So once you've found the website you want using your web browser, just copy all the text in the address bar. You'll be able to paste it right into Excel's New Web Query window.

2. **Open the worksheet you want to use, or create a new worksheet.**

   You can place as many Web queries as you want in a single worksheet, and you can freely mingle Web queries with ordinary data.

3. **Choose Data→Get External Data→From Web.**

   The New Web Query window appears, as shown in Figure 23-19. This window provides many of the conveniences of Internet Explorer, including an address bar and basic buttons for refreshing, stopping, and going forward and backward. The difference is that Excel automatically scans the page when you surf to it in the New Web Query window, looking for data that it can import. Every time it finds a table of information, it adds a tiny yellow arrow box next to the table. You click these arrow boxes to tell Excel what data you want to extract.

Tip: The New Web Query window is resizable, so stretch away if all the information on your web page isn't visible.

Figure 23-19: *Here's the New Web Query window with a page from Yahoo! Finance. Each arrow-in-a-box represents a table of info that Excel's capable of importing. To give Excel the go-ahead, click any arrow to transform it into a checkmark, like the one shown here next to the table of data about Microsoft's stock performance.*

1. **Enter the web address for the page you want to use in the Address box. Hit Enter or click Go to load the page.**

   Optionally, you can enter the address for a search engine like Google (*www.google.com*), perform a search, and then browse to the page you want. But the quickest and most reliable approach, as discussed previously, is to find the page in a standalone browser first (like Internet Explorer), and then copy that page's web address to the New Web Query window.

2. **Select the table you want to extract.**

   When the page appears in the New Web Query window, Excel adds yellow arrow boxes next to every table you can import. As you hover over each arrow box with the mouse, Excel draws a bold blue outline around the related table.

   Once you find the table you want to extract, click the arrow box (which then changes into a green checkmark). To deselect a table, just click it again.

3. **When you've finished selecting all the tables you want, click the Import button at the bottom of the New Web Query window.**

   The Import Data window appears.

4. **Choose where you want your data to go, and then click OK.**

   If you want to insert the extracted information into your current worksheet, select the "Existing worksheet" option, and then enter a cell reference. If you use the cell reference $A$1, Excel starts the first column of the first row in cell A1.

   If you want to insert the information into a new worksheet, select the "New worksheet" option.

   Once you click OK, Excel begins to fetch the information it needs. During this time, you'll see a "Running background query" message in the Excel status bar. If the query is taking a long time, you can click this message to cancel it. Otherwise, Excel retrieves the information and inserts it into your worksheet, as shown in Figure 23-20.

When you save a worksheet that uses a web query, Excel stores all the data that's now in the cells of your worksheet, along with some information about the web query's data source. You can grab the latest version of your information at any time by moving to any one of the linked cells and choosing Data→Connections→Refresh All.

Note: When you open a workbook that contains a web query, Excel turns off the Refresh All command and shows the message bar with a security warning, just to be safe. (This is the same watchdog behavior you learned about on page 669 with databases.) To turn on your web query, just click the Enable Content button to designate your workbook as a trusted document (page 804).

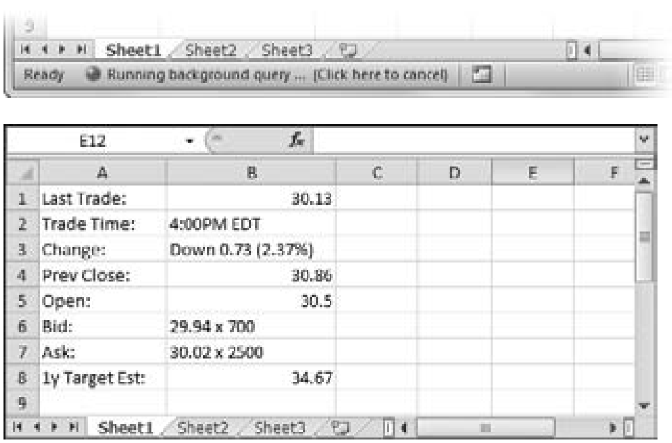

Figure 23-20: *These figures show a web query in action.*

Top: Here's what happens immediately after you start the web query. At this point, Excel's in the process of fetching the information it needs from the Web page.

Bottom: A few seconds later, the data appears. You can see that Excel successfully converted a formatted table of HTML into a plain text-and-numbers table of data that you can analyze to your heart's content.

Excel gives you some ability to respond to changes made on web pages from which you're drawing data. You can edit the data source at any time by following these steps:

1. **Choose Data→Connections→Connections.**

   The Workbook Connections window appears. It lists all the connections that link to outside data, including databases, XML files, and web queries.

2. **Find your web query connection in the list (which is probably the only connection you've created so far), and then click Properties.**

   The Connection Properties dialog box appears. Here you can give your connection a better name and description, and then set it to refresh your data periodically (either every time the file is opened or after a certain number of minutes). Page 686 has more information about this feature.

3. **Click the Definitions tab, and then click the Edit Query button.**

   This action opens an Edit Web Query window, which looks exactly the same as the New Web Query window. Using this window, you can choose a different table, or even enter a different web page URL.

4. **When you're finished, click Import. Click OK to close the Connection Properties dialog box, and then click OK again to close the Workbook connections window and return to your worksheet.**

The Research Pane–A Web Query Alternative

The Research feature is one alternative to a web query. While the web query tool lets you get information from any web page (and is very fragile), the Research pane gets information from a small set of web services that are designed explicitly for the purpose of sending real-time content to Excel over the Internet. To use the Research feature, choose Review→Proofing→Research. The Research pane appears on the right of the Excel window (see Figure 23-21).

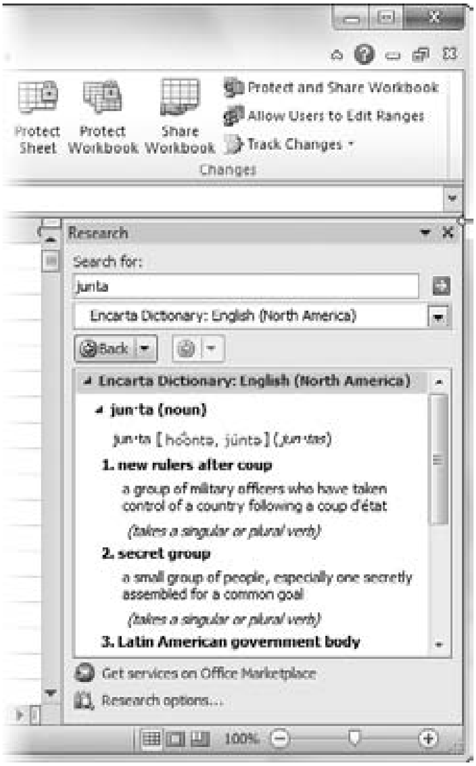

Figure 23-21: *Here's Excel's Research pane. Enter your search keyword in the first text box, choose the Research service you want from the pull-down menu below it, and then click the green arrow (circled) to start the search. When the results appear, you can read them in the results box. Depending on the service, you may also be able to copy some of the information into an Excel worksheet by clicking a button. In this example, the Encarta dictionary provides the definition for a confounding word.*

When you first install Excel, it's preconfigured with a set of Research services. These services include a dictionary that lets you search for detailed definitions, a thesaurus that shows you synonyms, and the Encarta encyclopedia, which provides detailed information on a range of topics. You can also get stock market quotes from the MSN Money website and perform broader web searches using Microsoft's Bing se-rach engine.

Overall, the Research task is a nifty way to get some free information (regardless of whether you want to put it in a worksheet). Microsoft hopes that third-party companies will develop services that plug into the Research framework—but currently, only a few exist (and they're all fee-based). Click the "Get services on Office Marketplace" link to browse a catalog with more useful services online.

CHAPTER 24

Protecting Your Workbooks

So far, you've created spreadsheets that are a bit of a free-for-all. You (or anyone else) can open them and change absolutely anything, from the most minor formatting detail to the most critical formula. For everyday Excel use, this freedom makes perfect sense. However, if you're thinking about sharing your carefully crafted work with other people—like colleagues who need to review your numbers—some caution is in order.

In this chapter, you'll consider two tools that can help you build bulletproof spreadsheets: *data validation* and *worksheet protection*. Data validation catches incorrect values. Worksheet protection locks down your worksheets so they accept only certain types of changes in certain areas. Using these features, you can make your workbooks impervious to error (and deliberate fudging).

Understanding Excel's Safeguards

Excel's data validation and worksheet protection give you a number of ways to keep your workbook on the right side of the law. Using them, you can:

- Prevent people from changing a worksheet's structure (inserting or deleting cells, columns, or rows).
- Prevent people from changing a worksheet's formatting (including the number format or other formatting details like column width and cell color).
- Prevent people from editing certain cells.
- Prevent people from entering data in a cell unless it meets certain criteria.
- Provide additional information about a cell in a pop-up tip box.

- Prevent people from editing—or even seeing—the spreadsheet's formulas.
- Prevent people from moving to cells they don't need to edit or inspect.

You may have different reasons to apply these restrictions. Often, you'll use them to make sure people don't tamper with data. For example, perhaps you have a workbook with a carefully compiled list of sales totals, expenses, and profit calculations. You want to let others update the expense information, but they shouldn't be able to fudge the sales records. Or maybe you want to lock down *all* the data, and let others play only with the summary tables and charts. In both these situations, worksheet protection can prevent unauthorized changes, so that the data in your spreadsheet is just as reliable after it passes a round of revisions as it was when it first left your hands.

In other cases, you might use data validation and worksheet protection to prevent errors, particularly if you're sharing your work with a less experienced Excel patron. For example, imagine you need to give a copy of your timesheet workbook to all the employees on your team. At the end of every month, they fill out their own copies and pass the finished workbooks on to a manager. Unfortunately, an ordinary workbook is a small minefield for someone who's new to Excel. An Excel novice can accidentally delete or overwrite a formula just by pressing the wrong key, and it's almost as easy to put the wrong information in a cell (for example, by entering a date incorrectly so that it's interpreted as text). With data validation, you can lock out certain types of errors and guide the people using your workbook to make sure they fill in the right information.

Tip: Worksheet protection and data validation also make great additions to Excel templates (Chapter 16). Using these features, you can make sure the people who use your templates put the right information in the right places.

Data validation and worksheet protection are two of the most powerful yet underused Excel features. Once you master them, you'll be able to safeguard your spreadsheets before you share them.

Data Validation

With data validation, you can easily prevent people from entering the wrong data in a cell (or at least warn them when they do). Data validation also helps make Excel a little kinder and gentler for a novice by letting you create custom error messages, and add helpful pop-up tips. You need a little time to set up data validation, so usually you'll use it only on your worksheet's most important cells (Figure 24-1).

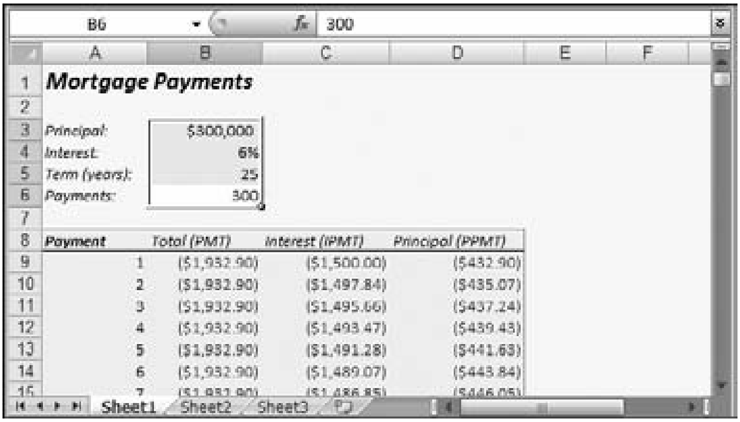

Figure 24-1: *In this worksheet, which calculates mortgage payments, it makes sense to use data validation on the cells you expect people to change–like the loan amount (cell B3) and the interest rate (cell B4). You can then use worksheet protection, as discussed later in this chapter, to prevent another person using the spreadsheet from modifying other cells altogether.*

To apply data validation, move to the appropriate cell, and then choose Data→Data Tools→Data Validation. A Data Validation dialog box appears with three tabs: Settings, Input Message, and Error Alert. You fill in the settings in these three tabs, and then click OK to put the rule into action. The following sections explain each tab of the Data Validation dialog box.

Tip: You can apply validation to a number of cells at once. Just select all the cells before you choose Data→Data Tools→Data Validation.

Settings

Use the Settings tab of the Data Validation dialog box (shown in Figure 24-2) to specify what values you want to allow in the cell. You have two methods at your disposal. First of all, in the Allow list box, you can set the *type* of value that's allowed. The simplest possible types include whole number, decimal, date, time, and text. (Two other types of values—custom and lists—are covered later in the chapter.) For example, if you select "Whole number" from the Allow box, and then try to input a value of 4.3 into the cell, Excel shows an error message and prevents your input.

Once you've chosen the data type, you still need to set the data *range*. Do so by choosing a comparison from the Data list box, and then specifying the values that you want to use for your comparison. For example, if you want to restrict input to a whole number from 0 to 5, choose "between" and set a minimum value of 0 and a maximum value of 5. Other comparisons you can use include less than, greater than, greater than or equal to, and so on. You set the data range for all data types, including dates, times, and text (in which case you set limits on how many characters can go in the text).

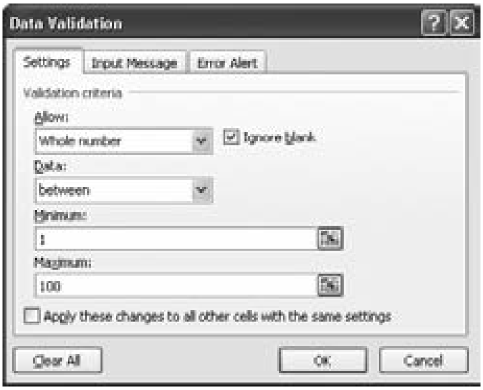

Figure 24-2:
The options selected here force the person using the workbook to enter a whole number from 1 to 100 (or leave the cell blank). In addition, the "Ignore blank" checkbox in the top-right corner is turned on. This setting tells Excel to allow empty values, so that it doesn't try to validate the cell if it doesn't contain any data.

It's easy to remove your validation rule later on. Just select the appropriate cells, choose Data→Data Tools→Data Validation, and then click Clear All.

Note: Data validation begins policing a cell only *after* it's been applied. If you apply a validation rule and the cell already contains invalid information, Excel doesn't complain.

TIMESAVING TIP

Changing the Validation in Many Cells at Once

If you have several cells with the same validation settings, you can change them all at once. For example, imagine you have several cells in an expense report that allow prices from $0 to $100. In a fit of generosity, you decide to raise the upper limit of the range and allow prices from $0 to $200. To do this, just select all the price cells, choose Data→Data Tools→Data Validation, and make your changes. (This trick doesn't work if the cells have different validation settings. In this case, Excel warns you that it will need to erase the current validation settings in the selected cells if you continue to insist on modifying them all at once.)

This feature is handy, but Excel has an even niftier way to make sweeping validation changes. If you have several cells with the same validation settings, and you want to change every single one of them, you can get Excel to select the cells automatically. To do this, click one of the cells, choose Data→Data Tools→Data Validation, and then switch on the "Apply these changes to all the other cells with the same settings." If you look at your worksheet now, you'll see that Excel has automatically found and selected every cell with the same validation settings. All you need to do is make your changes, and then click OK to perform a mass update.

Input Message

Once you've added data type rules, it's a nice touch to give the person using your workbook (or template) fair warning. You can do so by adding a pop-up message that appears as soon as somebody moves to the corresponding cell.

Tip: Use the Settings tab to prevent bad values. Use the Input Message tab to add some helpful information. Sometimes, you'll want to use just one of these two tabs, and other times you'll need them both. The Input Message tab really shines when you're sharing a workbook that someone's likely to copy and reuse in the future, or when you're building a template. And it's particularly handy for giving Excel newbies some guidance.

An input message has two components: a title (displayed in bold) and a more detailed message. In the Input Message tab of the Data Validation dialog box (shown in Figure 24-3), you can enter both of these pieces of information—just fill in the Title and "Input message" boxes. (Also, while you're there, make sure the "Show Input message when cell is selected" checkbox is turned on. Otherwise, Excel doesn't show your message at all.)

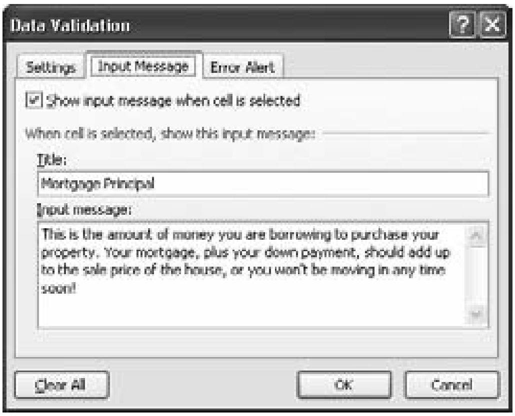

Figure 24-3:
Here's how to create a helpful and descriptive message to tell whoever's using your workbook what the term "mortgage principal" really means. When you define an input message, you choose a title for the pop-up box (the title appears in bold) and a descriptive message. Figure 24-4 shows the message in action.

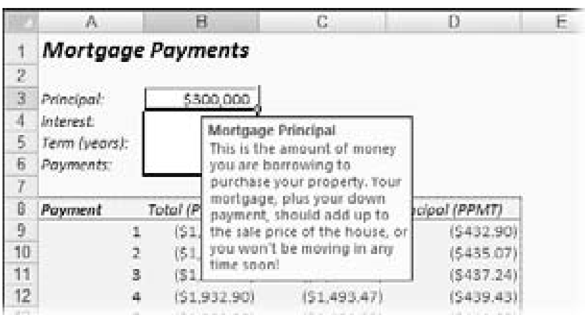

Figure 24-4:
When the person using the workbook moves to a cell that has an input message, a yellow tip box appears displaying the message. Anyone using the spreadsheet can drag the box to another location on the worksheet if it's obscuring some important information.

Input messages should contain more than a description of your data validation settings. Ideally, your input message *explains* a little bit about the data that the cell is looking for. In fact, depending on who'll be using your workbook, you can include information about the type and format of information, what restrictions they need to be aware of, and even how they should enter the information.

For example, for an invoice date cell, you might want a message that says, "This is the date your invoice was submitted to your manager. When you enter a date, use the format day-month-year (as in 29-1-2010 for January 29, 2010), and make sure you don't enter a date earlier than 1-1-2010." For a payment code, you might include a message like "This is the code from the top-right corner of your pay stub. All pay codes start with the letters AZ."

Error Alert

Despite your best attempts, someone, somewhere probably will still type the wrong information in a cell, defying your input messages. In this case, you need to respond by politely explaining the problem. Unfortunately, Excel's standard error message—displayed when someone breaks the data validation rules set down in the Settings tab—leaves a lot to be desired. It's unnecessarily harsh and confusingly vague. Use a better approach, with the Error Alert tab to define your own error message.

To do so, head to the Alert tab in the Data Validation dialog box (Figure 24-5). Begin by turning on the "Show error alert after invalid data is entered" checkbox. This setting tells Excel to monitor the cell for invalid information. If anyone enters the wrong data, a dialog box appears with an error message. The buttons that Excel uses in that dialog box depend on the style of error message you choose.

Figure 24-5:
Helpful error messages, like this one, tell anyone using the spreadsheet what they need to do to fix their mistakes.

To choose an error message style, select an option from the Style list. Excel novices will appreciate it if you use the friendly Information icon, as shown in Figure 24-6, instead of the red X alert icon. But keep in mind that different icons have different effects on whether Excel tolerates invalid input. Here are your choices:

- **Stop**. Excel displays the error message along with a Retry and Cancel button. The person using the workbook must click Cancel to reverse the change (which returns the cell to its last value) or Retry to put the cell back into edit mode and try to fix the problem. The Stop option is the only Style choice that completely prevents the person using the workbook from entering invalid data.
- **Warning**. In this case, the error message includes Yes and No buttons that let the person entering the data decide whether to go ahead with the input. Clicking Yes makes Excel accept the data entered into the cell, even if it breaks the validation rules.
- **Information**. The error message comes with Cancel and OK buttons. Clicking OK enters the new (invalid) data in the cell; Cancel leaves the cell unchanged.

Figure 24-6: *Because this example uses the Information style, the error message box shows a friendlier icon–the "i" inside a circle.*

After you set the icon, specify a title, and then type a descriptive message indicating the error message you want. Remember, you won't know exactly what went wrong, so it's best to reiterate the data type rules you applied, or use a phrase that begins with something like "You probably…" to identify a common problem.

Note: If you set a data validation rule but turn off the "Show error alert after invalid data is entered" checkbox, the person entering data won't see any error messages and won't have any idea if she's entered the wrong type of information. You'd switch off the "Show error alert after invalid data is entered" checkbox only if you ever want to temporarily turn off your validation rules but not remove them, so that you can apply them again later.

GEM IN THE ROUGH

Quickly Spotting Every Error

Auditing circles are an often-overlooked Excel troubleshooting tool. When you choose the Data→Data Tools→Data Validation→Circle Invalid Data command, Excel draws a red circle around every cell that breaks a validation rule on the entire worksheet (Figure 24-7). You can click Data→Data Tools→Data Validation→Clear Validation Circles to remove this information.

You don't need the auditing circles when you're using strict data type validation (in other words, when you're rejecting *all* errors with the Alert message type), because you never end up with invalid data. However, it comes in very handy when you're allowing errors and only showing a Warning or Information message. You might choose to apply these invalid data circles if you want to let one person fill out a workbook, but let another person review it. In this scenario, the auditing circles can help the second person in the chain automatically find the mistakes that need to be corrected.

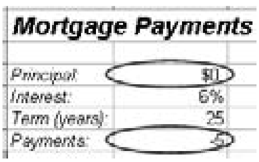

Figure 24-7:
Validation circles, automatically added by Excel when you use the Data→Data Tools→Data Validation→Circle Invalid Data command, help you spot the troublemaking cells.

Data Validation with Formulas and Cell References

You can create more advanced cell restrictions by using formulas and cell references. Imagine you want people to enter the current date on a form they're filling out. You want to make sure that the date is no *earlier* than the current date. In this scenario, you can't use a literal date in your data validation rule, because the date limit needs to change depending on when someone makes the edit.

However, you can use the TODAY() function to help you out. To do so, select Data→Data Tools→Data Validation. In the Settings tab, select the Date data type, and require values to be greater than or equal to *=TODAY()*. In other words, when someone types information into that cell, Excel runs the TODAY() function and compares the result against the cell value. This example shows one way you can embed a function within a data validation rule.

You could also use a formula that contains a cell reference. You might want to make sure that an expense cell always contains a value that's equal to or less than a corresponding budget cell. In this case, you can't put the budget limit directly into the data validation rule, because you don't know what the budget will be until the person using the workbook fills it in. You need to create a data validation rule that uses a formula that references the value of the budget cell. You could specify that the value in the expense cell must be less than or equal to the formula =C3, assuming the budget value's in cell C3.

For an even more powerful approach, from the Allow list box, choose the Custom data type. When you use Custom, you must supply a *conditional* formula in the formula box. A conditional formula is simply one that responds with a value of either *true* or *false*. (Chapter 13 tells you all about how to create conditional functions.) If the result is true, Excel allows the cell entry. If it's false, Excel displays your custom error dialog box (assuming the "Show error alert after invalid data is entered" checkbox in the Error Alert tab is switched on). The neat thing about conditional formulas is that you can combine as many unrelated conditions as you need, using the conditional functions AND() and OR().

Conditional validation rules are also useful if you need to compare the current cell against the value returned by a function. The following fairly intimidating-looking formula prevents the person using the workbook from entering a date that falls on a weekend. B3 is the cell containing the validation rule, and the WEEKDAY() functions verify that B3 doesn't represent a Saturday (a value of 7) or a Sunday (a value of 1). This formula forces B3 to adhere to both these restrictions by using the AND() function:

```
=AND(WEEKDAY(B3)<>1, WEEKDAY(B3)<>7)
```

Another time you might use a conditional formula in a data validation rule is when you need to monitor a group of cells and make sure their total value doesn't exceed a total that you specify. If you don't want the series of expense items in cells B2 to B7 to total more than $5,000, you select all these cells, choose Data→Data Tools→Data Validation, choose the Custom data type, and then supply the following formula:

```
=(SUM($B$2:$B$7)<=5000)
```

Figure 24-8 shows an example that introduces an improvement on this formula. Instead of using an exact budget limit in the data validation rule, another cell supplies the budget limit.

Note: You may notice that the SUM() formula uses absolute cell references (for example, $B$2) rather than ordinary references (like B2). You can apply the exact same SUM() formula to multiple cells at once. If you don't take this step, Excel modifies the cell references in each subsequent cell, which isn't what you want. See page 252 for a refresher on the difference between absolute and ordinary cell references.

You may also use the Custom data type and write a conditional formula to prevent duplicates in a range of cells. For example, the next formula checks to see that there's no other instance of B3 in the range of cells from B2 to B7. This validation rule goes into cell B3:

```
=COUNTIF(B2:B7,B3)<=1
```

Figure 24-8:
Here, a custom data validation rule polices the total value of a group of cells. Figure 24-9 shows the result.

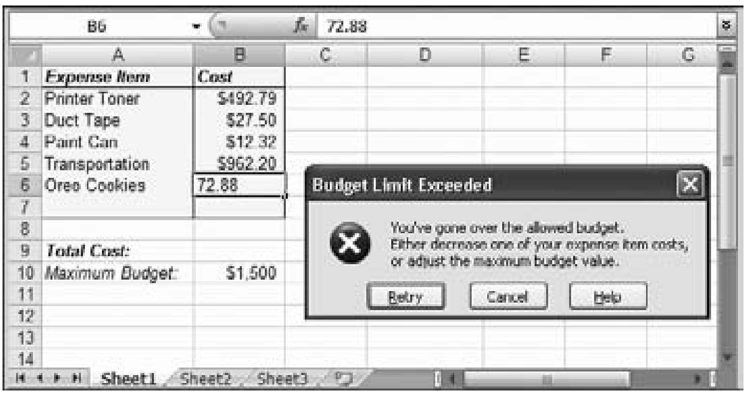

Figure 24-9:
If the value of cells B2 through B6 is greater than the Maximum Budget cell value (B10), the rule rejects the entry. Of course, you can combine this example with the techniques shown later in this chapter to lock the Maximum Budget cell, preventing other people from changing how much money they're allowed to spend.

This formula isn't quite as convenient as the summing formula because the SUM() formula applies to all the cells in a range. The COUNTIF() formula needs to be tweaked for each cell. The formula shown above is what you'd use to validate the contents of cell B3, but in order to perform the same check for duplicates in the other cells in the specified range (B2 to B7), you'd need to modify the formula (replacing the reference B3 with whatever cell you wanted to check).

Data Validation with Lists

The only other data type choice you have (when filling out the Allow field in the Settings tab) is the List option. The List choice is interesting because it doesn't just restrict invalid values, it also lets you add a handy drop-down list box that appears when anyone using your spreadsheet moves into that cell. The person who's entering data can use the list to quickly insert an allowed value, without needing to type it in. You can also type values in by hand, but Excel assumes that if the value you enter doesn't match one of the entries in the list, your entry's invalid (and it shows an error message depending on your Error Alert settings).

To create this list, choose the Data Validation dialog box's Settings tab, and then choose List from the Allow text box. You have two choices for supplying a list in the Source box: You can type in a list of comma-separated values with no spaces in between them (like *1,2,3* or *blue,black,red*), or you can use a cell range that contains a list. If you want the person using the list to be able to choose the entry from a drop-down list of options in the cell (which is a slick and convenient touch), make sure you keep the "In-cell dropdown" checkbox turned on. Figure 24-10 shows an example that modifies the lookup worksheet used in Chapter 12. You can create an invoice by choosing products from a drop-down list.

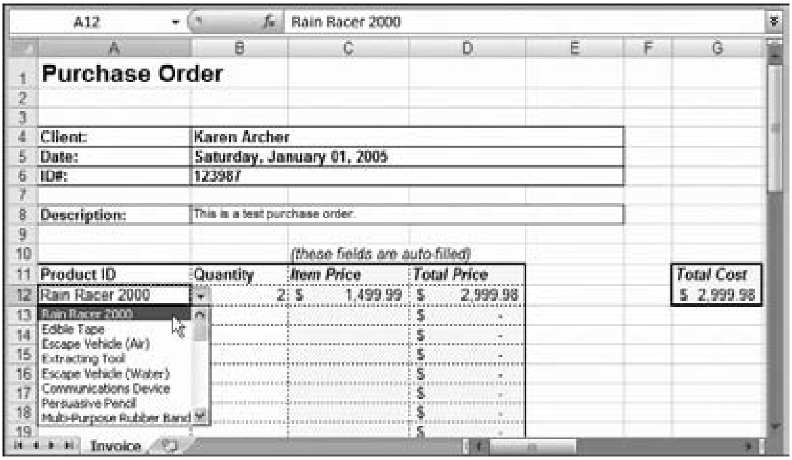

Figure 24-10: *This worksheet uses list validation. This approach's advantage is that people using the spreadsheet don't need to remember the name or ID of each product. Instead, there's always a complete list of possible choices available at their fingertips. The only disadvantage is that the list you use must be on the same worksheet. Here, that means the entire product catalog is off to the side of the current worksheet, so that it's available for the drop-down list and lookup functions. Figure 24-11 shows the data validation settings that make it work.*

Locked and Hidden Cells

Excel's data validation tools help make sure funky data doesn't end up in your worksheets. But they don't protect your worksheets against things like accidentally deleted formulas, mistakenly scrambled formatting, and "unintentionally" modified Maximum Deficit Spending values. To defend against these dangers, you need to use Excel's *worksheet protection* features.

Figure 24-11:
The list validation settings in this dialog box tell Excel to generate a list of product choices drawn from the range of cells indicated in the Source box. See Figure 24-10 for what this looks like on the spreadsheet.

To understand how worksheet protection works, you need to know that each cell can have one of two special settings:

- **Locked**. When a cell is locked, you can't edit it.
- **Hidden**. When a cell is hidden, its contents don't appear in the formula bar. The cell still appears in the worksheet, but if the cell uses a formula, you can't see the formula.

You can use these settings individually or together. When a cell is both locked and hidden, you can't edit it, or view it in the formula bar. On the other hand, if a cell is hidden but not locked, people can edit the cell but can never tell whether the cell uses a formula, because Excel keeps that information secret.

The most important thing you need to understand about locked and hidden cells is that these settings come into effect only when the worksheet is *protected*. If the worksheet isn't protected (and every worksheet begins its life without protection), Excel doesn't use these settings at all. In other words, you need to take two steps to building a bulletproof worksheet. First, you start by specifying which cells are locked and hidden, and then you finish up by protecting the sheet so your settings take effect.

Interestingly, every cell in your worksheet starts off in an unhidden and locked state. If you switch on worksheet protection without making any changes, your whole worksheet becomes read-only. Excel uses this approach for a reason. Typically, you'll use worksheet protection to make sure the person using your workbook can edit only a few select cells. It's much easier to designate the few cells that are editable than it is to try to select every single cell that needs to be locked. (Remember, every Excel worksheet boasts millions of cells, most of which are empty.)

Note: Once you've protected a worksheet, you can't change the protection settings of any cells. Of course, you can unprotect a worksheet, as explained below, and *then* change the cells' protection settings.

Protecting a Worksheet

Here are the steps you need to follow to protect your worksheet:

1. **First, *unlock* all the cells into which you want people to type information.**

   You can perform this procedure one cell at a time, or you can select an entire range of cells. Once you make your selection, right-click it, and then choose Format Cells. The Format Cells dialog box appears.

2. **Click the Protection tab. Then, turn off the Locked checkbox, and then click OK.**

   Next, you need to hide formulas that you don't want the person using the workbook to see.

Tip: You can lock or unlock cells without going to the Format Cells dialog box. Just select the cells, and then choose Home→Cells→Format→Lock.

3. **Select the cell or cells with the formulas you want to hide, right-click the selection, and then choose Format Cells again. This time, click the Protection tab, turn on the Hidden checkbox, and then click OK.**

   If you want, you can change both the Hidden and Locked settings for a cell or group of cells at the same time. Once you finish unlocking and hiding to your heart's content, it's time to protect the sheet.

4. **Select Review→Changes→Protect Sheet. In the Protect Sheet dialog box that appears, turn on the checkbox labeled "Protect worksheet and contents of locked cells."**

   In addition to protecting the contents of unlocked cells, this dialog box lets you toggle on or off a list of Excel actions that you want to let people using your worksheet perform, as described in the next step.

5. **From the "Allow all users of this worksheet to" list, turn on the things that you want people using your worksheet to be able to do.**

   Excel's standard approach is to restrict everything except for cell selection (the first two options). Here's a setting-by-setting breakdown of your choices:

   - **Select locked cells**. Turn off this checkbox if you want to prevent people from moving to locked cells.

Note: Keep in mind that if you can't select a locked cell, there's also no way to copy and paste the information in the cell to another worksheet or program. If you want the people using your workbook to be able to do this, you should keep the "Select locked cells" setting switched on.

- **Select unlocked cells**. Turn off this checkbox if you want to prevent people from moving to unlocked cells. You won't use this setting very often, but you might use it in conjunction with the "Select locked cells" setting to lock someone out of the worksheet entirely.
- **Format cells, Format columns**, **and Format rows**. Turn on these checkboxes if people need to be able to format individual cells or entire columns and rows. If you allow row and column formatting, Excel also permits people to hide rows and columns. However, Excel never lets anyone change the locked and hidden settings of a cell while it's protected.
- **Insert columns** and **Insert rows**. Turn on these checkboxes if you want to let people insert new rows or columns.
- **Insert hyperlinks**. Turn on this checkbox if you want to let people insert hyperlinks in unlocked cells. This setting can be dangerous because a hyperlink can point to anything from another worksheet to a malicious Web page. See page 90 for more information about hyperlinks.
- **Delete columns** and **Delete rows**. Turn on these checkboxes to bestow the ability to remove columns or rows. Use this setting at your peril, because it lets people decimate your worksheet—for example, removing entire ranges of data even if they contain locked cells.
- **Sort**. Turn on this checkbox to let people sort unlocked cells, while keeping locked cells impervious to sorting.
- **Use AutoFilter**. Turn on this checkbox to let people use filtering on any tables in the worksheet. See Chapter 14 for more information about tables.
- **Use PivotTable reports**. Turn on this checkbox to let people manipulate any pivot tables in the worksheet. See Chapter 22 for more information about pivot tables.
- **Edit objects**. Turn on this checkbox to let people edit or delete embedded objects in the worksheet. These objects can include data from other programs or, more commonly, pictures (Chapter 19) or charts (Chapter 17).
- **Edit scenarios**. Turn on this checkbox to let people edit or delete what-if scenarios. See Chapter 21 for more information about scenarios.

6. **If you want to stop other people from unprotecting the worksheet, specify a password in the "Password to unprotect sheet" text box (see Figure 24-12).**

   Once you protect a worksheet, anyone can unprotect it. All a person needs to do is select Review→Changes→Unprotect Sheet. This behavior makes sense if you're just using protection to prevent people from making casual mistakes.

However, if you're worried about deliberate tampering, or if you want to create a truly invulnerable worksheet, it's a good idea to set a password. If you do, no one can unprotect the sheet without supplying the password.

Note: Worksheet protection isn't designed to protect your work from data theft and malicious tampering. In fact, there are well-known exploits that let hackers quickly break worksheet protection, no matter what password you use. Instead, worksheet protection is designed to stop ordinary people from changing something they shouldn't. If you need a higher grade of security, consider encrypting your entire workbook file with a password, as described on page 50.

Figure 24-12:
The Protect Sheet dialog box lets you set a password that prevents anyone else from tampering with your worksheet.

7. **Click OK.**

   The protected worksheet doesn't look any different, but if you try to edit a locked cell, you get an error message explaining that the cell's locked (but explaining that you can unlock the worksheet by selecting Review→Changes→Unprotect Sheet).

Tip: Protected worksheets have a nice feature: You can quickly find all the unlocked cells. Just press the Tab key to move from one unlocked cell to the next. When you reach the last unlocked cell, Excel automatically jumps back to the first unlocked cell. If you want to make it even easier to navigate a dense worksheet, consider turning off the "Select locked cells" setting in the Protect Sheet dialog box.

Protecting the Entire Workbook

You can use one more level of protection: Excel gives you the power to protect an entire workbook. When a workbook's protected, Excel prevents people from inserting, moving, or removing worksheets (tasks that were covered in Chapter 4).

POWER USERS' CLINIC

Exerting Even More Control with IRM

Worksheet protection isn't the only way you can lock down a spreadsheet. Microsoft Office also has another technology for controlling what people can do. It's called Information Rights Management (IRM).

IRM is a powerful feature that can exert Big Brother-like control over your spreadsheets. For one thing, it can prevent people from copying data to the clipboard or printing a workbook. It can also automatically "expire" your workbook, so that it can't be accessed after a date you choose. Most importantly, IRM is user-specific, which means you can give different permissions to different people.

To run IRM on your own, you need to have some seriously high-end software, including a company rights management server and Active Directory. Unless you're at the helm of a fairly large company, you probably don't have this software on hand (or you don't want the monumental task of setting it up to use IRM). As an alternative, Microsoft lets people sign up for its free IRM service, which works over the Internet.

Microsoft's IRM service relies on the Windows Live ID authentication system. That means that as long as you have a Windows Live ID (which you do, if you have a Hotmail email account), you can use your email and password to log on to Microsoft's IRM service. If you don't have a Windows Live ID, you need to create a new account or get a Windows Live ID for an email address you already have. You can accomplish either task at *http://home.live.com.*

Overall, IRM makes most sense as a specialized tool for big business. But if you want to experiment with it, start by choosing File→Info. Click the Protect Workbook button, and then choose Restrict Permission by People→Restricted Access. Follow the prompts to install a small Office update that adds the IRM features, and sign up for Microsoft's free IRM service. You can then control who can use your workbook (by specifying their email address) and what they're allowed to do.

When you open an IRM-protected workbook, you need to supply your email address and password. Excel then communicates with the IRM service to find out what you're allowed to do. If you aren't in the list of authorized people, you're out of luck. And if you try to open an IRM-protected workbook on a computer that doesn't have an add-in, Excel prompts you to download it from the Web.

Workbook protection works hand-in-hand with worksheet protection. If you use workbook protection but not worksheet protection, people can still edit all the cells in your worksheets. However, they can't delete the worksheets or add new ones. On the other hand, if you use workbook protection *and* worksheet protection, people can't tamper with your data or the structure of your workbook.

To turn on workbook protection, select Review→Changes→Protect Workbook. The awkwardly named Protect Structure and Windows dialog box appears, as shown in Figure 24-13.

The Protect Structure and Windows dialog box provides two checkboxes:

- **Structure**. When you turn this option on, the people using your workbook can't insert or delete worksheets. They also can't rename an existing worksheet, hide it, or move it from one place to another.

Figure 24-13:
The Protect Structure and Windows dialog box looks a little like the Protect Worksheet dialog box, but it provides fewer options. You still have the same chance to lock out miscreants by using passwords, however.

- **Windows**. When you turn this option on, Excel doesn't let anyone change the size or position of your workbook window inside the main Excel window. (Usually, this setting has no effect because the workbook window's *maximized*, which means it's given the largest size that fits in the Excel main window.) In addition, you can't use or change other fancy viewing settings, like frozen columns and split windows (both of which are described in Chapter 7).

Note: In rare cases, you might use the Windows protection setting to make sure that your workbook viewers don't change specialized view settings you've created. However, most of the time it's not worth the trouble, because different people using different computers (with different monitors, different display resolutions, and different eyeglass prescriptions) may have good reasons to want to adjust these settings.

You can also supply a password to prevent people from unprotecting your workbook. You can use the same password that you used to protect a worksheet, or you can choose a new password.

Once you're finished, click OK to apply your protection settings. You can remove workbook protection by once again choosing Review→Changes→Protect Workbook, although if you used a password, you'll need to have it handy.

Protecting Cell Ranges (with More Passwords)

As you've already seen, you can use passwords to lock up individual worksheets or the entire workbook. For most Excel workers, this is as far as they want to go. But if you're hungering for more protection settings, you do have another option. You can lock up individual *ranges*. The ranges you protect can be anything from a single cell to a larger group that encompasses entire rows and columns.

Note: The ability to protect individual cell ranges sounds great. However, it's often more trouble than it's worth. After all, do you want to manage a workbook with dozens of different password-protected areas, each with different passwords? If you can, save yourself some aggravation by separating the data you need to protect and placing it in a dedicated worksheet.

Here's how to add protection to a specific range of cells:

1. **Select the cells you want to protect.**
2. **Choose Review→Changes→Allow Users to Edit Ranges.**

   The Allow Users to Edit Ranges dialog box appears.
3. **Click New to create your first range.**

   The New Range dialog box appears (Figure 24-14).

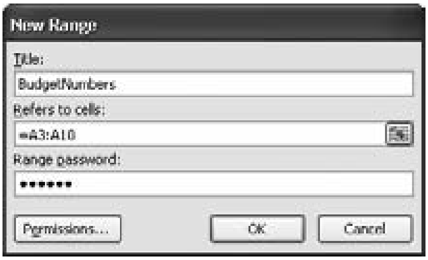

Figure 24-14:
The New Range dialog box lets you supply three key pieces of information for your range–the title, cell references, and password.

4. **Fill in a title and password for your range.**

   You don't need to explicitly set the "Refers to cells" box because Excel automatically fills it in with the cells you selected in step 1. You just have to supply a descriptive title (like "Budget Numbers") and a matching password that unlocks the range for editing.
5. **Click OK.**

   Excel adds your range to the list in the Allow Users to Edit Ranges dialog box (Figure 24-15).

   You can click New again to create more protected ranges. Or, you can click OK to close the Allow Users to Edit Ranges dialog box.

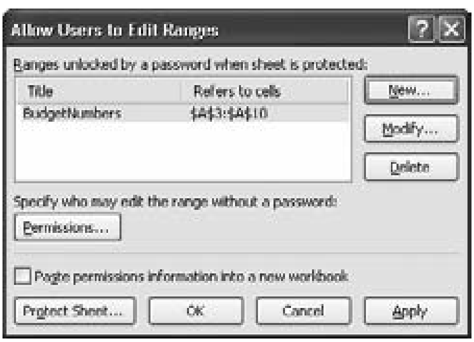

Figure 24-15:
The workbook associated with this dialog box has a single protected range. Using this dialog box, you can modify a range (click Modify), or remove it (click Delete). You can also add Windows permission settings (discussed in the next section).

6. **Choose Review→Changes→Protect Sheet, enter a master password (that only you know), and then click OK.**

   Remember, your ranges don't become password-protected until you turn on protection for the worksheet.

   Now, if people need to edit certain sections of your worksheet, you can give them the appropriate range password. However, you won't share the password that turns off worksheet protection with anyone. That's for your use only.

Figure 24-16 shows what happens when you start typing in a protected range.

Figure 24-16:
A nice feature about protected cell ranges is that the person using the workbook doesn't need to explicitly turn off the protection. Instead, when someone starts typing in one of the cells in the protected range, Excel pops up a dialog box asking for the password, as shown here. Fill it in, and you're ready to go.

Allowing Specific Windows Users to Edit a Range

When protecting cell ranges, there's one more wrinkle. As with worksheet and workbook protection, you can restrict access using a password, which the person using your workbook must supply before changing one of the protected cells. Alternatively, you can allow people based on their Windows *user account*. (User accounts are the login system Windows uses to let multiple people share a single computer.)

This option's a bit trickier—in order for it to work right, the person who uses your workbook needs to log on to the same computer you use, or the same network server. In other words, this technique can work for a small team of people working on a company network, but it's no good for more independent Excel fans.

Keep one other point in mind. The user account permissions work *in addition* to the cell range password. In other words, you follow the same process described in the previous section to lock up a range of cells using a specific password. Then, you give some special people the ability to edit the cell range straight away—no password required.

Note: If you never intend to actually use the cell range password, you can keep it a secret. However, the password's a good fallback if people need to edit your workbook when they aren't connected to the company network (for example, if someone takes it home on a laptop). At that point, Excel can't recognize the user, so the person who's using the spreadsheet needs to use the cell range password to unlock the cells.

Here's how to apply user account permissions:

1. **Choose Review→Changes→Allow Users to Edit Ranges.**

   The Allow Users to Edit Ranges dialog box appears.

2. **Select the range you want to use.**

   If you haven't created the range yet, you can click the New button to add it now. Follow the steps in the previous section to create the range.

3. **Click Permissions.**

   A Windows permission dialog box appears. You may have seen a window like this before if you've ever modified the permissions on your files. The permission dialog box lists all the people who are allowed to access the range, and can bypass the password. At first, this list is empty.

4. **Click Add to enter your first user.**

   The Select Users or Groups dialog box appears (Figure 24-17).

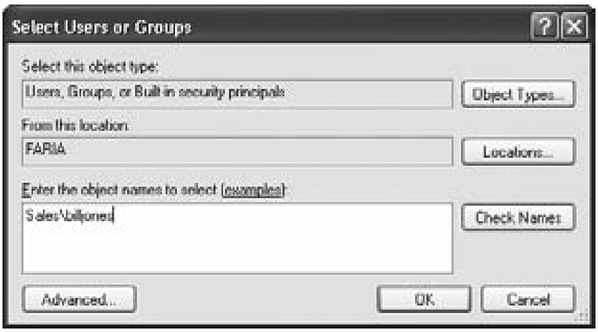

Figure 24-17:
You can safely ignore all the buttons in this dialog box. All you need to do to add a user is enter the user name in the "Enter the object names to select" box. In this example, the user billjones (who logs into the domain Sales) is being added.

5. **In the big text box, type in the user name.**

   Excel needs to know two things: where the user logs in (on the local computer, via a network server, or Windows domain), and who the user is (the user name). To fill in a user name correctly, you need to supply both pieces of information, separated by a backslash (\). For example, if billjones logs into the Sales domain, the full user name is *Sales\billjones*.

Tip: Forgot your user name? One easy way to get your full user name is to pay attention when you log onto your computer. You'll find all the details in the Windows login box. Or, if in doubt, talk to your friendly neighborhood network administrator.

Instead of using a user name, you can use a group name to save time and avoid headaches. The box on page 713 explains how this works.

TIMESAVING TIP

Put Your Users in Groups

It's easy to get in over your head when giving permissions to specific users. Before you know it, a workbook that you initially planned to share with one person is being passed around an entire company, and you're spending every lunch hour feverishly editing the list of allowed users.

Skilled network administrators use *groups* to simplify this kind of task. For example, *billjones* and everyone else on his team might be assigned to the same *SalesEmployees* group. That way, you can configure your Excel workbook so it allows the entire SalesEmployees group, rather than each person. This approach is quicker to set up and easier to maintain. Best of all, when new hires are placed in the SalesEmployees group, they automatically get access to your workbook. And if Bill goes on a bender and the company gets rid of him, the network administrator can pull him out of the SalesEmployees group. He won't have access to your workbook any longer, and you don't need to waste a second changing your protection settings.

6. **Click OK to add the user.**

   If you've made a mistake—for example, typed in a user who doesn't seem to exist or isn't defined on the computer or domain you picked—you'll wind up at the Name Not Found dialog box, where you can try to correct the user name.

   If you've entered a valid user range, Excel adds it to the list of allowed users, as shown in Figure 24-18.

Figure 24-18: *When you add a new user to the list, that user automatically gets a checkmark in the Allow column, which indicates that he's allowed to edit the range without a password. You can also use the Deny column to explicitly lock out specific users, but that approach rarely makes sense, because it introduces too many potential security holes.*

7. **You can now return to step 4 to add somebody else. When you're finished, click OK.**

   You're returned to the Allow Users to Edit Ranges dialog box.

8. **Click OK to get back to your worksheet.**

   Remember, you need to choose Review→Changes→Protect Sheet to protect the worksheet in order for your settings to take effect.

Tip: If you want to get a full breakdown that lists all your protected ranges, their passwords, and the people who have access to them, choose the "Paste permissions information into a new workbook" setting before you click OK. Excel creates a new workbook in a separate window, and copies this information into that workbook.

Once the worksheet is protected, it behaves a little differently. Now, the allowed users don't have any idea that they're typing in a restricted region. For example, if *billjones* starts editing one of the protected cells, Excel quietly checks the user account, sees that it matches your list, and lets him type away without any disruption. On the other hand, if someone else tries to make a change to one of the restricted cells, Excel pops up a dialog box asking for the password, just as it did before (as shown in Figure 24-16).

CHAPTER 25

Worksheet Collaboration

When Microsoft first created Excel, the personal computer was a piece of standalone hardware, capable of some remarkable feats but cut off from the rest of the world. Today, with local networks and the ever-growing Internet, you have the ability to share information, exchange ideas, and collaborate on projects with a large group of friends or a team of coworkers. Excel joins the party with a set of three useful collaboration features that lets groups of people work together to edit spreadsheets. Specifically, you can:

- **Add comments**. Excel's comments feature lets you insert questions, suggestions, or other miscellaneous notes that point to specific cells (like "This number's wrong" or "Please boost the sales estimate so we can impress the boss"). The person who created the spreadsheet can respond to these comments by modifying the data accordingly.
- **Track changes**. Change tracking is the real muscle in Excel collaboration. Change tracking lets you keep track of the edits made by multiple people. You can then choose to apply or reject some or all of the changes. If multiple people make changes to different copies of the same document, you can even merge all their changes back into the original file in one step, saving countless headaches.
- **Protect your spreadsheets**. In Chapter 24, you learned how to lock down worksheets so people can make changes only in specific cells. This feature is handy when you're sharing workbooks and you want to make sure no one futzes with your masterpiece.

In this chapter, you'll take a close look at comments and change tracking. (For a refresher on worksheet protection, see Chapter 22.) You'll also consider an even more radical (and risky) proposition—letting multiple people change the same workbook simultaneously.

Your Excel Identity

Excel knows more about you than you might think. Behind the scenes, all of Microsoft's Office applications store some basic information about you, including your user name. This information is set when Office is first installed (even if someone else installed it), and it usually matches your Windows user name (so if you log on as joeZhang, your Excel user name is also joeZhang).

For most tasks, it doesn't matter who Excel thinks you are. But making sure you have the right identity becomes much more important when you need to collaborate with other people. For example, if you're adding comments regarding someone else's work, you want to make sure they know that it's Tim Smith making the suggestions (and not "Excel User 1" or "SalesComputer012").

Note: If your computer has more than one user account, remember that every Windows user has a distinct set of personal settings. So when someone else logs onto the same computer, the Office applications use a different user name.

The best idea is to make sure that Excel has the correct name stored for you before you start to use comments and change tracking. To see the name Excel is using, select File→Options. The name Excel uses is stored in the "User name" text box at the very bottom of the Excel Options dialog box. You can edit the name now, and then click OK to apply your change. Spaces are perfectly acceptable, so feel free to replace "smacdo1123" with "Sue MacDonnell."

Note: Any changes you make to the user name in Excel affect all the other Office applications, including Word, which has its own review and collaboration features.

Preparing Your Workbook

Before you go ahead and release your workbook into the wild, you may want to perform a few last-minute tasks (things like giving your spreadsheet a descriptive name, checking whether it's encrypted, and so on). Excel helps you out by grouping some of the most common collaborative-prep tools in backstage view. To see them, choose File→Info (Figure 25-1).

The following sections walk you through the most important options in the File→Info page and show you how they can prepare your spreadsheet for a wider audience.

Note: The only section you won't consider is the Versions section, which lets you recover old versions of a workbook after you make changes. This feature, which doesn't have anything to do with document sharing, is discussed on page 82.

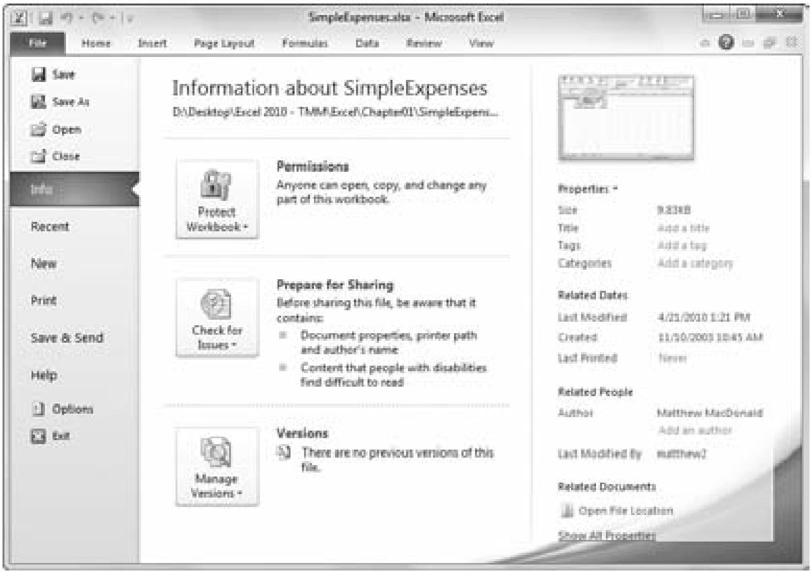

Figure 25-1: *Before your workbook goes out the door, you may want to consider switching on document protection, scouring your spreadsheet for potential issues, and changing its properties. File→Info page in backstage view is a springboard for all of these tasks.*

Workbook Protection

The File→Info page is split into several sections, each with a large button. The first button, named Protect Workbook, helps you to prevent other people from changing things they shouldn't in your workbook.

When you click the Protect Workbook button, a menu pops up with a list of options. You've seen some of these options in previous chapters. They are:

- **Mark As Final** designates your document as "final," meaning no more changes are needed (or allowed). Once you take this step, Excel's editing commands switch off, you can't type in any cell, and a special icon appears in the Status bar (it looks like a rubber stamp and a sheet of paper). However, it's important to realize that the Protect Workbook→Mark As Final command is just for descriptive purposes. In other words, it tells people the document is finished, but it doesn't prevent them from "unfinalizing" the workbook and making changes. In fact, anyone can make the spreadsheet editable simply by choosing Protect Workbook→Mark As Final again.
- **Encrypt With Password** lets you scramble your document so absolutely no one can open it without supplying a password. For more information, refer to the detailed steps on page 50. It's also worth noting that the Protect Workbook→Encrypt With Password command adds a "password to open" but not a "password to modify." If you want the latter, which lets you create a document that unauthorized people can view but not change, refer to page 50.

- **Protect Current Sheet** and **Protect Workbook Structure** let you apply worksheet protection, as you learned in Chapter 24. Worksheet protection is a more selective restriction tool than encryption; it can limit edits to certain cells or permit only certain types of modifications.

Tip: Encryption and protection play complementary roles, and you may need both. If you're worried about industrial spies cracking into your files to steal data or create fake copies of your workbook, you need encryption. But if you want to control the kind of changes that authorized people can make, you need worksheet protection.

- **Restrict Permission by People** lets you add advanced restrictions to your workbook using an add-in called IRM (Information Rights Management). For example, IRM can prevent people from copying worksheet data to the clipboard or make a workbook inaccessible after a certain date. However, IRM is also more complex to set up, and it depends on Microsoft's web servers to perform authentication by asking people for their Windows Live IDs. Page 708 has more.
- **Add a Digital Signature** lets you add a special marker to your spreadsheet file that "proves" you created it. The downside's that in order for a digital signature system like this one to be meaningful, you need to set up a complex system of personal certificates and certificate *authorities* that vouch for a given certificate's validity. In other words, unless you're working in a large company (or you're ready to buy a certificate from an online certificate authority), your certificates won't be worth much. (If you *are* ready to buy a certificate, head to the website of a certificate authority like *www.verisign.com* or *www.thawte.com*.)

Checking for Issues

Also in the File→Info page, just under the Protect Workbook button, the large Check for Issues button can help you catch any lingering problems before you share your work with the world. When you click this button, you'll get a menu with just three options:

- **Inspect Document** runs the Document Inspector, a tool that helps you track down scraps of information you may not want others to see (Figure 25-2).
- **Check Accessibility** runs the Accessibility Checker, which flags content that people with visual disabilities might find hard to read. For example, people who use screen-reading software won't get much out of pictures unless the images have descriptive text defined (known as alternate, or just alt text). Similarly, they won't find standard worksheet names (Sheet1, Sheet2, and so on) much help when navigating a workbook to find the content they need.
- **Check Compatibility** runs the Compatibility Checker—a tool that flags potential problems if you plan to share your document with people using Excel 2003 or older. Page 41 has the full story.

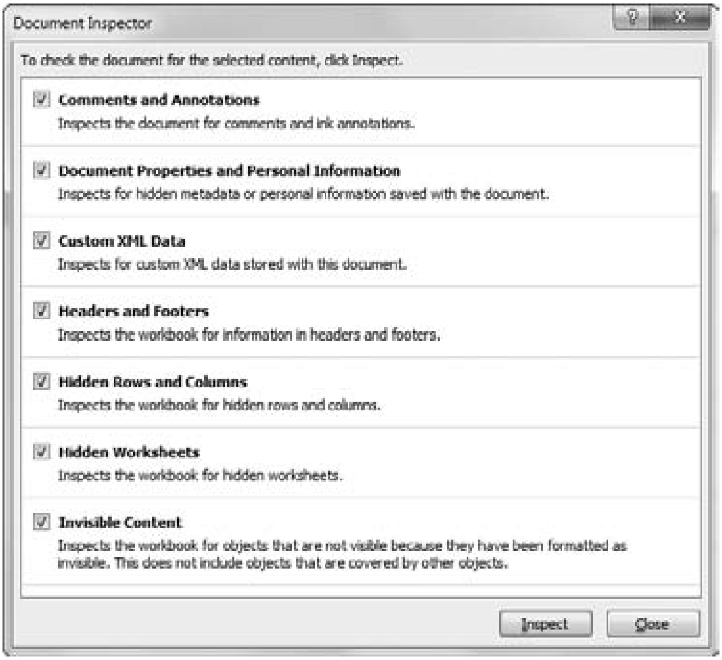

Figure 25-2: *The Document Inspector is a tool for the truly paranoid. Turn on the checkboxes that you want the Inspector to sniff out, and then click Inspect. If it finds that type of content, it gives you the option of removing it sight unseen by clicking a Remove All button (which appears after you've run an inspection, if the Inspector's found something). Typically, you'll use the Document Inspector tool if you're creating a stripped-down copy of a workbook, and want to make sure there isn't personal information lingering in a comment, page header, or somewhere else. But make sure you save a new, separate version of your document, with a different file name, because you may also lose some information you need! (Like that comment reminding yourself to ask your boss for a raise.)*

Document Properties

At the far right of the File→Info page, you'll find a panel of document properties that lists key bits of information about your workbook, like its title and author.

Some of the document properties can't be edited, because Excel tracks them automatically. These properties include details like the file size and the last revision date. You can change other properties, like the title. To edit one of these details, simply move your mouse over the property. The property changes into a text box, which you can then click to start editing. The only exception is the author information, which requires the two-step editing process shown in Figure 25-3.

Tip: To set even more descriptive details, click the Properties heading, and then choose Show All Properties. Excel adds entries that let you fill in details like status, subject, company, and comments.

Figure 25-3: *Left: To change the author, begin by clicking and typing in the "Add an author" text box. This adds a second author to the list.*

Right: To remove the original author, right-click that name, and then choose Remove Person. Or, leave the original author if you want to list multiple people.

When you create a new workbook, Excel uses your Excel user name (page 716) to set the author name. It leaves most of the other details—for example, title, category, company, tags, and so on—blank. It's up to you to add these details if you want them. All of this raises an excellent question: now that you know how to set the document properties, what do you actually do with them?

The primary role of Excel's document properties is to supply helpful information to other people who might need it. This information turns up in a few interesting places. You'll see it when you hover over a file icon or investigate file properties in Windows Explorer (Figure 25-4). You can also add it to the footers and headers you use in your printouts (page 220). Finally, searching tools (like the one built into Windows) will search the information in the document properties first. In fact, that's the whole purpose of the tags property: you add descriptive keywords, like "budget" and "2010," that someone might use when performing a file search later on. That way, they can locate relevant files quickly.

If you find yourself changing document properties often while you work, you don't need to keep switching to backstage view. Instead, you can show the Document Properties panel, which sits just above your worksheet and provides text boxes for quick property editing (Figure 25-5). To show the Document Properties panel, choose File→Info to switch to backstage view, and then click the Properties heading (at the far right side of the window) and choose Show Document Panel.

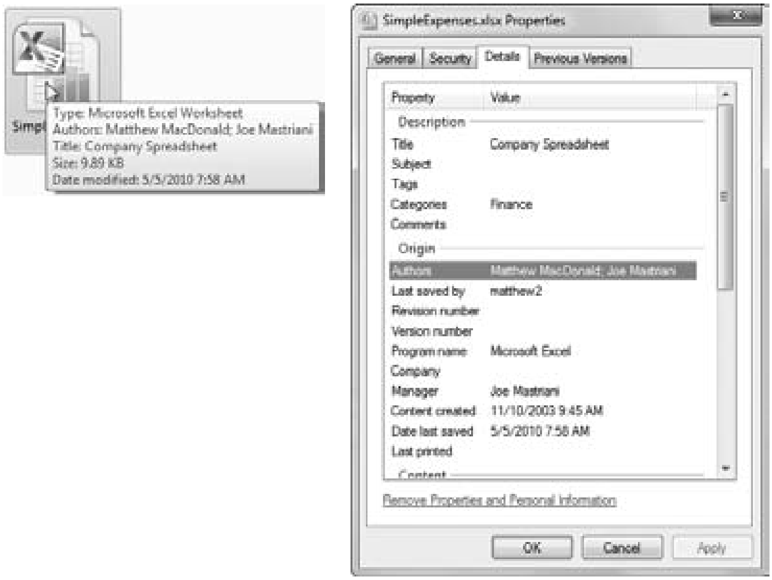

Figure 25-4:
Left: When you hover over an Excel file on the desktop, a tooltip pops up that lists some of the document properties (assuming they're not blank).

Right: Right-click an Excel file, and then choose properties to get the familiar multi-tabbed Properties dialog box. If you choose the Details tab, you'll see a slew of details about the file, including important dates (when it was last saved, when it was last printed, and so on), and its document properties.

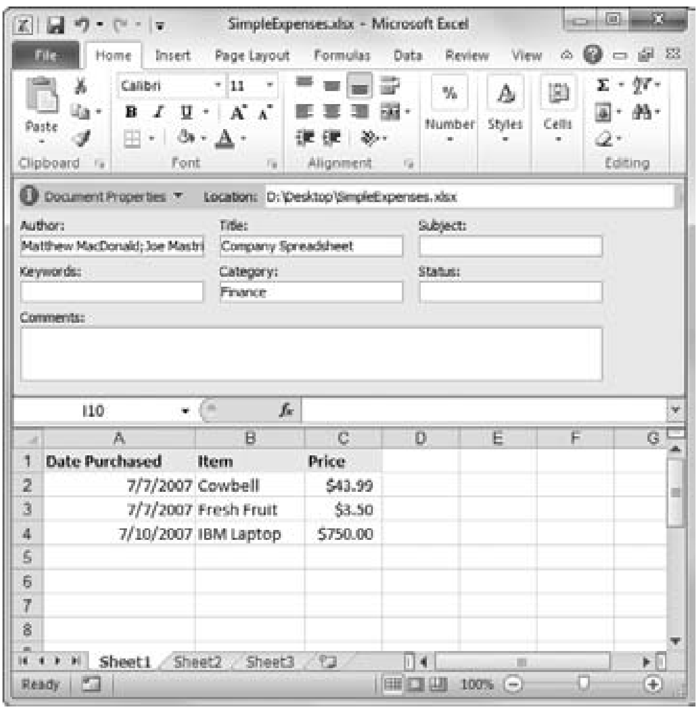

Figure 25-5:
The Document Properties panel shows key documents properties, including the title of the workbook, its subject, extra comments, and so on. To fill in this information (or change it), click the appropriate box, and then type away. To quickly close the Document Properties panel, click the X at the far right. Most people find that the Document Properties panel takes up too much room, and they don't keep it around for long.

Distributing a Document

Once your document is ready to go, you have plenty of ways to share it with others. You can distribute your workbook just as you'd distribute any other type of file. Possibilities include attaching the workbook file to an email, burning it to a CD, copying it to a memory stick so you can carry it around, or copying it to a network drive that others can access. All these techniques work perfectly well, as thousands of Excel fans will attest.

But if you're really in the Excel zone and don't want to emerge from the comforting shelter of your worksheet grid, you'll be interested in Excel's document sharing page, which can send your workbook on its way without forcing you to take the extra step of firing up a Web browser or email program. To try it out, choose File→Save & Send, which brings you the page in backstage view that has all the options (Figure 25-6).

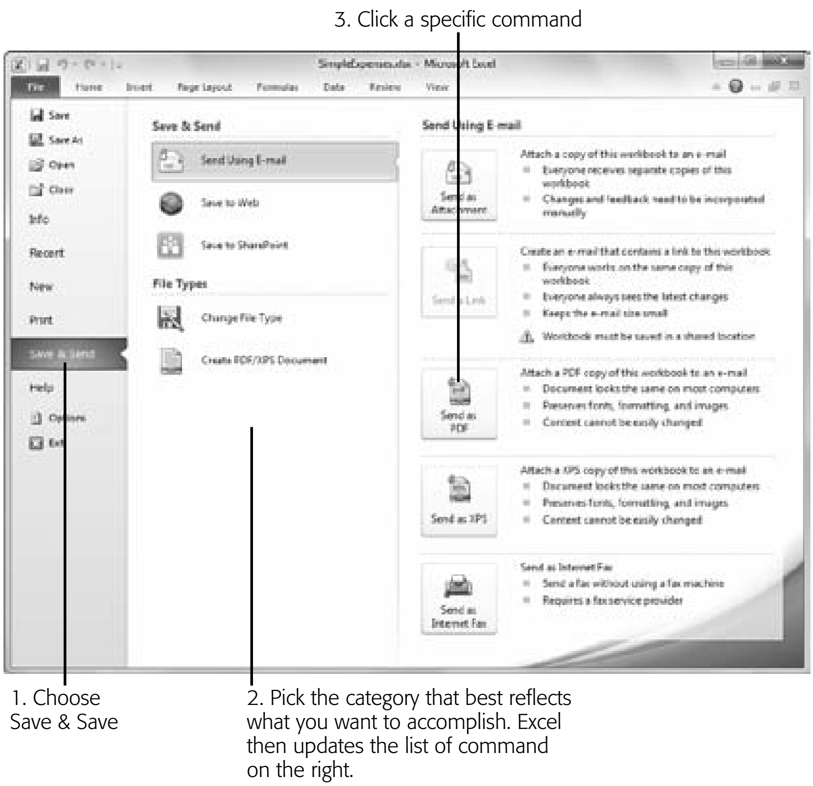

Figure 25-6: *Excel's backstage view gives you several ways to distribute your document, including by email, Web, and SharePoint. Start by choosing one of the options in the middle portion of the window, and finish by clicking a related button on the right. For example, if you choose Send Using E-mail in the middle, you'll see the options for attaching your workbook in its original glory or PDF, as shown here.*

Sending by Email

The Send Using E-mail option lets you send your workbook as an attachment to an email message. On the right are several choices:

- **Send As Attachment** copies your workbook and attaches it to a new email message. To create the email message, Excel uses the email program installed on your computer, like Microsoft Outlook. It won't let you use a Web-based email account like Hotmail or Google Mail. When you click the Send As Attachment button, Excel launches your email program and opens a new message with your workbook file attached, as shown in Figure 25-7.

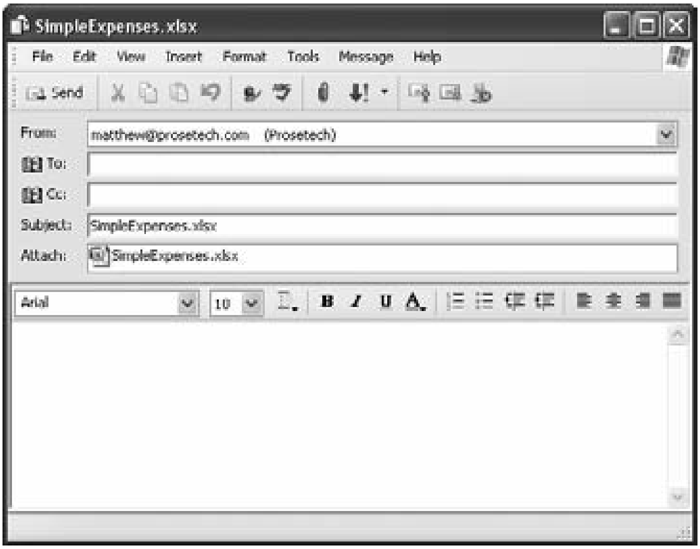

Figure 25-7: *The Excel workbook file SimpleExpenses.xlsx is attached to this message. Fill in a recipient, subject, and message, and then click Send. Or, you can close the message window to cancel the message.*

Note: Excel uses the email program specified by your computer's Internet Options control panel. Usually, that's the same program you use on a regular basis, but to make sure, from the Windows Control Panel, select Internet Options, choose the Programs tab, and then see which program is specified in the E-mail box. If you'd like to pick a different email program, choose one from the drop-down list.

- **Send a Link** creates an email message with a link that points to your workbook. The idea is that the recipient of that message can click the link to open the original copy of the document, rather than getting a duplicate. However, this button works only if the worksheet is stored somewhere public, like a shared folder on a networked computer. It also assumes that the shared location is accessible to the recipient's computer.

- **Send as PDF** and **Send as XPS** work like Send As Attachment, but they create a copy of the workbook in PDF or XPS format, respectively. This is handy in two cases: if your recipient doesn't have Excel, or if you don't want your recipient to edit your work. Page 46 has more about the PDF and XPS standards.
- **Send as Internet Fax** lets you send the document to a fax number, just as if you'd printed it up, and then fed it into a fax machine. The drawback is that you need a *fax provider*, which is an Office plug-in that can transmit the electronic data to a real phone number, with no modems or telephone connection required. The first time you click Send as Internet Fax, Excel opens a Web page that lists a variety of fax providers you can consider.

Uploading to the Web

Email isn't the only way to route your workbook to other people. You have dozens of other options, including copying it to a network, uploading it to a website, placing it on a USB stick and walking over, and more.

The Save to Web option lets you upload your spreadsheet file to Microsoft's free document-hosting service, called SkyDrive. SkyDrive isn't just a useful place to stash your spreadsheet files; it's also the starting point for browser-based Excel editing with the Excel Web App. You'll learn to use this feature in Chapter 26.

The Save to SharePoint option lets you upload your spreadsheet to SharePoint, a document-sharing tool built into Microsoft's Windows Server operating system. Using SharePoint, you can create team websites where people can share Office documents, discuss their documents, and collaborate. Best of all, you can assign different permissions to different people, ensuring that only the right people can make changes or review sensitive data. Big businesses (with deep pockets) use SharePoint to coordinate Office workflows, making sure the right people review the workbook you create, in the right order.

Tip: SharePoint is outside the scope of this book, but you can learn more by checking out a dedicated book or heading to one of Microsoft's many websites. Visit *www.microsoft.com/office/preview/servers/sharepointserver* for the marketing spiel and *http://msdn.microsoft.com/office/server/moss* for heavy duty tech talk about setting up SharePoint systems.

Adding Comments

Comments are the simplest of Excel's collaboration features. Excel displays comments in a floating yellow box that points to a single cell (as shown in Figure 25-8). Place whatever descriptive text you want inside the comment box. For example, you can use comments to flag an error, raise a question, make a suggestion, or praise a particularly brilliant formula.

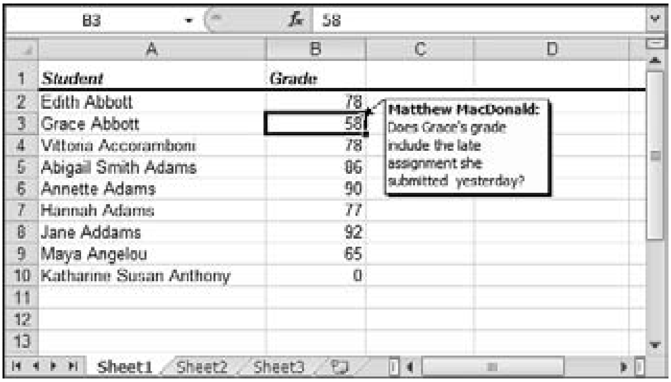

Figure 25-8: *This worksheet includes a single comment attached to cell B3. You can clearly see the arrow that connects the floating comment box to cell B3. Additionally, a tiny red triangle in the top-right corner of the cell indicates it has an attached comment.*

The beauty of comments is that you're free to include as much information as you want without modifying the worksheet data. For that reason, comments work perfectly when you're sending a workbook out for review. An employee can send an expense report to a manager, and the manager can add feedback using comments without altering the original information. Once the manager sends the workbook back, the employee can then decide whether to heed the comments and make some changes, keep the comments for later consideration, or just remove them altogether.

Inserting a Comment

Every workbook can include thousands of comments. The only limitation is that each individual cell can have only one attached comment.

To create a new comment, just follow these steps:

1. **Move to the cell where you want to place the comment.**

   Every comment anchors to a single cell in a worksheet. You can move the comment to any location you want after you create it, but it always points to the same cell (using a long arrow).

2. **Select Review→Comments→New Comment.**

   A new comment box appears next to the cell. Excel fills in your name on the first line, and positions the cursor inside the text box so that you can start typing.

Tip: Excel automatically uses the name that it has stored for you. You can edit this part of the comment, but rather than change the name every time you add a new comment, it's far easier to change your name information in Excel's settings, as explained on page 716.

3. **Enter the text for your comment.**

   You can space your comment out over multiple lines by pressing Enter to jump to the next line. The comment box scrolls down if you enter more text than can fit in the visible area.

4. **When you're finished entering the comment text, click the original cell to return to the worksheet, or press Esc twice.**

   Excel marks commented cells with a tiny red triangle in the top-right corner. To see the comment, hover over the cell with the mouse. Figure 25-8 shows an example of a worksheet with a single comment.

Showing and Hiding Comments

Ordinarily, comments appear when you hover over the commented cell with the mouse. When you move somewhere else, they politely disappear from view (although Excel still gives you a clue which cells have comments, as shown in Figure 25-9). This behavior makes sense if you're dealing with a lot of comments from multiple reviewers. Otherwise, you'd see so many comments that they'd obscure each other—or even important worksheet data.

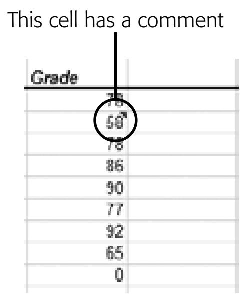

Figure 25-9:
When comments are hidden, you still know they exist, thanks to the small red triangle in the top-right corner of the cell with the comment.

Sometimes, though, you want to make sure a comment is clearly visible and that no one will overlook it. In that case, you can set a comment so that it's *always* visible, no matter where your mouse is. To do this, move to the cell that has the comment, and then choose Review→Comments→Show/Hide Comment. Use the same command a second time to tuck the comment back out of sight.

Tip: You can use the Review→Comments→Show All Comments command to display every comment in your worksheet at once. Comments always appear wherever they were last situated.

Fine-Tuning Comments

Once you've created a comment, you can manipulate it the same way that you can manipulate other floating objects like charts and graphics. To start off, move to the cell that has the comment. Then, choose Review→Comments→Edit Comment. (Or, if you've already used the Review→Comments→Show/Hide Comment command to make the comment visible, you can just click the comment box to select it.) You'll see resizing handles appear around the comment box. You can now perform the following tasks:

- **Move the comment box**. It's quite likely that the place where the comment box first appears isn't exactly where you want it. The comment may obscure important information on the worksheet. Fortunately, you can easily drag the comment out of the way. Just move the mouse pointer over the border of the comment so that it changes into a four-way resize arrow, and then drag the comment box to its new location. Excel automatically adjusts the arrow that connects the comment box to the original cell.
- **Resize the comment box**. To resize a comment, click one of the resizing handles and drag the box's edge or corner. (The resizing handles look like circles at the edges of each corner and at the middle of each side of the comment box.) You may want to resize a comment box to enlarge it so it can show all the comment text at once. Even though you can scroll through comment box text using the cursor, it's not always obvious that some of the text is out of sight because the comment box doesn't show any scroll bars.
- **Edit the comment**. Just click inside the comment box, and then start editing. Or, right-click the cell with the attached comment, and then choose Edit Comment.
- **Delete the comment box**. To delete a comment, click the border of the comment box to select the whole box, and then press Delete. Alternately, you can right-click any cell with an attached comment, and then choose Delete Comment from the pop-up menu.

Tip: To delete multiple comments at once, select all the cells that have the attached comments, and select Review→Comments→Delete. But be warned, Excel doesn't ask for confirmation before it removes your comments.

You can also format the text in a comment box. Often, you'll take this step to make your comment stand out (by enlarging or bolding the font) or to fit more text in (by shrinking the font). You may also use a specific color if you know that multiple reviewers will add comments to the same worksheet, and you want to be able to distinguish the comments from different authors at a glance.

To format a comment, select the text you want to format, switch to the ribbon's Home tab, and then use the Font and Alignment sections. You can apply a new text color, a new font and size, and different alignment options. Figure 25-10 shows a couple of examples.

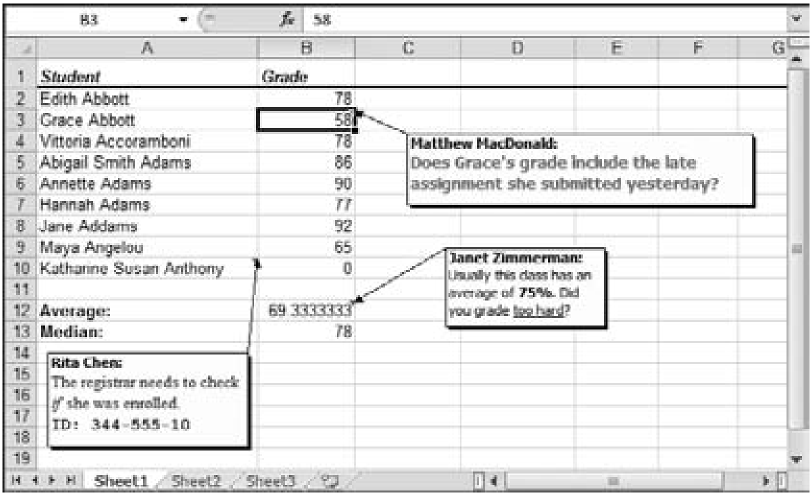

***Figure 25-10:** Formatting helps you tweak all or a portion of your comment text. You can use it to highlight in boldface a piece of important information in a comment (like the student ID in Rita's comment or the percentage in Janet's comment). Or, it can help distinguish the comments from different reviewers.*

Reviewing Comments

If it's your job to review everyone else's comments (and make the requested changes), you'll be interested in Excel's buttons for *comment navigation*. These buttons let you move through all the comments in a worksheet, one at a time. Best of all, this feature doesn't force you to unhide any comments. Instead, you can deal with them one at a time, which keeps your screen clutter-free and your sanity intact.

To move through your comments, start at the first cell of your worksheet (A1), and then choose Review→Comments→Next. Excel scans the worksheet starting from the current cell, and then moves to the right, one cell at a time. If Excel doesn't find any comments in the current row, it scans the next row from left to right, starting in the first column. When Excel finds a comment, it stops the search and selects that comment. If the comment was hidden, it now appears on your worksheet.

To keep moving through the worksheet, click Review→Comments→Next again (or use Review→Comments→Previous to move backward). As you move on, previously hidden comments are concealed once more.

Printing Comments

Excel's standard behavior is to ignore all comments when it prints a worksheet. Your printed document won't show the comment text or even indicate that a cell has a comment.

If you'd like a printed record of your comments, Excel gives you two options:

- **You can print the visible comments on your worksheet**. In this case, Excel draws the graphical comment boxes in the printout exactly as they appear, potentially obscuring other worksheet information. Hidden comments don't show up.
- **You can print all the comments on a separate page**. In this case, Excel creates a list of comments. Each entry in the list indicates the cell reference, and the comment text.

To change the option you're using, follow these steps:

1. **Head to the Page Layout→Sheet Options section, and then click the dialog box chooser (the small square with an arrow in it) at the bottom-right corner.**

   The Page Setup dialog box appears, with the Sheet tab displayed.

2. **Make a selection from the Comments list box.**

   You can choose "None" (the default), "At end of sheet" (which creates a separate comment page), or "As displayed on sheet" (which shows the graphical comment boxes). Figure 25-11 compares your options.

Note: Excel doesn't let you print just the comment page. If you choose "At end of sheet," you must print the comments page *and* the worksheet data.

3. **Click OK.**

   Excel stores your comment options for this worksheet. You can now use the File→Print command to send your data to the printer.

Tracking Changes

Comments are a great way for people to leave messages for each other in a spreadsheet. However, they become awkward if the workbook requires substantial changes, and they're downright aggravating if more than one person revises a workbook. Imagine a worksheet that tracks a team's weekly progress or one that represents a communal effort to create a business plan. In such situations, where more than one person needs to make substantial additions or changes to the worksheet data, comments just can't handle it all.

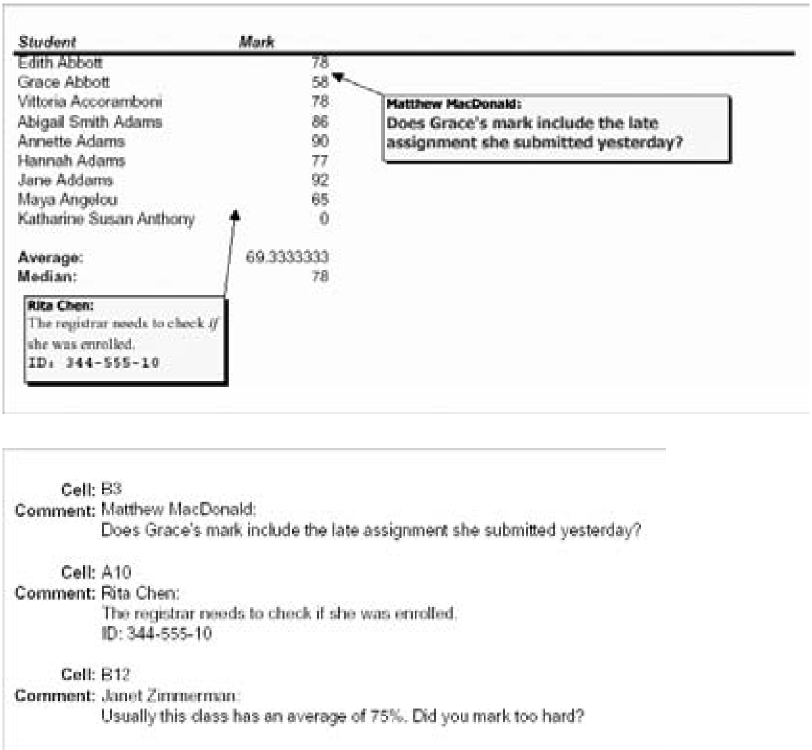

Figure 25-11: *Here are two different ways to print comments in Excel. The top example shows the result of the "As displayed on sheet" option, which prints only the currently visible comments and may obscure worksheet data. The bottom example shows the result of the "At end of sheet" option, which prints a separate comment page that's useful for at-a-glance review.*

Excel provides another tool that makes it easier for groups of people to work together—*change tracking*. Change tracking makes sure that the changes made by different people are carefully logged, giving you the power to inspect each person's changes individually and reverse them if you choose. It's a little like keeping several versions of the same worksheet in a single spreadsheet file.

You may consider using change tracking for several reasons:

- **You want to send a workbook to another person for review or editing**. However, you want to be able to quickly spot the changes they make.
- **You want to have the last word on other people's changes to your workbook**. In other words, you not only want to review changes, but also want to discard them if they aren't correct.

- **You want to distribute copies of your spreadsheet to several people at once**. Once everyone has made their changes, you want to merge all these changes back into the original copy.

Note: In one of Excel's most irritating quirks, you can't use change tracking if your workbook includes a data table (like the ones described in Chapter 14). Your only recourse is to move into the table, and then convert it to a group of ordinary cells by choosing Table Tools | Design→Tools→Convert to Range.

Switching On Change Tracking

To turn on change tracking, follow these steps:

1. **Select Review→Changes→Track Changes→Highlight Changes.**

   The Highlight Changes dialog box appears.

2. **Turn on the "Track changes while editing" checkbox.**

   When you turn on change tracking, you also automatically switch on workbook sharing (a separate feature that's described on page 741). There's no way to turn on change tracking without also turning on workbook sharing.

3. **Ignore the Who, When, and Where checkboxes for now, and click OK to return to your worksheet.**

   The Who, When, and Where checkboxes don't affect how change tracking works. They just configure what changes Excel highlights in your document. You'll learn about these settings a little later in this chapter.

   Excel automatically saves your workbook at this point. If you're turning change tracking on for a brand new spreadsheet, then you need to pick a file name and location. Otherwise, just click OK when Excel prompts you to save an updated copy of your workbook.

To see how change tracking works, just enter text in a new cell or edit an existing cell. Excel uses several visual indicators to show that a cell has been changed. Most obviously, it draws a blue outline around the cell and adds a small blue triangle to the cell's top-right corner. Additionally, Excel changes the column headers at the top of your sheet to use red text if that column or row contains changed data.

To get the specifics about the change, you can hover over the cell with your mouse. When you do, a yellow box appears that looks just like a comment box. However, you can't edit the text in this box—instead, it contains a message that indicates who last changed the cell, the date and time the change was made, and what the change was. This message disappears as soon as you move the cursor somewhere else. Figure 25-12 shows a closer look.

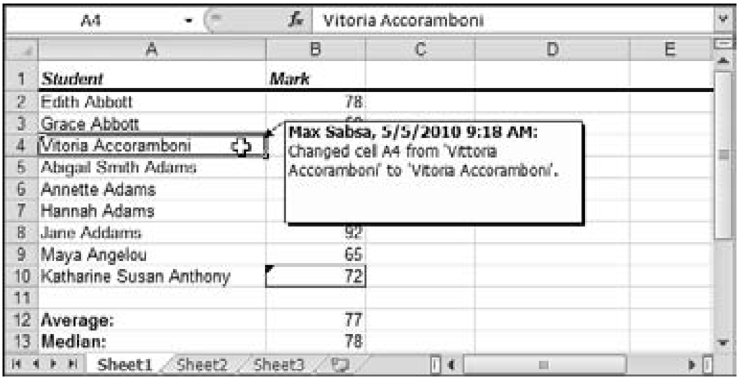

Figure 25-12:
Cells A4 and B10 have the telltale blue triangles indicating they've been changed. To examine a change, simply position your mouse pointer over the correct cell. Using this technique, you can see the name in cell A4 was modified by removing an extra "t" in the word "Vittoria."

Note: Change tracking applies to the entire workbook. When you switch on change tracking, Excel tracks changes for each cell in every worksheet.

FREQUENTLY ASKED QUESTION

Features That Change-Tracking Affects

I turned on change tracking, and now a bunch of other features don't work–what gives?

Unfortunately, when change tracking is on, Excel automatically prevents you from taking certain actions. This side effect is unavoidable, because change tracking requires workbook sharing (page 741). Workbook sharing always applies a few basic safety restrictions.

As a result, once you turn on change tracking, each of the following changes are off limits:

- Creating a data table, picture, chart, or pivot table.
- Deleting a worksheet.
- Merging cells or splitting merged cells.
- Adding or changing conditional formats.
- Adding or changing data validation.
- Inserting or changing hyperlinks.
- Protecting or unprotecting worksheets.
- Creating, changing, or viewing scenarios.
- Grouping or outlining data.
- Recording, editing, or assigning macros in the current workbook.

In all these instances, you're limited in what you can create or change. If these features already exist in your workbook, they'll continue to function. Once you turn on change tracking, you can't add a hyperlink. However, if you've already created a hyperlink, it still works as it did before.

Understanding the Change Log

Excel's change tracking is a little quirky, and it often confuses reasonably intelligent people. Even though it's superficially similar to the change tracking features used in other programs (like Microsoft Word), it's subtly different in many ways. In this section, you'll learn how it really works.

Once you turn on change tracking, Excel keeps a *change history log*, which it stores along with your workbook file. The change history log lists every modification you've made since you first turned on change tracking. It also records who made the change and when.

Excel records all changes on a per-cell basis. So if you change the text in cell A2 from *John Smith* to *Adam Bergman*, Excel adds one change item to the log, indicating the old and new values. If you change the cell text twice, Excel adds two entries. Excel records the change as soon as you commit it (either by pressing Enter or moving to a new cell).

Excel doesn't track every possible change. In fact, it ignores the following changes:

- Formatting changes (like when you change the font or background color for a cell).
- Hiding or unhiding rows or columns.
- Adding, changing, or deleting comments.
- Inserting or deleting worksheets. (However, if you add a new worksheet, Excel does track all changes you make to *that* worksheet.)

Excel also doesn't worry about cells that change indirectly. For example, if you change the value in cell A1 and that causes a formula in cell A2 to display a new value, Excel records a change only for cell A1. If you think about this approach, it makes perfect sense, because the real content in cell A2—the formula—hasn't changed at all. The only difference is the displayed result.

Changes don't remain in the change history log forever. In fact, once 30 days have passed since a change was made, Excel discards it from the log. Excel checks the date and discards old entries every time you open a file, so make sure you aren't sharing workbooks between two people with drastically different dates set on their computers, or some changes will go missing!

Tip: If you decide you need more time, you can adjust how long Excel hangs onto the change log. To do so, select Review→Changes→Share Workbook. Select the Advanced tab, and then change the number next to the "Keep change history for" option. You've told Excel the number of days that changes stay in the log.

You can clear changes from the change history log another way—by turning off change tracking altogether. Just choose Review→Changes→Track Changes→Highlight Changes to show the Highlight Changes dialog box, and turn off the checkmark in the "Track changes while editing" checkbox. You can turn change tracking back on later, but you can't recover the information about the changes you made earlier.

Highlighting Changes

One of the most confusing aspects of change tracking is the difference between how a change is *tracked* and how it's *highlighted*. As you've learned so far, Excel tracks every change you make in a cell. However, it doesn't necessarily highlight these changes so that you can *see* them.

To witness an example of this behavior in action, try saving and closing a workbook that uses change tracking, and then reopening the workbook. You'll immediately see that all the blue triangles that flag the changed cells have disappeared. Even though Excel is still tracking these changes, it automatically switches off change highlighting when you open the workbook, because it doesn't know what changes you really want to see.

Confusingly enough, you use the exact same dialog box to configure change highlighting as you do to turn on change tracking—the Highlight Changes dialog box (Figure 25-13). To show this dialog box, choose Review→Changes→Track Changes→Highlight Changes.

Figure 25-13:
The Highlight Changes dialog box lets you switch change tracking on or off with the "Track changes while editing" checkbox. You can also configure what changes Excel highlights (with blue triangles and pop-up boxes) by using the When, Who, and Where lists.

What you do next depends on the types of changes you want to see. Here are some of your choices:

- **Show all changes**. In the When list, choose All, which is the standard choice that Excel selects when you first switch on change tracking.
- **Show recent changes**. In the When list, choose "Since I last saved" to show all the changes you've made since the last time you saved the workbook (or, if you haven't saved it yet, since you first created or opened it). Excel automatically selects this choice when you open a workbook that uses change tracking.
- **Show changes since a specific date**. In the When list, choose "Since date…"; Excel automatically fills in the current date, but you can edit this date by hand. When you use this option, Excel highlights the changes that were made any time on or after the specified date.

- **Show changes that haven't been reviewed**. In the When list, choose "Not yet reviewed." This option works best when you use Excel's change revision feature to examine each individual change. When you use this approach, Excel highlights only the changes that you haven't yet reviewed and confirmed. For information about this technique, see page 737.
- **Show changes made by a specific person**. In the Who list, choose the name of the person whose changes you want to see. Excel automatically fills the Who list with the name of each person who's made a change in the workbook. You can also choose to show changes by other people and hide your own changes by choosing "Everyone but me" from the Who list. If your worksheet has changes from several different people, you may not see the familiar blue triangles. Instead, Excel tries to color-code the comment flags, so that you can tell at a glance which person left which set of comments.
- **Show changes made in specific cells**. Click the Where box, and then drag the worksheet to select the appropriate range of cells. Alternatively, you can type the cell references in range notation (A1:B3) or as individual cell references separated by commas (A1, A2, A3, B1, B2, B3).

Tip: You can also use a combination of settings in the When, Who, and Where checkboxes. In this case, a change must meet all of the criteria in order for Excel to highlight it. For example, you may want to find changes made to a certain range of cells by a certain person after a certain date.

Once you make your selections, ensure that the "Highlight changes on screen" checkbox is turned on, and click OK. If Excel can't find any changes that match the criteria you've chosen, it displays a warning message. However, Excel still flags any new changes you make, provided they meet the highlight criteria.

Note: The When, Who, and Where checkboxes determine what types of changes Excel *highlights* in your worksheet, but they don't affect what types of changes Excel *tracks*. Once you switch on change tracking, Excel tracks every change you make–no matter who you are or what cell you modify–until these changes expire and get removed from the log.

Examining the Change Log

Change highlighting can get a little crowded. If you have a heavily edited document, you'll quickly run into two problems:

- There's no way to distinguish the most recent changes from ones made earlier. You can change the highlighting settings so you don't see all the changes at once, but it's still difficult to figure out the order in which changes were made.

- When you hover your cursor over a changed cell, Excel shows a box with information about the most recent change for that cell. If a cell was modified more than once, you can't see the information about any of the *earlier* changes—in fact, you have no way of knowing whether they were made.

If you need to find out a little more about the changes made to a document, you can create a *change history report*. This report presents an ordered list of changes. It's not very useful if you want to see the final product of all the revisions, but it's a great tool if you want to find out what's been modified along the way.

To create a change report, just follow these steps with any shared workbook:

1. **Save your workbook.**

   The change report takes into account changes only in the current saved version of your workbook. If you've made changes but haven't saved the workbook, these changes won't appear in the change report.

2. **Choose Review→Changes→Track Changes→Highlight Changes.**

   The Highlight Changes dialog box appears.

3. **If you want to create a change report that shows only certain changes, you can adjust the Who, When, and Where checkboxes, as described earlier.**

   If you want to create a change report that shows all the changes made by everyone (anywhere in the entire workbook), select All from the Who list, and then make sure the When and Where checkboxes aren't turned on.

4. **Turn on the "List changes on a new sheet" checkbox.**

   This setting tells Excel to create the change report worksheet. This worksheet is new, and Excel copies all the information from the change history into it.

5. **Click OK.**

   Excel adds a new worksheet named History to your spreadsheet file. The History worksheet contains a list of the changes ordered from oldest to most recent (see Figure 25-14). By reviewing the change report, you can quickly see a list of all the changes made in the worksheet. Using Excel's change highlighting, by contrast, you see only the information for the most recent change. The bottom of the list has a message indicating when Excel generated it.

Excel doesn't let you delete the History worksheet. However, as soon as you save the workbook, Excel automatically removes the History worksheet. Excel takes this rather drastic step to ensure that you never end up with an out-of-date change report. If you decide you want to keep a change report for a longer amount of time, you can select the cells in the History worksheet and copy them to another worksheet.

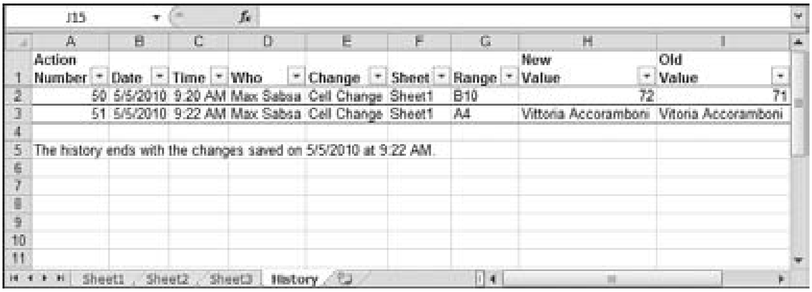

Figure 25-14: *This change report shows three changes. The change report indicates who made the change, when it was made, where it was made, and what the old and new values were.*

Accepting and Rejecting Changes

So far, you've learned how change tracking helps you to see the history of a spreadsheet, letting you examine who modified it and when the changes were made. Change tracking is also useful if you want to review changes and make final decisions about whether to keep or discard each change. This review process is called *accepting and rejecting changes*.

When a change goes into the change history log, Excel flags it as a new, unconfirmed edit. If you decide not to perform a review of the workbook, then the latest changes remain in effect. However, if you review the workbook, Excel lets you independently examine each change and decide whether to make the change permanent or discard it. If you decide to discard, or *reject* a change, Excel restores the previous value for the cell and removes the change from the history log. If you decide to keep, or *accept* a change, it stays in the history log, but it's flagged as a confirmed edit. Excel will never ask you to accept or reject that change again.

To review the changes in a workbook, follow these steps:

1. **Choose Review→Changes→Track Changes→Accept or Reject Changes.**

   If you haven't saved the document yet, Excel prompts you to do so now. In Excel, you can't review changes until you've saved them to the workbook file.

   Once you've saved the document, the Select Changes to Accept or Reject dialog box appears (see Figure 25-15). This box lets you set filter conditions to indicate the types of changes you want to review. These filter conditions are the same as those that you find in the Highlight Changes dialog box.

2. **Adjust the When, Who, and Where lists to indicate the changes you want to review.**

   If you want to review every change that hasn't been reviewed yet, select "Not yet reviewed" from the When list, and don't make any other selections. (For more information about the When, Who, and Where lists, see page 734.)

Figure 25-15:
When you start the review process, you can choose the type of edits you want to examine by turning on the checkboxes and clicking the drop-down menus in the Select Changes to Accept or Reject dialog box. Usually, you'll choose to take a look at everything that hasn't already been confirmed or rejected, as shown in this example.

No matter what selection you make, Excel ignores changes that you have already accepted.

3. **Click OK to start the review process.**

   If Excel can't find any unconfirmed changes, it informs you that no changes are available and ends the review process. Otherwise, Excel begins scanning the document from your current position. It scans from left to right, and then top to bottom, examining one cell at a time. When it reaches the last cell, it starts over at the top of your worksheet to scan the remaining cells, and then continues its search with the other worksheets in your workbook.

   Whenever Excel finds a changed cell, it shows a message that offers you the opportunity to review the change, and then choose to either accept it or reject it, as shown in Figure 25-16.

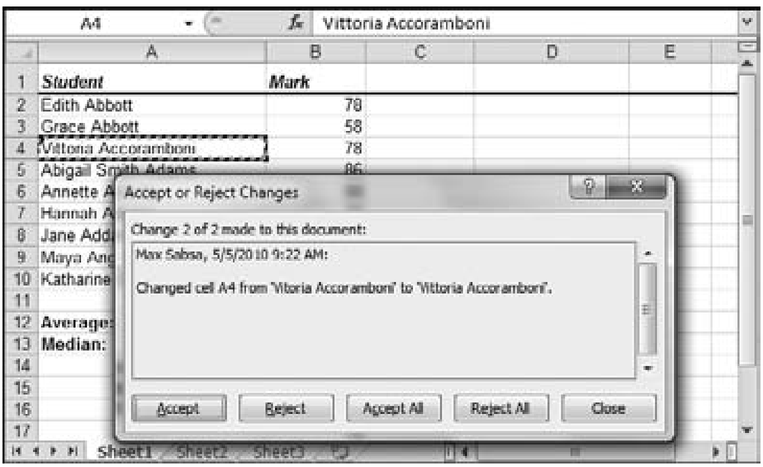

Figure 25-16:
Excel has found a change made to cell A4. The Accept or Reject Changes dialog box lets you choose to accept this correction–in this case, a minor correction of a misspelled name–or reject it and move on.

4. **Click Accept to confirm the edit, or Reject to discard the edit and remove the change from the log.**

   If you know that you want to reverse all the changes in a workbook, you can click Reject All to cancel everything in one fell swoop. Similarly, Accept All confirms all the outstanding changes. Be careful, though—if you use these buttons, you can't undo your action afterward.

   After you make your choice, Excel continues its search. When it can't find any more unconfirmed changes, the review process ends, and you return to your worksheet.

Note: Even after you've accepted a change, it remains in the change log until the time limit expires. This system differs dramatically from the change tracking you'll find in a program like Microsoft Word, which tracks changes only until you review them.

Merging Multiple Revisions into One Workbook

Change tracking works great when you need to send a workbook from one person to another. Each person can add her insights, and the original author can review all the changes before preparing a final version. Of course, life rarely works so smoothly. More commonly, several people need to work on the same workbook at the same time to get the document revised as quickly as possible. In this case, each reviewer ends up with a separate copy of the original workbook. When the revisions are finished, the original author needs to consolidate all the changes from each copy into one workbook file.

This situation probably seems like a guaranteed recipe for hours of manual Excel labor. However, Excel provides a special merge feature that neatly solves the problem of multiple reviewers. You can combine all the changes into one workbook. Unfortunately, this feature is one of the more specialized commands that don't appear on the ribbon. If you want to use it, you need to customize the Quick Access toolbar. Here's how:

1. **Right-click the Quick Access toolbar (in the upper left of the Excel window), and then choose Customize Quick Access Toolbar.**

   The Excel Options window opens, with the Customize section highlighted.

2. **In the "Choose commands from" list, select "Commands Not in the Ribbon."**

   This action shows the commands that Excel has hidden from sight.

3. **In the list, select Compare and Merge Workbooks, and then click the Add button.**

   This action moves the Compare and Merge Workbooks command to the list on the right, which details all the commands currently in the Quick Access toolbar.

4. **Click OK.**

Once you've added the Compare and Merge Workbooks command to the Quick Access toolbar, you're ready to go. Here's how to combine two spreadsheets and find the differences:

5. **Make sure you turn on change tracking (page 731) before you distribute your document.**

   The compare and merge feature uses the change history. If you want to compare more than one copy of a spreadsheet, and the changes were made without change tracking turned on, you're out of luck.

Tip: You can't merge a document if some of the changes have expired from the change history log. To avoid this problem, it's a good idea to make sure the change history tracks changes for longer than 30 days before you send the workbook out for review. Just select Review→Changes→Share Workbook, and then select the Advanced tab. In the "Keep change history for" setting, enter a really large number like 500 days. The only disadvantage to keeping change entries for a longer amount of time is that your workbook file becomes larger as it retains additional data.

---

6. **Make sure each file you want to merge has a different file name.**

   Excel can't open files with the same name at the same time. Usually, you'll give each file a new name based on the reviewer. For example, if your workbook is called expenses.xlsx, you might want to create copies of it with names like expenses_rr.xlsx (for someone whose initials are R.R.), expenses_mj.xlsx (for someone whose initials are M.J.), and so on.

7. **Put all the workbook copies you want to merge into the same folder on your computer.**

   By having all the workbook files in one location, you'll be able to merge all the files in one step.

8. **Open your master copy of the workbook.**

   Before continuing, you should also make sure none of the workbook copies are currently open. If they are, you can't merge them into your master workbook. You may also want to make a backup copy of your original workbook. That's because Excel automatically saves the workbook after the merge process is finished, and you can't undo the changes.

9. **In the Quick Access toolbar, click the Compare and Merge Workbooks button.**

   A standard file dialog box appears.

10. **Browse to the folder containing the files, and then select the copies you want to merge.**

    To select more than one file at a time, press the Ctrl key while you click the file name.

11. **When you have selected all the copies, click OK.**

Excel compares each copy and applies all the changes to the current workbook.

During the merge process, Excel may need to decide between two conflicting changes. If two people have modified the same cell, for instance, or if you modified a cell *after* sending the workbook out for review, you have a conflict. In this situation, Excel uses the most recent change. However, Excel carefully enters every change in the change history log. After you've merged all the workbooks, you can use the Accept or Reject Changes command to review all the changes, including cells where there were conflicts.

When you find a cell where multiple changes were made while reviewing a document, you see a message like the one shown in Figure 25-17. You can then choose which values you want to keep.

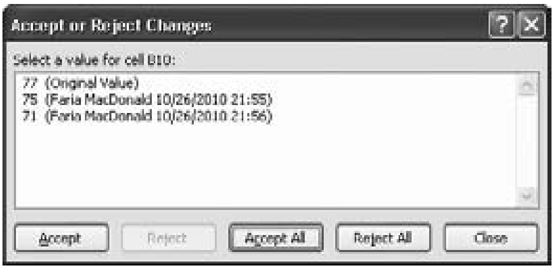

Figure 25-17: *Occasionally, you may find that more than one unconfirmed change has been made to the same cell. In this situation, Excel adds them all to the change log, and uses the most recent value in your worksheet. When you accept and reject changes, Excel shows a dialog box that shows all the changes, ordered from oldest to most recent. You must choose which change in the list "wins."*

Note: You can also use Excel's workbook merging tool to combine comments from several different versions of the same file. Excel automatically adds the comments into your master workbook when you perform the merge. Remember, in order for this plan to work, you need to have change tracking turned on in each file.

Sharing Your Workbook

Change tracking is really just one part of a larger feature known as *workbook sharing*. Workbook sharing makes it possible for more than one person to modify an Excel document at the same time. In order for this sharing to work, the workbook file needs to be somewhere that both people can access. (The most typical location is somewhere on a company network.)

Workbook sharing is a little risky (as you'll soon see), but it gives you some unique collaboration abilities that you wouldn't have otherwise. To get some perspective on how workbook sharing works, you need to understand what happens when two people fight over a workbook file that isn't shared. Read on.

Multiple Users Without Workbook Sharing

Ordinarily, only one person can open an Excel workbook file at a time. Excel enforces this restriction to prevent problems that can happen when different people try to make conflicting changes at the same time and someone's changes get lost.

Figure 25-18 shows what happens when you try to open an Excel workbook that someone else is already using. Excel warns you about the problem but gives you the chance to open a copy of the document that you can save with a new name. If you need to update the version that's currently in use, you can ask Excel to notify you when your collaborator closes it and the workbook becomes available again.

Figure 25-18:
Top: If you try to open a workbook that someone else is using, you receive the message shown here. You can click Read Only to work with a copy of the file, which you'll need to save with a new name. If you do so, Excel opens a copy of the file and lets you start working with it.

Bottom: Once the first person closes the original workbook, Excel alerts you with a new message. At this point, you can click Read-Write to open the file for editing (at which point no one else will have access to the file). Excel automatically closes the copy you were previously working with and applies all the changes you made in your copy to the current version.

Taking turns editing a file works well enough, but it has one significant problem: You never know when the file will become available. If someone opens a workbook, and then leaves for a two-week vacation without shutting off his computer, you'll be completely locked out of the workbook. Worst of all, you have no way to know when the workbook will finally be available.

And if you make changes to a copy while waiting for someone else (using the same file) to release the original workbook, more problems happen. The problem is that when you make your changes, you're actually working with a separate copy of the workbook file. If the other person makes any changes to the original workbook file, you'll find yourself in a sticky situation. Excel forces you either to save your copy as a separate file or discard all your hard work before letting you open the original copy (see Figure 25-19). You have only two solutions—discard your changes completely, or save the workbook with a new name, and then open the other person's version.

For all of these reasons, it usually makes more sense to use workbook sharing if you know more than one person needs to access or change the same workbook at the same time.

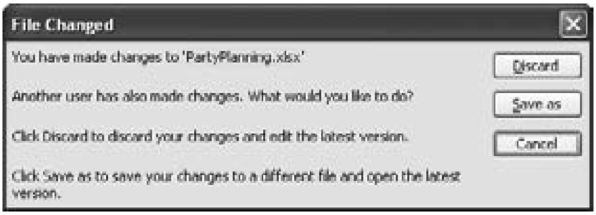

Figure 25-19:
While waiting for PartyPlanning.xlsx to become available, the person using this file made changes to the copy. Now that PartyPlanning.xlsx is available, Excel notices that the first person also changed the original PartyPlanning.xlsx file at the same time.

Turning On Workbook Sharing

It's quite easy to turn on workbook sharing. In fact, it takes only a couple of steps:

1. **Select Review→Changes→Share Workbook.**

   The Share Workbook dialog box appears.

2. **Turn on the "Allow changes by more than one user at the same time" checkbox.**

3. **Select the Advanced tab.**

   Figure 25-20 shows the Advanced tab in action.

Figure 25-20:
The Advanced tab lets you configure whether change tracking is used, when separate workbook copies are synchronized, and how Excel handles conflicting changes. The bottom two checkboxes ("Print settings" and "Filter settings") indicate that Excel stores some private information that it doesn't share with the document—namely, your personal printer settings and any view settings (like the position of the window at the current zoom magnification). These settings can't be synchronized because they may apply only to your computer.

4. **If you don't want to use change tracking, turn off the checkmark next to the "Don't keep change history" option.**

   When you share a workbook, Excel assumes you want to use change tracking. That way, if a conflict happens and one person overwrites another person's changes, you can still find the discarded change in the log if you need it. If you don't want Excel to save changes, take this opportunity to turn change tracking off.

5. **In the "Update changes" section, choose how often you want Excel to inform you about changes.**

   When workbook sharing is switched on, everyone editing the workbook has a separate copy of the workbook. The difference is that each person can save changes to the original workbook file. When you save your copy of the workbook, Excel examines the changes that your collaborators have saved, and then refreshes your copy of the workbook accordingly.

   Alternatively, you can tell Excel to save your workbook periodically to make sure it stays up to date. To switch on this behavior, change the "When file is saved" option to "Automatically every 15 minutes." (If you want to update more or less frequently, you can change the number of minutes.) Then, select the "Save my changes and see others' changes" option.

   Finally, you can choose to refresh your document periodically without actually saving your changes. Simply change the "When file is saved" option to "Automatically every 15 minutes," and select the "Just see other users' changes" option.

6. **If you wish, select one of the settings for "Conflicting changes between users" setting.**

   Here's your chance to tell Excel what to do if the changes from more than one person conflict with each other. Keeping the standard "Ask me which changes win" setting usually makes sense because Excel explains any conflicts and shows you the changes that were made earlier. If you choose "The changes being saved win," Excel automatically overwrites conflicting changes without even warning you that important data could be lost or alerting you that a conflict happened.

7. **Click OK.**

   Excel saves the workbook automatically. You can tell that it's shared because the word [Shared] appears in the title bar.

When you're working with a shared workbook, Excel keeps track of the number of people involved. To see who's working on the workbook at any given moment, select Review→Changes→Share Workbook. Figure 25-21 shows an example.

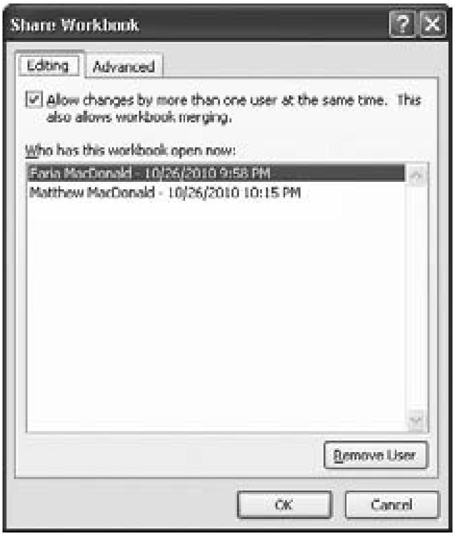

Figure 25-21:
In this example, two people have the workbook open at the same time. One of Excel's oddest features is the ability to stop people from saving a shared workbook if you select them in the list, and then click the Remove User button. This feature has no limits, which means that, in this example, Matthew can remove Faria and, similarly, Faria can remove Matthew if she gets to the setting first. When people are removed from the list, they can't save their changes to the shared workbook. Instead, Excel lets them save a copy only under a new file name. This removal process is temporary, however. If Matthew removes Faria, there's no reason that Faria can't open the workbook after Matthew closes it. This anything-goes attitude can make life a bit chaotic if you let multiple people edit the same workbook at once.

Note: Remember, once you turn on workbook sharing, Excel completely turns off certain features (like inserting charts or pictures). For the full list, see the box on page 732.

Workbook Sharing in Action

When workbook sharing is switched on, there's no limit to how many people can open the same workbook at once. However, each person actually has a separate copy of the workbook. Excel automatically synchronizes your copy when you save changes (or at fixed intervals, depending on the "Update changes" option described on page 744).

This approach works well when people change different portions of the workbook. However, it's entirely possible that two people could choose to edit the *same* cell at once. Sooner or later, both people will try to save their changes.

Imagine that Faria changes cell A1 at the same time that Ricardo modifies A1. Faria saves her changes and doesn't experience any problem. Shortly after, Ricardo tries to save his changes. At this point, Excel warns Ricardo that his changes conflict with the changes someone else already made. Ricardo now needs to decide whether to overwrite Faria's changes or discard his own. Figure 25-22 shows an example of the message you may see in this situation.

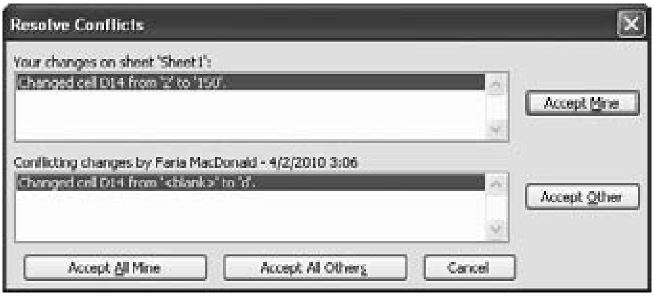

Figure 25-22:
In this example, Ricardo has just tried to save his copy of a shared workbook. However, one of his changes conflicts with a change Faria made. Ricardo can choose Accept Mine to discard Faria's change or click Accept Other to leave the workbook as it is.

Here's a clear problem—the last person to make changes gets to make decisions about someone else's work. There's no way to avoid this limitation of shared workbooks. Excel doesn't provide any way to stop one person from obliterating a co-worker's edits.

In fact, even if changes don't directly conflict, simultaneous edits can still clash. Imagine you have two cells (A10 and B10) that don't quite match up the way they should. (A10 may have the price of a product and B10 may have a discounted price, which—oddly enough—is higher.) One person tries to fix the problem by modifying A10. At the same time, another person tries to resolve the same problem by modifying B10. The end result? Both A10 and B10 get changed, Excel doesn't complain, but they still don't match up. These kinds of things make workbook sharing a high-stakes, life-in-the-fast-lane approach.

Tip: Workbook sharing works best when multiple people need to read a workbook and only one person needs to edit it. In this situation, the chance of conflicting changes is low, and people who just want to print the data aren't inconvenienced while it's in use. As a rule of thumb, try to separate the workbook so that different people will update the data on different worksheets, which won't cause a conflict.

FREQUENTLY ASKED QUESTIONS

The Best Way to Share Workbooks

The Excel Web Application has a collaborative editing feature. Does this replace workbook sharing?

In the next chapter, you'll learn about a whole other way to allow workbook sharing—using collaborative edits in the Excel Web App.

In some respects, the Excel Web App takes a simpler approach than workbook sharing. Edits synchronize almost immediately, while ordinary workbook sharing waits until you save the document or takes place just every few minutes (using the "Update changes" option described on page 744). However, the Excel Web App isn't as good at handling conflicts when they happen. If by some chance you do change the same cell at roughly the same time as someone else, one person's work will be overwritten, without any warning or notification. (This sounds scary, but it isn't a big concern for many people. Microsoft research shows that in most collaborative situations, different people work with different data and edit different cells.)

Truthfully, the choice between ordinary workbook sharing and the Excel Web App depends on a number of factors. If you're working on a home or company network, and if you want to prevent any chance of accidental overwriting, ordinary workbook sharing is the safest choice. But if you want to perform collaborative editing with people anywhere, you aren't worried about clashing edits, and you don't mind that it has significantly fewer features than the desktop version of Excel, the Excel Web App may be a better choice.

If you're somewhere in between, the decision probably depends on where you want to perform your work—in the full-featured Excel window, or in a slick but less powerful Web application that runs anywhere, even on computers that don't have an installed copy of Excel. For more information about the Excel Web App, keep reading.

CHAPTER
26

Using Excel on the Web

Almost as long as the Internet has existed, Excel fans have been trying to put their work online. In Chapter 1, you learned about Excel's Save as HTML feature, which is the simplest way to get a workbook ready for web viewing. But as you already know, just looking at a spreadsheet misses the real fun. Life is a lot more interesting when you can *edit* your work (for example, adding new details to a company spreadsheet) or *manipulate* a spreadsheet (for example, rearranging a pivot table or changing the filter settings in a table to home in on the most important information). Unfortunately, once you've converted a workbook to HTML, you're left with a fixed, inert web page, and you can't perform either of these tasks.

For most of Excel's history, Excel fans accepted these limitations as a simple fact. Excel was for spreadsheet authoring, and the Web was for viewing some of that information, some of the time. But when slick web applications like Gmail and Google Docs began to appear, people realized that there just might be a way to take some of the tasks that normally happen inside the Excel window and move them online.

Office 2010 delivers with Web-based versions of its core programs, including Word, Excel, and PowerPoint. These web apps let web surfers edit documents right in their browsers, even if they don't have any Office software installed. Each web app differs in its exact capabilities (but, obviously, none matches the extensive functional of the desktop software). The Excel Web Application is particularly powerful. Not only does it let you create tables and build formulas, but it also adds a new live collaboration feature that lets a team of people change the same spreadsheet at the same time. It's no exaggeration to say that the Excel Web App is the biggest innovation in Excel 2010.

FREQUENTLY ASKED QUESTIONS

Browser Compatibility

Does the Excel Web App need Internet Explorer?

Here's the unexpectedly good news: the Excel Web App works with most modern browsers. It officially supports the following:

- Internet Explorer 7 and 8
- Firefox 3.5 on Windows, Mac, and Linux
- Safari 4 on Mac

But this isn't the final story. If you're using a browser that's not on this list, it could quite possibly still work. So give it a try.

Occasionally, a feature might not work in all browsers. One example is Firefox, which doesn't let Web pages use the clipboard. When you try to use the copy feature in the Excel Web App (for example, by choosing Home→Clipboard→Copy from the ribbon), you'll get an informative message that explains the problem.

Lastly, the Office Web Apps may get just a get a bit better if you have Silverlight installed. Silverlight is a free, lightweight browser plug-in from Microsoft (see *http://silverlight.net*) that lets developers create rich applications that run in the Web browser, from animated games to streaming video players. (For example, you might have used it to watch live events in the recent 2010 Winter Olympics.) However, Silverlight is entirely optional for all of the Office Web Applications. If you do have Silverlight installed, you may notice that the Word Web App has clearer, more detailed text, and the PowerPoint Web App has smoother animations. The Excel Web App remains essentially the same.

Putting Your Files Online

The idea of running an Office program in a web browser seems simple enough. But working on the Web is different from putting spreadsheets together on your own computer, and it raises quite a few new questions. Such as, where do I put my files? What sort of browser should I run? And what happens if several people try to change a spreadsheet at the same time?

First, before you can do any editing with the Excel Web App, you need to get your workbook online. More specifically, you need to transfer your document to a SharePoint server that has the Excel Web App software installed on it. These are the only files that the Excel Web App can access. It can't open the files on your local hard drive, the files you've emailed to a friend, or the files you've posted on your personal website.

Note: SharePoint is a Microsoft product used by big businesses to help people collaborate and share information. You don't need to install SharePoint to use its features. Instead, you connect to a server computer that's running SharePoint, in much the same way that you connect to a web server to buy a book on Amazon or run a Google search.

At first glance, this seems like a Very Big Problem. After all, plenty of Excel fans want to use the Excel Web App, and most of them don't have a spare SharePoint server kicking around. Fortunately, there's an easy solution: you can use Microsoft's SharePoint servers. In fact, you can do it for free using a Web-based service called SkyDrive.

Introducing SkyDrive

SkyDrive began as a handy way to put big files online—for example, so you could back up your computer or share a document with your friends. But SkyDrive's support for the Office Web Applications means that you can upload an Excel spreadsheet, and then start editing it in your browser with the Excel Web App. In this way, SkyDrive actually does two things. First, it stores the Excel files you upload. Second, it hosts the Excel Web App, which lets you (and other people) edit your files.

To use SkyDrive, you need a Windows Live ID. If you use the Microsoft's Hotmail service for web-based emailing, you already have one. If not, you can sign up and get a Hotmail address at the same time, or you can sign up using an email address that you already have. Either way, your starting point is *http://home.live.com*. Click the Sign Up button, fill in your personal information, and you'll be set up in seconds.

Once you have a Windows Live ID, you can use SkyDrive. To sign in directly, point your web browser to *http://skydrive.live.com*. Or, if you're logged into another Windows Live service (like Hotmail), look at the bar of links at the top of the page, and then click More→SkyDrive (Figure 26-1). Figure 26-2 shows you the starting point of a new SkyDrive account.

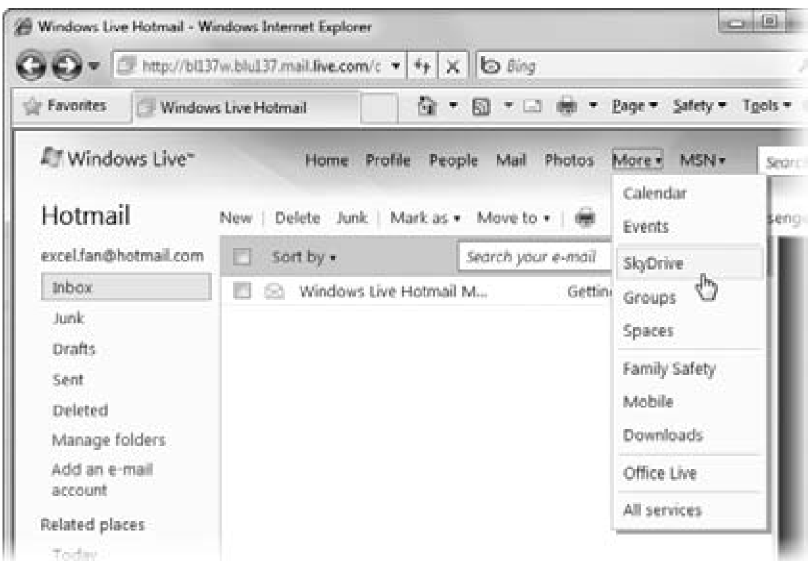

Figure 26-1: *Here's how to get to SkyDrive if you're already signed in to another Windows Live service (like Hotmail).*

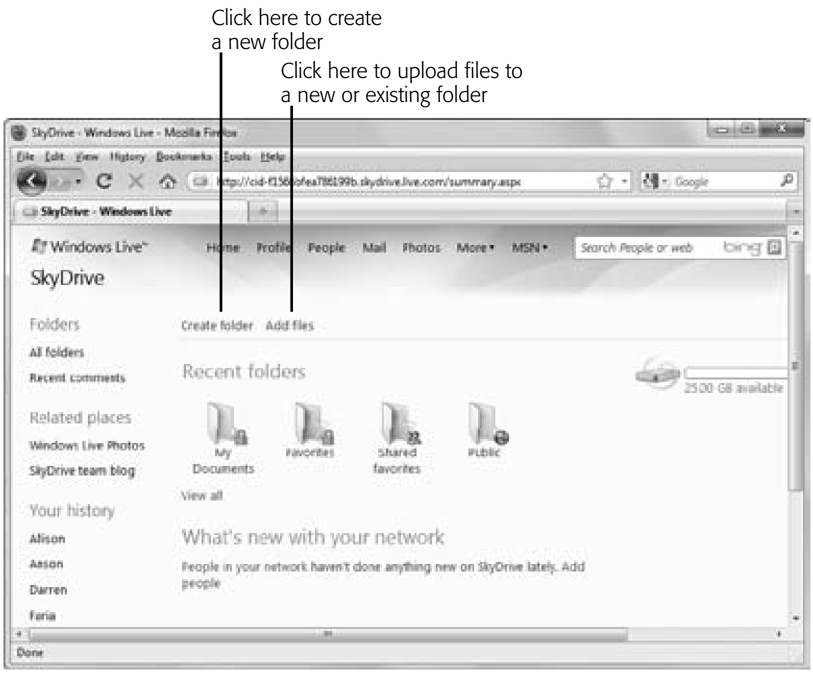

Figure 26-2: *This new account has four SkyDrive folders. To start browsing, click one of folders. Or, you can use the "Create folder" and "Add files" links to make new folders, and then upload files from your computer–two tasks you'll perform in the coming pages.*

A new SkyDrive account starts out with four folders. Two of them ("Favorites" and "Shared favorites") are for storing favorites (Internet bookmarks), so you can access them on different computers. To use this feature, you need the Windows Live Toolbar. The other two are reserved for your documents. The first, "My Documents," holds personal documents that you aren't sharing with anyone else. The second, "Public," holds files that any other web surfer can access, even without a Windows Live ID.

Of course, you can create as many new folders as you need. When you do, you choose the exact folder permissions. For example, you can make a folder completely public, or limit it to just a few people in your contact list. You'll see how on page 755.

Note: At the time of this writing, SkyDrive gives every user the same amount of space: a whopping 25 GB (which is equivalent to 4,000 songs in MP3 format, five movie-length DVDs, or well over 100,000 Excel workbooks). And if you run out of space, you can simply open a second account.

Your Web Workflow

Before you start uploading Excel spreadsheets, you need to give a bit more thought to your *workflow*—that is, the way you want to manage your spreadsheet files and their revisions. Two common approaches incorporate the Excel Web App into your Excel workflow:

- **Use the Web to share a *copy* of your work.** This approach makes sense if you're the owner and chief maintainer of a spreadsheet that has some important information that other people want to see. You periodically upload a copy of your work, but you keep the main version and revise it on your own computer. In this way of doing things, you're really just using the Excel Web App's viewing capabilities. Other people may edit the copy with the Excel Web App, but there's really no point, as you'll just overwrite it when you upload your next copy.
- **Let your document live online.** In this case, the document that you upload *is* the central version. You (and other people) can edit this document using the Excel Web App or the desktop version of Excel, but changes are always saved to the online file. The creators of the Excel Web App think that this scenario is the most interesting one, because it makes it easier to share information and collaborate on a spreadsheet without needing to designate one person as the official spreadsheet owner. Often, marketing types describe this as putting your documents "in the cloud."

You might want to use both strategies for different files. For example, you could use the first strategy to share copies of a company financial plan, so that no one else can fudge the numbers. But if you want to compile a company expense summary, you use the second approach to share your workbook with all the department heads, so they can enter their numbers in the same workbook.

Note: Online workbooks seem like a perfect place to use document protection (Chapter 24), the Excel feature that lets you control who can do what with a workbook. Document protection is fine-grained, which means you can give different people access to different worksheets or different parts of a worksheet, and you can limit the types of changes they're allowed to make. Unfortunately, the Excel Web App doesn't support document protection in this release. That means the only control you have is to set who can and can't access a workbook, and everyone on the "can" list is allowed to make unlimited changes.

Uploading a File to SkyDrive (Using Your Browser)

Now that you have a bit of background about the Excel Web App and SkyDrive, you're ready to upload a file. You can transfer a file to SkyDrive in one of two ways: using your browser (in which case you can upload virtually any type of file) or using Excel (in which case you can upload only workbooks).

Usually, you'll find it a bit quicker to upload your workbooks from inside Excel. However, some tasks require your web browser, like creating a new folder and setting folder permissions. In this chapter, you'll learn to use SkyDrive in both ways.

The following steps show you how to upload a file to SkyDrive with your browser:

1. **Point your web browser (any one will do) to *http://skydrive.live.com*.**

   Log in with your Windows Live ID email address and password.

2. **Click "Add files."**

   SkyDrive shows you a list of folders in your SkyDrive account (Figure 26-3).

Figure 26-3: *Before you can upload a file, you need to pick the folder where it should go (or choose "New folder" to create a new one).*

3. **Choose "New folder" to create a new folder. Or, if you want to use a folder that already exists, click the folder, and then skip to step 6.**

   If you choose to create a new folder, the "Create a folder page" appears (step 4). If you choose an existing folder, the upload page appears (step 6).

4. **Fill in the information for your new folder.**

   To create your new folder, you need to supply two pieces of information: the folder name (for example, "Excel Files") and your "Share with" setting.

   Your "Share with" setting determines who else can access the files you put in this folder. Choose "Just me" to keep your files to yourself, or choose "Everyone" to make your folder public, which throws the doors open to every web surfer, even those who don't have a Windows Live ID. Your other two options are somewhere in between: use "My network" to share the files with everyone that you've added to your Windows Live network, or "Select people" to pick and choose (Figure 26-4).

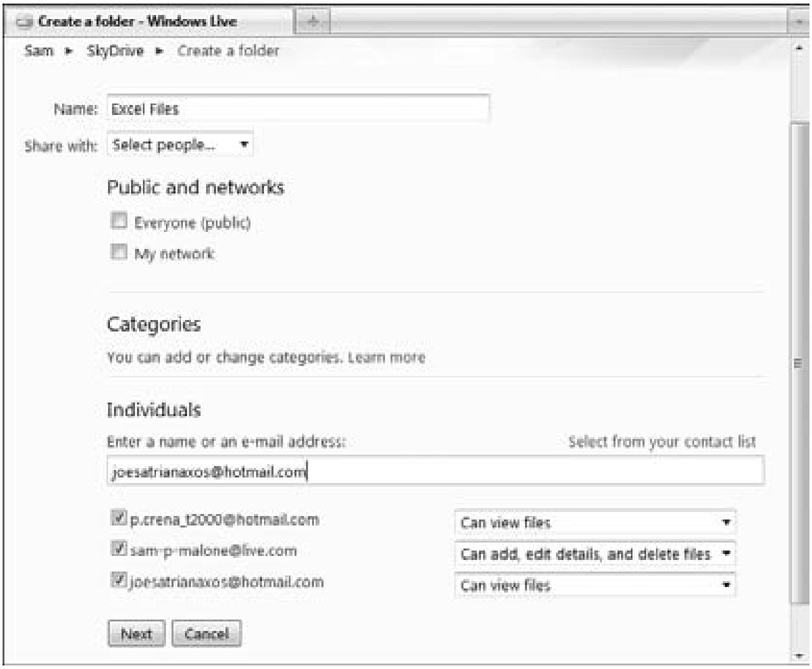

Figure 26-4: *Here, a new "Excel Files" folder is being created. This folder is accessible to the SkyDrive account owner and a select group of three people, who are identified by their Windows Live ID email addresses. To add a new person to the list, type the email address in the text box shown here and press the Enter key. Once added, you can give each person one of two permission levels. You can give a person full control to add, edit, and delete files in this folder (like sam-p-malone), or restrict a person to viewing files only (like joesatrianaxos).*

Tip: You can create a folder without uploading a file by clicking the "Create folder" link on the main SkyDrive page. This takes you to the same figure that's shown in Figure 26-4.

5. **Click Next to create your new folder.**

   This brings you to the upload page, where you'll see several text boxes, one for each file you want to upload at a time (Figure 26-5).

6. **To add a file, click Browse, select the file on your hard drive, and then click Open.**

   It's up to you whether you upload a single file at a time or up to five at once. But if you have several, clicking the Browse button, and then picking each one individually can get tiring. In Internet Explorer, you can use a shortcut: install a handy plug-in that lets you drag files out of a folder on your computer, and then drop them onto a big box on the upload page. This makes it quicker and easier to upload a whole whack of files at once. To give it a whirl, click the "Install the upload tool" link, which appears above the list of text boxes with the files you want to upload.

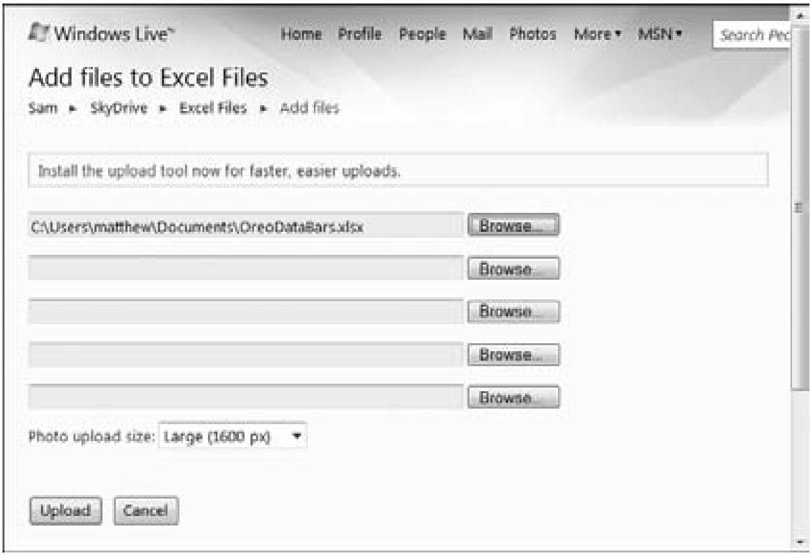

Figure 26-5: *Here, a single file (OreoDataBars.xlsx) has been chosen for uploading. To actually start the transfer, you need to click the Upload button at the bottom of the page.*

7. **Once you've picked all the files you want to upload, click Upload.**

   Depending on the size of the file, there may be a delay while it's transferred to the SkyDrive servers. When the operation is complete, your browser shows you a list of all the files in the upload folder, including the one you've just added.

Uploading a Workbook to SkyDrive (Using Excel)

Excel has a nifty SkyDrive integration feature that lets you upload workbooks without leaving the comfort of the Excel workbook. Once you've created a SkyDrive folder with the right permissions (as shown in Figure 26-4), you'll probably do most of your uploading from inside Excel. Here's how:

1. **Open the workbook you want to share.**

   The Excel Web App doesn't support all of the features of the Excel desktop application, and it supports others only partially. It's worth reviewing these limitations (which are described on page 765) before you upload a new workbook.

2. **Choose File→Save & Send.**

   This brings you to Excel's backstage view.

3. **In the Save & Send list, choose Save to Web. Then, click the Sign In button.**

   A dialog box appears where you can enter your Windows Live ID.

4. **Enter your Windows Live ID email and password, and then click OK.**

   If you don't want to go through this process every time you upload a new Excel workbook, put a checkmark in the "Sign me in automatically" box. This tells Windows to remember your Windows Live ID.

Note: Even if you don't choose "Remember my credentials," Excel keeps you connected for some time—usually until you close Excel or let several hours pass without communicating with the SkyDrive server.

Once you've signed in, you'll see your SkyDrive folders appear on the right side of the window (Figure 26-6).

Figure 26-6: *This SkyDrive account has three folders. Under each folder is a description of its permissions—in other words, who can access the files that you place inside.*

5. **Click the folder where you want to upload your files, and then click the Save As button.**

   At this point, Excel exits from backstage view and begins to contact the SkyDrive server, which takes a bit of time. (During this process, you'll see a message in Excel's status bar that explains that it's busy contacting the server.) Once Excel has made its connection, the Save As dialog box appears (Figure 26-7).

If you want to create a new folder, you can, but it's not quite as convenient as it seems. When you click the New Folder button Excel launches your browser, which prompts you to log on to SkyDrive, and *then* shows the folder creation page (Figure 26-4). When you're finished, you need to click the Refresh button in Excel to make sure your newly created folder appears in the list.

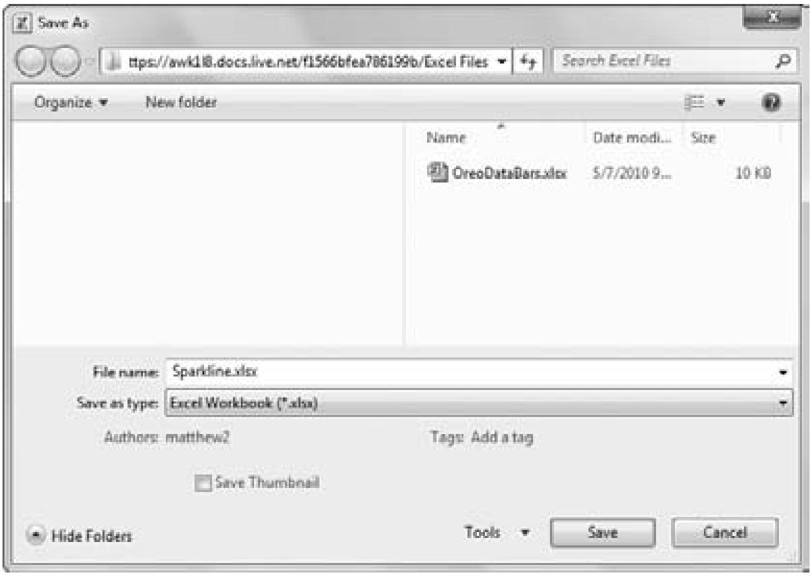

Figure 26-7: *This Save As dialog box looks a little different from the one you see when saving a workbook on your computer, because it's connected to the SkyDrive server. First, the folder pane is a blank box, because you can't switch between folders. Second, the full file path in the box at the top of the window includes a Web server name (like* https://wu5wsu.docs.live.net*) and a cryptic code (like 5a66c3bd5973893f). The code uniquely identifies your SkyDrive account. Finally, the file box on the right shows the files and subfolders in the current folder. You can double-click a subfolder here to move into it.*

6. **Make sure the name in the "File name" box is the file name you want to use for your uploaded file, and then click Save.**

   You should also make sure that you're using the right format. The Excel Web App uses the modern XML-based Excel formats that were introduced in Excel 2007 (and described on page 39), including .xlsx, .xlsm, and .xlsb. If you upload an older .xls file, you'll be able to view it with the Excel Web App, but not edit it.

   Once you click Save, Excel uploads your workbook. This typically takes less time than establishing the initial connection in step 5.

From this point on, the copy you're working with in Excel is the SkyDrive copy. If you change your workbook and save it again, you'll end up updating the online copy. (If this isn't what you want, close the current workbook, and then open the copy that's on your computer.)

Tip: Before you make any changes, make sure it's clear which is the master version of your spreadsheet: the local copy on your computer, or the online copy on SkyDrive.

Finally, you have to close the workbook before you (or anyone else) can edit it in the Excel Web App. This step is important because the desktop version of Excel takes exclusive control over the files it opens, locking out the Excel Web App.

Uploading a Workbook to a SharePoint Server

The examples in this chapter use SkyDrive. After all, it's free and works for everyone. However, if you're using an instance of the Excel Web App that's installed somewhere else—for example, a company SharePoint server—your experience remains essentially the same. The only difference is that you upload your work to a different place before you get started.

Here's how it works:

1. **Open the workbook you want to share.**

   Check the limitations on page

2. **Choose File→Save & Send.**

   This brings you to Excel's backstage view.

3. **In the Save & Send list, click Save to SharePoint, and then "Browse for a location."**

   If you've done this before, you may see the SharePoint location in the Recent Locations list. If so, you can click if there instead of choosing "Browse for a location," and save some time.

4. **Click the big Save As button.**

   Excel pops open the Save As dialog box. Now it's up you to find your company's SharePoint server. One common approach is to scroll down the list of folders until you find the Network group. If you know the SharePoint server by its computer name, you may find it here. Another option is to use a network shortcut that a computer administrator has created for you on your computer. Or, a network administrator may have given you specific instructions about what to type into the box at the top of the Save As dialog box to jump straight to the SharePoint server you need.

5. **Once you've found your SharePoint server, you'll need to browse to the right folder, and then click Save to transfer your workbook.**

   Remember, the original version of your workbook still exists on your computer. Depending on the workflow you're using (page 770), you may want to change the local copy or the online copy, but not both.

Viewing a File in Your SkyDrive Account

Once you upload a file to SkyDrive, it remains there, permanently, unless you choose to remove it.

To do something with your uploaded file, you need to leave the comfort of Excel. Launch your web browser, head to *http://skydrive.live.com*, and log on (if you haven't already). Then, browse your way to the file by clicking the right folder, and then click the file. At this point, a new page appears, with a variety of options.

If you're viewing a non-Office file—say, a set of vacation pictures or a ZIP file with your super-secret banana fudge recipes—you see SkyDrive's standard Download link. Because you're the owner of this SkyDrive account, you also get a few options that other people don't, like links for deleting, renaming, moving, and copying your file. (You can optionally give this ability to other people, as shown in Figure 26-4.)

Life gets much more interesting if you're looking at one of the SkyDrive-supported Office files, like a Word document, PowerPoint presentation, or Excel spreadsheet, in which case you see an additional message (Figure 26-8).

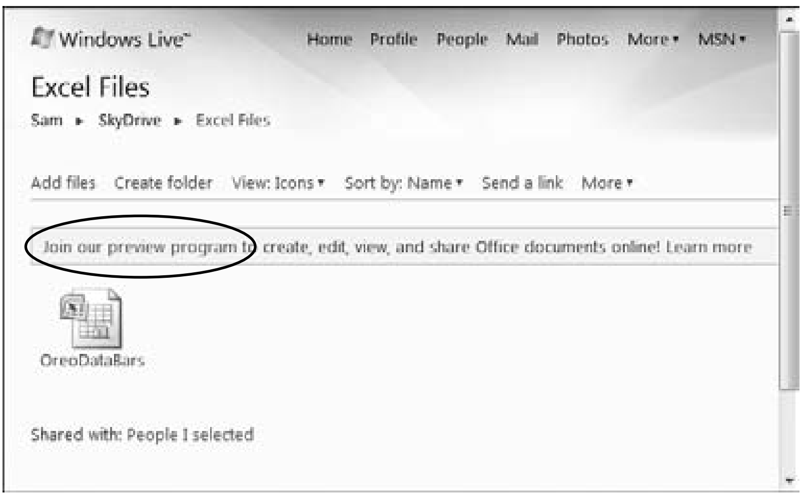

Figure 26-8: *At the time of this writing, the Excel Web App is still considered to be a preview feature on SkyDrive. The first time you view an uploaded Excel file, you'll see a link (circled here) that invites you to join the preview program. Click that link, and then click Accept to agree to Microsoft's licensing terms (which means you promise not to use the Excel Web App for nefarious purposes). Now, the next time you view an Excel file, you'll see the options shown in Figure 26-9.*

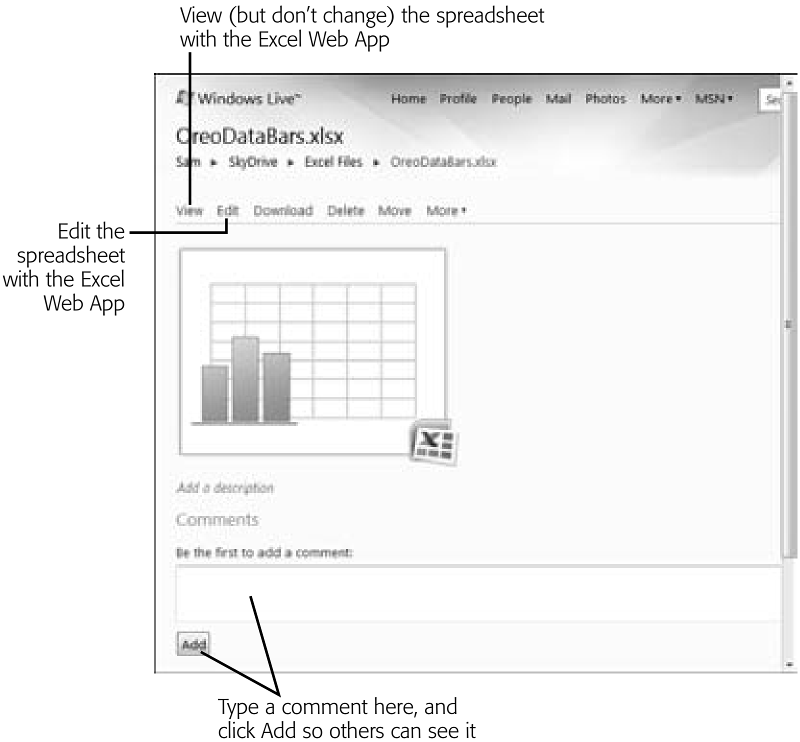

Figure 26-9: *When you click a file in SkyDrive (like OreoDataBars.xlsx), SkyDrive gives you two extra links: View and Edit. Both of these launch the Excel Web App. Underneath is a comment section where people can add their feedback.*

Tip: Being harassed by over-eager commenters? Although SkyDrive currently doesn't let you turn commenting off, you can delete any comment by clicking the "Delete" link next to it, assuming you're the SkyDrive account owner.

The View link lets you start the Excel Web App in read-only mode. You can look at your workbook, move from worksheet to worksheet, and scroll up and down the grid. However, you can't make any changes. Using view mode gives you two key advantages. First, you can open files that are locked for exclusive editing by the desktop version of Excel (page 770). Second, you avoid making accidental edits. This is a particular problem because the Excel Web App saves changes automatically, so you might not even realize that your careless keystrokes have spread beyond your browser session until it's too late.

The Edit link lets you run the Excel Web App in all its glory. You'll start your tour on page 764. But first, it's worth considering how you can get to an Excel file in someone else's SkyDrive account.

Viewing a File in Someone Else's SkyDrive Account

It's easy enough for you to view and download one of your own SkyDrive files. All you need to do is log on to your SkyDrive account, and then browse to the right folder. But what if you want a file from someone other people's SkyDrive accounts? How do you point your web browser to their files?

There are two basic techniques. If you want to view the files from someone in your contact list (for example, someone you send messages to with Messenger or someone you email regularly), it's easy enough, if a bit tedious. Here's how to do it:

1. **When you're logged on to a Windows Live service (like SkyDrive or Hotmail), look for the list of links at the top of the page. Click People.**

   A long, alphabetical list with all of your contacts appears.

2. **Find the person's email address in your contact list, and then click it.**

   An information page appears for the selected contact. This page lets you send a quick message to the person (if your contact is online) or an email (if your contact isn't online). The page also lists the email messages you've recently exchanged with your contact.

3. **Look for the "View profile" link, and then click it.**

   The exact name of the link depends on the person's name. If your contact is named Joe, it's "View Joe's Profile."

4. **In the column of links on the left, look for the SkyDrive link, and then click it.**

   Congratulations, you've finally arrived. Now you can click your way to the right folder and file, as though you were in your own SkyDrive account. If it's an Excel workbook, you can download, view, or edit the file (as you'll see shortly), but you aren't allowed move, delete, or upload anything.

Note: To view or edit an Excel document, you need to be part of the Excel Web App preview program, as described on page 760. However, SkyDrive doesn't prompt you to sign up if you simply view someone else's Excel file. Instead, you need to upload a spreadsheet of your own, view it (see Figure 26-8), and join the program before you'll be able to edit anyone else's Excel files.

If this approach takes too long, or if you want to access the files of someone who isn't in your contact list, you can go directly to the file detail page (like the one shown in Figure 26-9), provided you have the right link. This is the link shown on the right side of the file detail page, under the "Information" heading, in the "Web address" box. Figure 26-10 highlights what you're looking for.

This process requires the document owner's participation. For example, if you have a SkyDrive file you want to share with other people, you have to take the link to that file, paste it into an email message, and send that email message to everyone who needs to access the file.

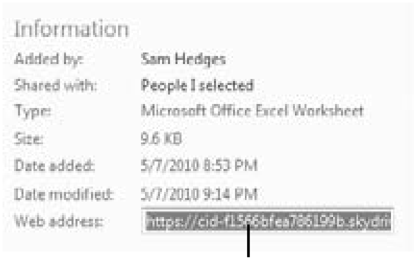

Click here and copy this link

Figure 26-10:
The link to this file incorporates a unique code that identifies your SkyDrive account, and the full path to the file. It's generally something hideously ugly, like https://cid-f1566bfea786199b.skydrive.live.com/self.aspx/Excel%20Files/OreoDataBars.xlsx. *To copy it, click the box, and then press Ctrl+C. You can then paste it into an email.*

Tip: If you need a nicer, shorter URL—one that someone can actually write down on a napkin and type into a browser, you can use a URL redirection service (like *www.tinyurl.com*). This service takes your long, ugly URLs and issues you a more manageable, shortened URL, like *http://tinyurl.com/yf82cbc*. When you type this URL into a web browser, it sends the browser off to the URL redirection service, which then automatically redirects the browser to the web location you really want.

Remember, the SkyDrive account owner chooses who can access a particular folder. If the file is in a public folder, everyone can access it, including people who don't have a Windows Live ID. But if the folder is restricted to specific people, you'll need to log in with your Windows Live ID before you can view the folder's contents or access any of the files inside.

GEM IN THE ROUGH

Creating a Workbook in SkyDrive

You've seen how to create a workbook in Excel, upload it to SkyDrive, and then access it there. This process takes a few steps, but it makes good sense. By creating your workbook in the desktop version of Excel, you get the chance to use some features that you can't access in the Excel Web App. For example, you can add charts, sparklines, and conditional formatting in desktop Excel, and then view them in the Excel Web App; even though you can't actually *create* any of them in the Excel Web App. (Page 770 has the full story.)

However, sometimes you might want to create a workbook for SkyDrive without using the desktop version of Excel at all. In fact, you might not even have desktop Excel installed on the computer where you're currently working. In this situation, you can create a blank workbook in SkyDrive by following a few simple steps:

1. Log on to SkyDrive.
2. Browse to the folder where you want to add the workbook.
3. Choose New→Microsoft Excel Workbook.
4. Fill in a file name, and then click Create.

Using the Excel Web App

You already know how to get started with the Excel Web App. Using a web browser, you surf to the SkyDrive location where the file is stored, select it, and then click Edit. Once you do, the Excel Web App starts and loads the workbook into a browser window (Figure 26-11).

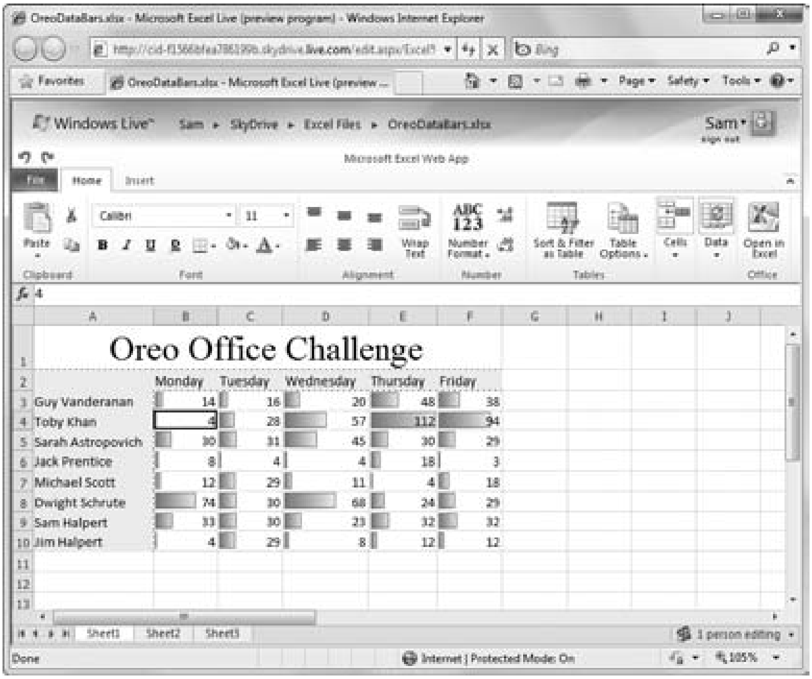

Figure 26-11: *The Excel Web App lets Web surfers edit spreadsheets directly in their browsers, even if they don't have Excel installed. Visually, the Excel Web App looks the desktop version of Excel. Functionally, it behaves like the desktop version (albeit with a reduced set of features). And its speed is surprisingly snappy, even when several people are collaborating on the same worksheet.*

Tip: At the time of this writing, the Excel Web App could be a bit quirky when loading a spreadsheet for the first time. If you click View or Edit and the web page doesn't appear, try refreshing the page (in Internet Explorer, that means right-clicking the page and choosing Refresh). This step usually solves any problems.

One of the first things you'll notice is that the Excel Web App looks strikingly similar to the desktop version of Excel, but with a drastically cut-down ribbon that features just two tabs and a File button. This resemblance is part of key feature that the creators of the Excel Web App call *visual fidelity*. Thanks to visual fidelity, what you see on your worksheet always looks the same, whether you're using the full desktop version of Excel or the Excel Web App. In fact, even if the Excel Web App doesn't fully support a feature, it can usually display it correctly. For example, the Excel Web App doesn't let you create charts. But if you open a workbook that already has a chart in it, that chart will still appear, and it will usually look the same as it does in desktop Excel. Similarly, the Excel Web App doesn't allow you to create data bars, but if you open a spreadsheet that already has them in place (like the OreoDataBars.xlsx file shown in Figure 26-11), the data bars appear, with all the right formatting. They even resize themselves appropriately if you change the underlying cell values.

Supported Features

When seasoned Excel veterans look at the Excel Web App, the first thing they want to know is which features work and which ones don't. Here's a quick rundown of what the Excel Web App can do:

- **Cell editing.** You can move around the grid of cells, skip from one worksheet to another, and enter text, numbers, and dates. You can also resize columns and rows. However, you can't add, delete, or rearrange the order of your worksheets.
- **Undo and Redo.** The world's two most useful buttons sit just above the File button. As with desktop Excel, the Excel Web App maintains a detailed Undo history, letting you step back through multiple changes, one change at a time. The only limitation is that you can't find out what action the Undo button will reverse until you click it. (The desktop version of Excel is nicer—when you hover over the Undo button, it shows a tooltip that describes the previous action.)

Note: The Undo and Redo features have a significant limitation. If more than one person is editing a workbook at once (page 768), the Excel Web App switches off both features, leaving you without a safety net.

- **Formula editing.** Just like in desktop Excel, the Excel Web App supports point-and-click formula creation (page 245), where you assemble a formula by clicking the cells you want to reference. The Excel Web App also supports Formula AutoComplete, which pops up possible function names as you type (Figure 26-12). You can also use named ranges in your functions, but the Excel Web App doesn't let you review, create, or change named ranges.

Note: The formula results that the Excel Web App calculates match those that you get in the desktop version of Excel. If you know a bit about web-page design, this might surprise you. After all, every browser is different, as are Macs and PCs, so it seems logical that what gets calculated on one system might not match the result in another. But the Excel Web App doesn't use your web browser to perform its calculations. Instead, it talks to the SharePoint server computer and asks it to do the work. This design also has a side effect: With complex formulas and a powerful server, the Excel Web App might get the answer faster than desktop Excel.

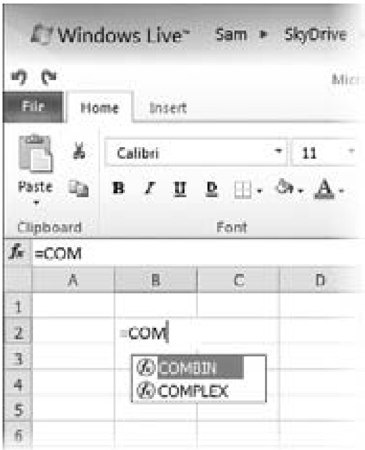

Figure 26-12:
Here, a drop-down list of suggestions makes it easier to find the right function. The Formula AutoComplete in the Excel Web App isn't quite as powerful as desktop Excel, as it can't tell you anything about the arguments a function uses.

- **Cell selection.** You can drag to select groups of cells, just like in desktop Excel. The only limitation is that you can't select noncontiguous groups of cells (page 99), a little-used feature that lets you select different blocks of cells on different parts of the spreadsheet at the same time.
- **Copy and paste.** Firefox doesn't support it at all. In Internet Explorer, you'll get a warning message the first time you perform a copy, asking you if you want to let the Excel Web App use the Windows clipboard. Assuming you agree, you'll find that cutting and copying works essentially the same in the Excel Web App as in the desktop version of Excel. However, the Paste feature is more limited, giving you just the four most common pasting options: Paste, Paste Formulas, Paste Values, and Paste Formatting. (Page 107 explains the difference.) As with desktop Excel, when you paste a formula with relative references, it automatically adjusts itself to point to the appropriate cells (page 252).
- **Inserting and deleting rows and columns.** You can use the Home→Cells→Insert and Home→Cells→Delete buttons to add and remove entire cells and columns, just as you do in desktop Excel.
- **Formatting.** The Home toolbar has all the buttons you need to adjust colors, fonts, and cell borders (in the Font section), text alignment and wrapping (in the Alignment section), and number formats (in the Number Format section). The Excel Web App even supports *live preview*, which means that when you hover over a new formatting option you'll see it applied to your worksheet, before you go ahead and make the actual change. However, the Excel Web App doesn't support applying or editing styles and themes.
- **Find.** The Excel Web App provides a stripped-down version of Excel's find feature (choose Home→Data→Find) that simply lets you scan cells forward or backward looking for a specific value.

- **Tables.** The Insert tab has just two buttons, and one of them (Insert→Tables→Table) lets you create a table, just as you did in Chapter 14. Figure 26-13 shows an example.
- **Hyperlinks.** The second button on the Insert tab (Insert→Links→Hyperlink) lets you place a link in a cell that points to another web location. You can choose the text to display and the URL.

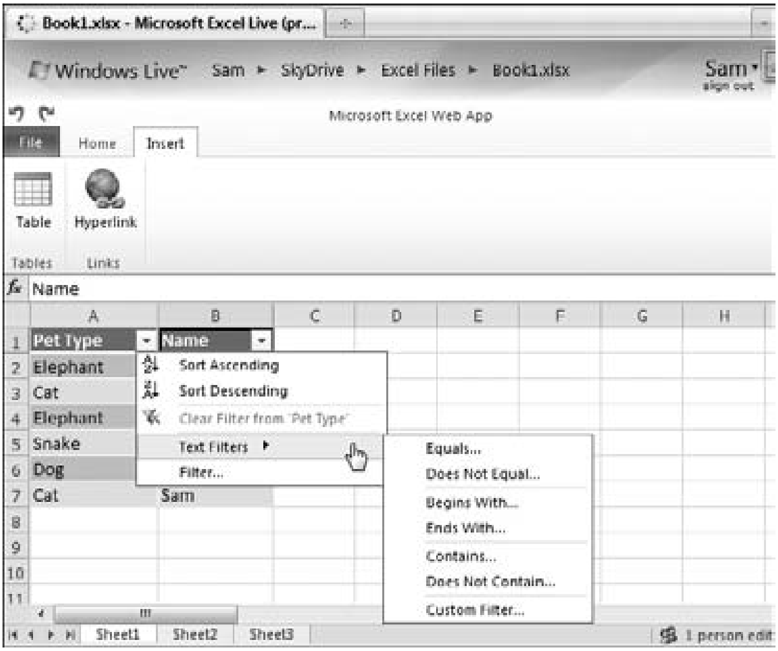

Figure 26-13: *As with the tables in desktop Excel, you can click the headers to apply a different sort order or a new filter.*

The common theme of the Excel Web App is this: it supports only a subset of Excel's features, but when it supports a feature, it duplicates the desktop approach as closely as possible. In future releases, more features will migrate from Excel into the Excel Web App, and its capabilities will grow.

The best way to get a feeling for what the Excel Web App can and can't do it is to give it a test drive. Enter some data, move it around, create a table, and then browse through the simplified ribbon. If you're looking for a simple way to share spreadsheets and keep working over the Web, the Excel Web App will impress you. But if you're a power user who relies on Excel's premier features, or you spend a lot of time creating charts and graphics, you'll probably miss your favorite features.

Note: The Excel Web App doesn't have a backstage view feature. When you click the File button, you'll see a short menu of options instead.

Saving Files

If you spend enough time poking around the ribbon and File menu in the Excel Web App, you'll notice that something is missing. No matter where you look, you won't find the Save command.

This omission reflects an important difference in the way that the Excel Web App works. Desktop Excel saves changes when you ask Excel to save them. But the Excel Web App saves changes continuously as you edit. This way, the Excel Web App ensures that you never lose your work, even if you lose your network connection or your browser crashes—two annoyances that are all too common in the wild world of the Web.

Of course, automatic saving has potential drawbacks too—namely, you can't make casual changes or try out new sample data. Any modifications you make have the potential to replace important data. The problem is compounded by Excel's support for collaborative editing (as described in the next section), because other people will see your changes right away.

You might think that the Undo feature can help you out, but it's not quite as protective as you think. If you close the workbook by navigating away, or if someone else starts editing the document at the same time as you, you'll lose your Undo history. So if you need to make potentially dangerous changes—ones that you might want to reverse, or ones that you don't want other people to see yet—the best bet is to open a copy of the workbook in desktop Excel, as described on page 770.

Collaboration: The Excel Web App's Specialty

The Excel Web App's niftiest feature is its support for pain-free collaborative editing. Unlike workbook sharing (page 770), you don't need to do anything to turn on collaborative editing in the Excel Web App. In fact, there's no way to stop it. Every file you upload to SkyDrive is automatically available to everyone who wants to edit it. While you're working on a spreadsheet in the Excel Web App, it keeps you informed about the other people who are editing the same file (Figure 26-14).

Note: The desktop version of Excel doesn't play along with the Excel Web App's collaborative editing. As long as the workbook is open in desktop Excel, it's not available for editing in the Excel Web App.

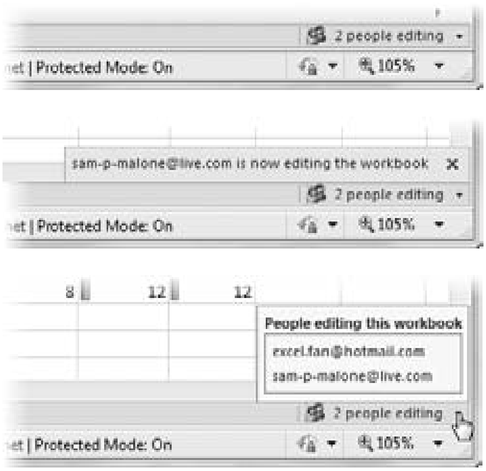

Figure 26-14:
Top: At any time, you can look to the bottom-right corner to find out how many people are working on the current file.

Middle: Excel alerts you when a new person opens the workbook for editing.

Bottom: When you click the bottom-right corner, the Excel Web App displays the full list of editors, identifying each person by his or her Windows Live ID email address.

UP TO SPEED

Collaborative Editing Scenarios

The collaborative editing feature of the Excel Web App isn't just a way to let a family of accountants review the same figures. It can also serve as everything from a data collection tool to a live chat application. Here are some ideas that show how you can use collaborative editing:

- **Collecting data.** The Excel Web App is perfect if you want to get information from a group of people in one place. To prevent the workbook from becoming a data-entry free-for-all (and risking overlapping edits), create a separate worksheet for each person that has information to contribute before you upload the document. That way, everyone can work on the same time, but on completely separate cells.
- **Keeping up-to-date.** If you need to have up-to-the-minute numbers at your fingertips, the Excel Web App may be the answer. You simply need to leave the workbook open in your Web browser, and the changes that other people make will appear on your desktop almost instantly. You could use this approach to monitor everything from current product prices to the status of different members on a team.
- **Chat.** Whether you want to exchange ideas in a meeting, or collaborate on a list of dishes for a company potluck, the Excel Web App makes it easy (in a freewheeling sort of way). And unlike a chat application, the Excel Web App can accommodate mindless banter *and* actual work—that is, numbers and formulas that you save for later.

Whenever more than one person is working on a document, it's possible for changes to clash. You saw this phenomenon with sharing files on a network (page 742), where the situation can quickly get messy. But the Excel Web App takes a much more relaxed approach to collaborative editing, called *last-in wins* updating. To understand this system, imagine what happens if two people start changing the same cell at once. When the first person finishes and moves to a new cell, the Excel Web App commits the change and updates the cell. But when the second person finishes, the Excel Web App commits the second change, wiping out the first one. The advantages of this approach are no locking and no cryptic error messages. The disadvantage is that your edit can be overwritten almost immediately, and the Excel Web App won't even notify you about the change.

Your opinion about the last-in-wins editing system depends on the way you plan to use the Excel Web App. In Microsoft's testing, they discovered that people who wanted to collaborate with the Excel Web App rarely made changes to the same cells. But if you need a strict versioning system to make sure no data goes missing in a critical workbook, the laissez-faire style of the Excel Web App probably isn't for you.

Taking a Workbook Back to Desktop Excel

The Excel Web App gives anyone an easy way to pop into an Excel document, add some information, and get back to work. It shines if you (or a group of people) need to edit an Excel workbook from multiple computers. It's especially handy for making quick edits on a computer that doesn't have Excel installed.

However, sometimes the Excel Web App isn't the best tool for the job. For example, you might want to use a feature that's available only in the desktop version of Excel (like spell check), or you might want to make changes that are quicker and easier to do offline (like searching and replacing text, or troubleshooting misbehaving formulas). In other cases, you might choose to take exclusive control of the workbook before you make major changes. That way, you're less likely to run into a problem with overlapping edits, and less likely to obliterate something important with an automatically saved change you didn't want to make.

You have several ways to deal with scenarios like these. If you're using the Excel Web App to provide a copy of an important workbook, you can simply change your local version of the file (which is the main version, after all), and then upload a new replacement copy. But if the main version is the online version, you need to open that in the desktop version of Excel. You accomplish this by finding the file in Excel's recent document list, or by choosing File→Open in Excel in the Excel Web App. Either way, the file is downloaded to your computer, opened in Excel, and locked. That means that while you're working on it, no one else can use the Excel Web App to change it (Figure 26-15).

When you're finished editing an online workbook in desktop Excel, simply choose File→Save (or press Ctrl+S) as you would with any ordinary Excel workbook. At this point, Excel uploads the updated file to its online location (Figure 26-16). Chose File→Close when you're finished to release the file and let other people edit it online.

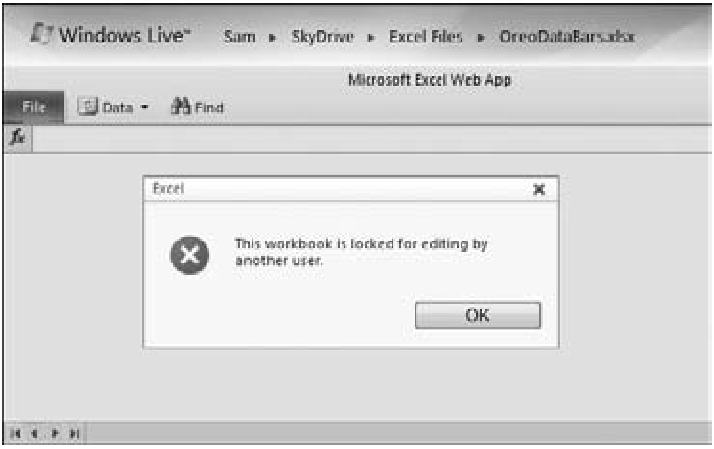

Figure 26-15: *This is the message you get if you try to edit a workbook in the Excel Web App that's open in the desktop version of Excel. If you just want to look at the spreadsheet without making changes, you can use the View command on the file detail page (Figure 26-9), or you can download a copy to your computer.*

Figure 26-16: *The status bar tells you when Excel is uploading a file. It does this in the background, which means you can get back to work viewing or editing the workbook before the upload is finished. If you close Excel while an upload is underway, Excel waits until the upload is finished before shutting down.*

In some situations, you may want to open a workbook in desktop Excel without locking it. You may simply want a copy of the workbook to use as the starting point for a different file. In this case, you have two options: you can click the Download link on the file detail page (Figure 26-9), or you can choose File→Download a Copy while editing the workbook in the Excel Web App.

Unsupported Features

As you've no doubt already noticed, the Excel Web App lacks the majority of Excel's less essential features. For example, if you want to debug a formula, run a spell check, or fill in missing values with the Solver add-in, you're out of luck. Removing these features is relatively harmless, if a bit inconvenient. It won't prevent you from opening or editing an existing workbook.

However, other features, once applied, *do* affect the way a spreadsheet works. If the Excel Web App can't support these features, it won't let you open the workbook at all. Here are the most commonly used forbidden features:

- Scenarios (page 594).
- Comments (page 724).
- Workbook protection, data validation, digital signatures, and IRM (Chapter 708).
- Ink annotations, which are usually "drawn" on a worksheet with a pointing device called a stylus.
- Embedded and linked objects (Chapter 27).
- XML maps (Chapter 23).
- Macros or custom functions (Chapter 28).
- Buttons and other Visual Basic controls (page 818).
- Workbooks that are saved with the "Show formulas in cells instead of their calculated results" settings switched on.

Tip: For the most comprehensive list, which includes rarely used features, unusual options, and legacy settings, see *http://tinyurl.com/34nk9nu*.

If you try to open a file with one of these features, you'll receive an error message like the one shown in Figure 26-17.

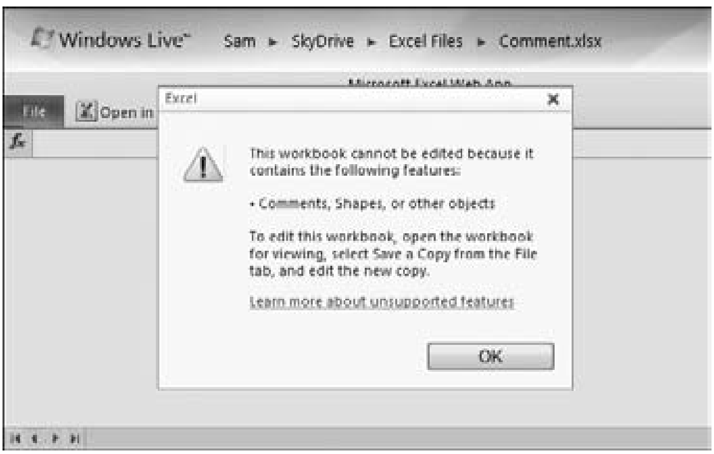

Figure 26-17: *Because this worksheet includes comments, you can't edit it. However, you can still view it in read-only mode by clicking the View link instead of Edit on the file details page. Incidentally, even in view mode, the pop-up comment boxes won't appear in the worksheet.*

Unfortunately, Excel doesn't include a tool for checking if a workbook is compatible with the Excel Web App. Instead, it's up to you to upload it, and then try editing the file in your browser.

Partially Supported Features

A second class of features can't be applied within the Excel Web App, but keep working if they're already present in a workbook. Here are the most notable examples:

- **External references.** If you have a formula that points to a cell in another workbook, the Excel Web App uses the most recent value in its place.
- **Grouping.** The Excel Web App doesn't let you group cells. But if the grouping has already been set up with desktop Excel, you can use it in the Excel Web App to collapse columns and rows out of sight.
- **Pictures.** They appear as the familiar floating picture boxes. However, you can't move, resize, delete, or add a picture with the Excel Web App. Furthermore, the Excel Web App doesn't support some types of picture effects, like rotations, shadows, cropping, and 3-D bevels. If your graphics use these effects, the Excel Web App still displays them, but without the extra fanciness.
- **Charting.** As you've already learned, you can't create charts in the Excel Web App. However, if there's already a chart in your workbook, it appears in a floating box. Like a picture, you can't move, remove, or modify it. However, if you edit the data that's displayed in the chart, the Excel Web App will regenerate the chart and refresh the chart picture. This means that charts can still be quite useful in the Excel Web App, as long as you remember to create them in desktop Excel first.
- **Conditional formatting and sparklines.** The Excel Web App won't let you create conditional formatting rules or sparklines. But if they're already in place, the Excel Web App will use them, and it will refresh the formatting and redraw the sparklines as the values in the linked cells change.

The clear theme here is that these features still work in the Excel Web App; they just need a helping hand from desktop Excel. You can use them, but you may find you have to switch back and forth between the Excel Web App and the desktop version of Excel fairly often.

The routine is simple enough. Every time you need to add (or tweak) the grouping, charts, sparklines, or conditional formatting in your workbook, you use the File→Open in Excel command to bring the workbook into desktop Excel. Once you're finished, you close the file to make it available online through the Excel Web App.

Note: In a sense, the Excel Web App already supports the most complex part of all these features. For example, it can accurately plot numbers and draw charts. Similarly, it can examine values and apply conditional formatting. They only missing ingredients are the windows that let you pick and choose your options. For this reason, many Excel fans speculate that future versions of the Excel Web App will support far more features–and may even begin to close the gap between the Web and desktop worlds.

CHAPTER

27

Exchanging Data with Other Programs

No program is an island. Sooner or later, you're probably going to want to take your carefully crafted Excel data and insert it into a completely different program. Maybe you want to bolster a presentation slideshow with some real data or give some heft to a report. No matter what the reason, you need a flexible way to share Excel tables and charts with other software programs.

Fortunately, Microsoft designed the modern Windows operating system with exactly that idea in mind. Windows lets you integrate different types of data through a pair of features called *linking* and *embedding*. Using linking and embedding, you can plant Excel data in other programs. You can also do the same thing in reverse, incorporating objects from other programs into your worksheets.

In this chapter, you'll learn how to use linking and embedding to integrate different types of content into one document. Once you've mastered these techniques, you'll learn how to export and import raw worksheet data.

Sharing Information in Windows

Every program has its own strengths and weaknesses. For example, Word provides the best tools for formatting long reports, while Excel shines when crunching numbers and charting trends. PowerPoint creates slick slideshows, while Access lets you store and search vast interrelated tables of information. Software developers realized long ago that no one could create a single program that was perfectly suited for every type of document.

Tip: Looking to polish your Word, PowerPoint, and Access skills? Check out *Word 2010: The Missing Manual, PowerPoint 2010: The Missing Manual,* and *Access 2010: The Missing Manual.*

However, it's not always realistic to separate different types of data. People need ways to integrate the strengths of *all* their favorite programs. Say you want to put an Excel chart in the middle of a Word document (or even attach a Word memo to the end of an Excel worksheet). To build this type of compound document, you need to rely on both Word and Excel. Fortunately, the programmers who built Windows had this type of data linking and sharing in mind from the very beginning, and they integrated it right into your favorite operating system's fabric.

Before you attempt any of these maneuvers, it helps to understand that you can transfer information between two programs in two different ways:

- **Embedding and linking objects**. Essentially, embedding and linking let you put a document from one program *inside* a document from another program. The trick is that even when you combine these two documents, they both remain in their original format. If you embed an Excel table in a Word document, you can still use Excel to edit the embedded table. Embedding and linking is the Windows way to share data. Also, if you use linking, you can refresh the copied information based on changes you make to the original source.
- **Importing and exporting data**. Importing and exporting are more traditional ways of sharing data, but they still make good sense in a lot of situations (for example, when the program you're looking to work with is *not* part of Microsoft's Office suite). When you import or export data, you convert a document written in one program to a format understood by another program. In some cases, one program may understand another's format so well that you don't need any real conversion. In other cases, the process of importing data changes the data's structure, removes its formatting, or even strips away some information.

So which approach is best? Embedding and linking is the way to go if you want to combine your data from two different programs into one document (like a report in Word containing a chart from Excel). Even though both programs still need to get involved, you can create the illusion that there's only one document. Linking is also the only option if you want to create a document that updates itself automatically when a linked file changes. If a PowerPoint presentation links to an Excel spreadsheet, the charts in the presentation automatically update when you modify the numbers in the spreadsheet.

On the other hand, importing and exporting makes sense when you need to process the information contained in one program with the tools provided by another program. For example, if you have a table of information in a Word document that you want to analyze using Excel, you'll probably want to convert this table to Excel's format. If you simply embed the table inside a worksheet, you can see the data, but you can't manipulate it with formulas, charting, and other Excel features.

In this chapter, you'll consider both these approaches. First up: embedding and linking.

Embedding and Linking Objects

Embedding and linking are two tools that let you build *compound documents*, which contain content from two or more programs. Maybe you have a Word file that contains an Excel worksheet. You can save this document as one file and print it as one document, but you need to use both programs to edit its content. Figure 27-1 shows an example.

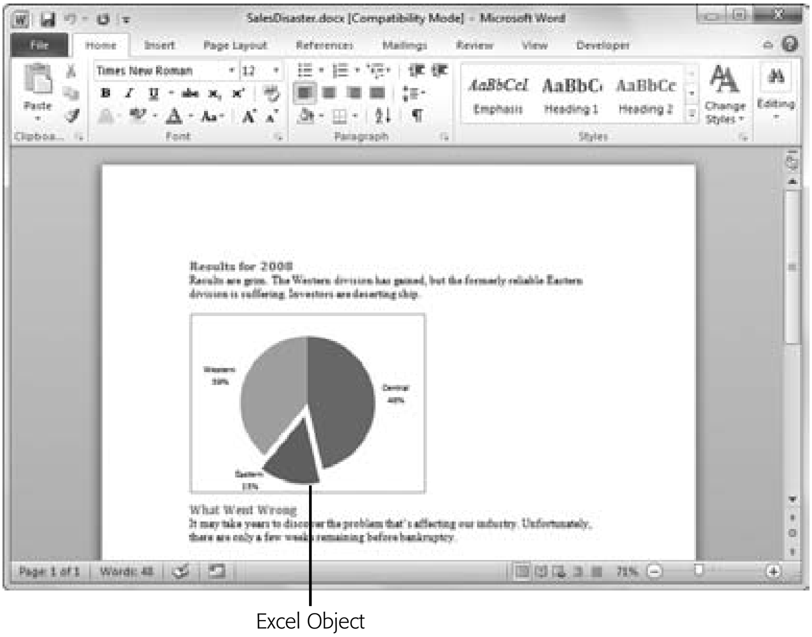

Figure 27-1: *If you place an Excel chart inside a Word document (as shown here), Word stores the chart in a floating box, and the document's text flows around it. The document looks just as it would if you had used Word's Insert→Illustrations→Picture command to place a picture of the chart in your document.*

Embedding and linking are really two different, but related, concepts:

- **Embedding** means that a copy of one document is physically stored inside the other. If you embed an Excel chart in a Word document, the Word (.docx) file contains all the Word document content *and* all the Excel worksheet data.
- **Linking** means that one document contains a *reference* to another document. If you put an Excel chart inside a Word document using linking, the Word file stores the Excel workbook file's name and file location information. A linked document still stores some of the information from the original source, so that you have something to show even if the link gets disrupted (if you delete the source file, for instance). However, a linked document's real power is that you can refresh the link whenever you need to—at which point, Excel copies the latest information from the source document.

Both embedding and linking generate the same results visually. Which one you use depends on how you plan to update the document. When you embed an object, you create a copy of the data inside the new document. You can edit that data separately, but it's no longer attached to the original source document. When you link a document, any changes you make to the source document automatically appear in the linked document. If you insert a chart inside a Word document, and then modify the chart in your worksheet, then the Word document also gets updated.

Note: Overall, embedding is easier to manage, but linking is your best approach if the source data changes frequently and you want to make sure the compound document always has the latest and greatest information. Linking also makes sense if you want to keep several documents synchronized with the same data (you want to do something like show an important worksheet table in four different Word reports). Embedding usually leads to larger files, because more information is placed into one file (the compound document that contains the embedded objects).

Exporting Charts Out of Excel

The best way to understand how linking and embedding work, is to try them yourself. Follow these steps to transfer an Excel chart into another Office program like Word:

1. **Open a workbook that contains a chart.**

   If you don't have a chart yet, create it now using the skills you picked up in Chapter 17, or download a workbook containing an example chart from the Missing CD page at *www.missingmanuals.com/cds*. (The chart workbooks are with the content for Chapter 17.)

2. **Select the chart, and then choose Home→Clipboard→Copy. Or, just press the shortcut key Ctrl+C.**

   The Windows clipboard is the key to transferring data between all Windows programs. Once you transfer data to the clipboard, you can retrieve it in another program.

3. **Switch to the program where you want to insert the data. Open the target document (or start a new document).**

   In this case, you need to switch to the Office program (like Word or PowerPoint) where you want to place the chart. If you don't have a suitable document, just create a new one (you can use the File→New command in Word or PowerPoint).

4. **Use the Paste Special command, if the program provides it. Otherwise, press Ctrl+V to perform a normal paste.**

   Many Windows programs provide an Edit→Paste Special command. When you use this command, you see a Paste Special dialog box (like the one shown in Figure 27-2) that lets you decide whether to link or embed the object, and what format it uses. If you're using an Office program with a ribbon (like Word 2010), use the Home→Clipboard→Paste→Paste Special command.

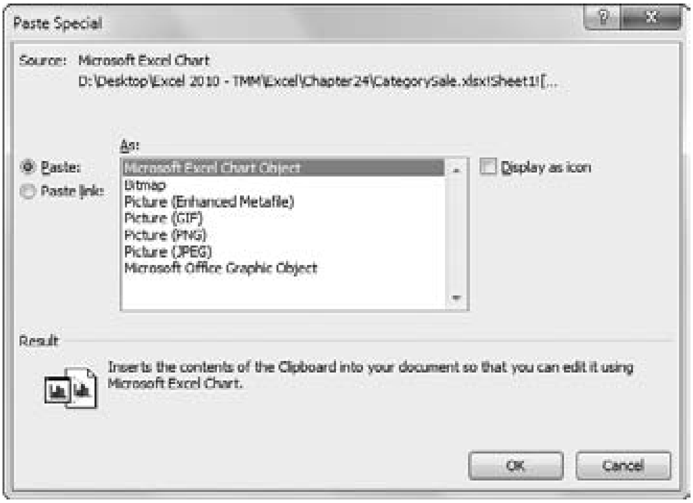

Figure 27-2:
Many Windows programs provide a Paste Special command. When you use Paste Special, you can choose to embed or link the object (depending on whether you select "Paste" or "Paste link"). You also have the option of converting the object into any of the formats listed in the As list box. In this example, you can paste the Excel chart as a Microsoft Office Excel Worksheet Object, which lets you edit the chart data later on. Or, you can paste your chart using one of a variety of different picture formats.

5. **Choose either "Paste" to insert an embedded object, or "Paste link" to insert a linked object.**

   Remember, an embedded object is a *copy* of the source data that's inserted into your document. A linked object is simply a reference that points to the original file.

6. **Choose the format you want to use.**

   The exact format options you get depend on the type of object and the application where you're performing the paste. For a chart, you'll typically see three main formats:

   - **Excel Worksheet Object.** Choose this to paste the Excel chart object, complete with all the behind-the-scenes Excel worksheets. This option gives you unlimited flexibility to change the chart later—for example, you can edit the worksheet data and even save the embedded workbook as a separate Excel file.

   - **Picture.** This takes a snapshot of your chart as a picture in one of several possible graphic file formats. This is the best option if you don't need to change your chart, because it takes less space and it won't bloat your document. If you decide to modify the chart later on, you'll need to open the original workbook, edit the chart, and paste a new version of the picture into your document.

Tip: The PNG format is a good choice for chart pasting, because it's smaller than a bitmap and does a better job of preserving fine detail than a JPEG.

- **Microsoft Office Graphic Object.** This only appears when you're pasting the chart into another Office application, like Word. This option is similar to Excel Worksheet Object, because it pastes both the chart image and the underlying worksheet information. However, there's a difference—when you use this option, you can show only the chart in your compound document (not the Excel grid of data). This is almost always the result you want anyway, which is why this is the best choice—when it's available.

Note: The Microsoft Office Graphic Object format avoids one quirk that can occur with the Excel Worksheet Object format. Sometimes, when editing an embedded worksheet object, the wrong part of the Excel workbook appears in your compound document. For example, you might paste a pie chart into Word, edit the pie chart, and then find that Word is showing the grid of cells from your worksheet instead of the actual pie chart. This issue never occurs with the Microsoft Office Graphic Object format.

7. **Click OK to paste the object.**

   The object appears inside the target document in a floating box.

UP TO SPEED

The Clipboard's One-Object Rule

Windows imposes one restriction when you copy data. You can have only one object on the clipboard at a time. When you copy a new object, any objects you previously copied become unavailable. It's surprisingly easy to stumble into this problem, because many different programs use the clipboard to store copied information. Windows Explorer's file-copying feature and Word's text-copying tool both use the Windows clipboard and displace whatever's currently there.

Fortunately, Excel and other Office programs don't always need to suffer from this limitation, since they have the Clipboard panel, as described on page 111. The Clipboard panel shows a list of recent objects you've copied in any Office program. Of course, the Clipboard panel can't help you if you're copying from information between Excel and another non-Office program, because the Clipboard panel is available only to Office programs.

Note: When you save a document that contains embedded objects, the file itself doesn't look any different. A Word document with an embedded Excel chart retains the standard .docx file extension. In fact, if you look at the file in Windows Explorer, you can't tell that it contains embedded Excel data in addition to its regular Word content.

To the untrained eye, there's no way to tell if your pasted chart is a homegrown image (stored as raw picture data) or an embedded object that's owned by a specific program. Either way, it looks like a floating box that you can move around your document as you see fit. You'll notice a difference only if you want to *modify* your chart. If you pasted an ordinary picture, you can't change the chart directly. Your only option is to delete the picture, edit the chart in Excel, and then paste in the new one. On the other hand, if you pasted full-fledged Excel data, you can modify that data to alter the chart (as described in the following sections).

Note: There's no reason you can't embed one object inside another, and then take that compound document and embed it as an object inside yet another program. This "Russian doll" approach probably won't do anything other than confuse everyone who uses the document, but it's completely possible.

UP TO SPEED

How Special Is Paste Special?

Different programs differ slightly in how they work with linking and embedding. In Office programs, the Paste Special command gives you the most flexibility, and most Windows programs provide this menu option. However, a few programs don't provide a Paste Special feature, in which case you have to rely on the less nimble Paste command.

Consider the simple Microsoft Paint drawing program, which comes with just about every version of Windows ever created. (You can run Paint by going to the Start menu, and then selecting Programs→Accessories→Paint.) If you go to the Paint menu, and then select Edit→Paste while an Excel chart is on the clipboard, Paint inserts a picture of your chart. Paint doesn't support object linking and embedding at all.

On the other hand, if you use WordPad (from the Start menu, select Programs→Accessories→WordPad), and then select Edit→Paste, then WordPad inserts an embedded Excel object. These examples just go to show that different programs make different choices when you use the Paste command. To control exactly what happens, look for a Paste Special command instead.

Editing a Linked Object

If you've inserted a linked object, you can easily change the linked content. Just open the original file in Excel, and then tweak the numbers. For example, if you pasted a linked chart object into Word, the next time you open the Word document you see the newly modified chart. The best part is that you don't need to do anything to tell Word to update the chart.

You can also use the program that has the linked object (in this example, Word), to open the linked object in the appropriate program (in this case, Excel). The exact process for doing this depends on the program that contains the linked object. In many Windows programs, you need to right-click the linked object, and then choose Linked Worksheet Object→Edit Link (or just double-click the object). This action pops up a new Excel window with the corresponding workbook.

Sometimes, when you change the worksheet data, the changes don't appear in the linked object right away. To trigger a refresh right away in a non-Office program, right-click the chart, and then look for an Update Link command. Or, just close the document that contains the linked object, and then open it up again to force it to get the latest data.

Tip: Linking works particularly well if you want to insert only a small part of a much larger workbook. Linking doesn't increase the size of the target document, because it stores only a link to the worksheet containing the data you're using.

Editing an Embedded Object

Working with an embedded object is sometimes a bit trickier, because the Excel object is actually stored *inside* the other file (for example, inside a .docx Word document). As a result, you can't just open the workbook directly in Excel. Instead, you need to open the host program (like Word), and then tell it to fire up Excel for you.

Different programs vary in how they work with embedded objects. Usually, you need to right-click the object, and then choose a command like Worksheet Object→Edit. Or, you can double-click the embedded object. At this point, a remarkable transformation takes place. The embedded object turns into a miniature Excel window, and the Excel ribbon appears (Figure 27-3). When you're finished, you simply need to click another section of the compound document, and the embedded Excel data reverts to an ordinary object.

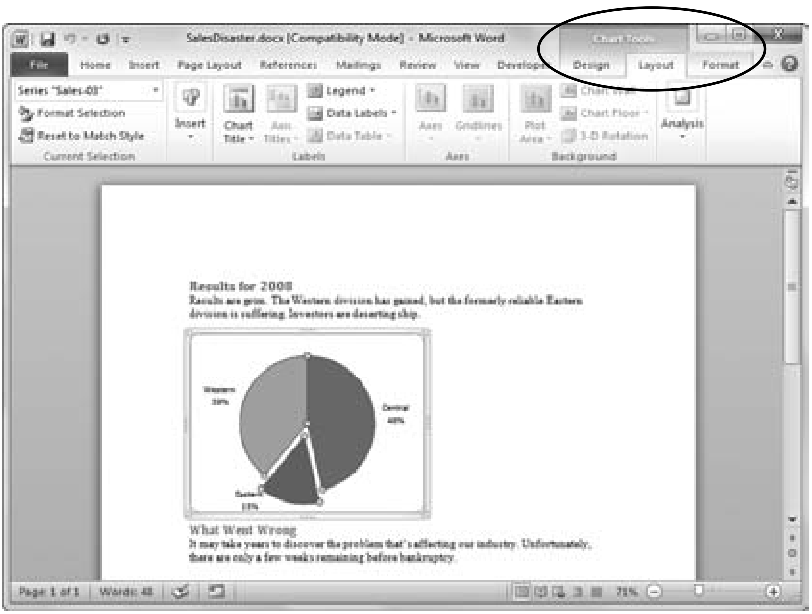

Figure 27-3: *Thanks to object embedding, you can do some chart tweaking right inside Word.*

If you find this arrangement a bit cramped, you can open an embedded Excel object in a separate Excel window. Usually, you do this by right-clicking the object, and then choosing a command like Worksheet Object→Open. A new Excel window appears with your embedded chart. When you're finished making your changes, choose File→Close and Return to close the "virtual" workbook and automatically insert the new data back into your compound document.

Note: If you chose to use the Microsoft Office Graphic format instead of the Excel Worksheet Object format (see page 780), you must right-click the chart and choose Edit Data to modify the numbers. This command opens the workbook in a separate Excel window, just like Worksheet Object→Open.

Interestingly, you'll notice that your workbook doesn't look exactly the same as when you first copied it. That's because when you paste an embedded chart using the Excel Worksheet Object format, Excel moves the chart to a separate worksheet in the document (Figure 27-4). This transformation doesn't happen if you use the Microsoft Office Graphic format.

When you edit an embedded workbook in an Excel window, it seems like you're editing a standalone workbook file. However, it's just a clever illusion. In reality, all your Excel data is embedded in another file that belongs to another program (like Word or PowerPoint).

Note: Remember, when you embed an object in another document, you're creating a copy of your workbook. The original Excel workbook file lives on, with the same data. It's important to realize that these two documents are completely separate. If you modify one, it doesn't affect the other. In fact, you can even delete the original workbook file and still use the embedded object.

FREQUENTLY ASKED QUESTION

Super-Size Compound Files

My compound files are gigantic–is there anything I can do to keep the size down?

An embedded chart object contains more than just the chart–it also has the original worksheet's data. You need some of the data to build the chart. To be on the safe side, Excel embeds the entire workbook in the document. Other objects, like pictures and additional charts, also come along for the ride. These accessories are among the reasons a compound document can quickly bloat to an uncomfortably large size.

Before you embed an object, it's a good idea to create a new Excel workbook that doesn't have any unnecessary data or objects. If you want to embed a chart, start by saving a new version of the workbook that contains the chart. Then, delete all the extra information in the workbook. Once you're finished paring the workbook down, you're ready to embed the chart in another program. This technique cuts down on the amount of information that the target program needs to store.

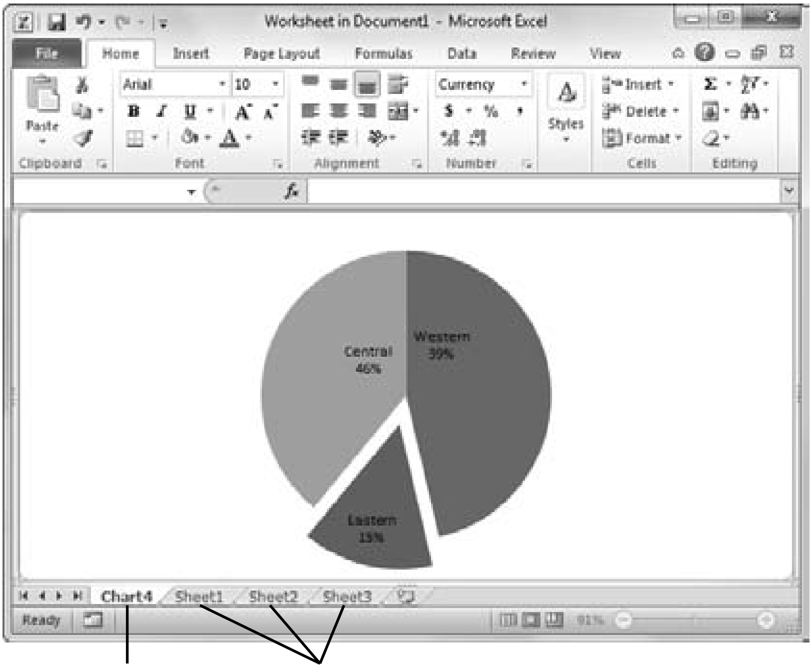

This tab has the chart that shows up in the Word document

These other tabs represent your orginal workbook, which was copied into your compound document

Figure 27-4: *To change your data, you need to find the worksheet where the chart's data resides. This task isn't too difficult, because an embedded chart object actually contains your complete workbook. You can switch from one worksheet to another using the familiar worksheet tabs. (Here, the original chart data is in Sheet1.) However, make sure you return to your chart worksheet before you close Excel. That's because the current worksheet is the only worksheet that appears in your compound document.*

Importing Objects into Excel

You can use the same process described in the previous section to copy objects *out* of other programs and place them inside Excel. In Excel, embedded objects appear as floating boxes (like pictures and charts). You can position these floating boxes wherever you want on your worksheet.

Embedding doesn't make sense for all types of objects. Although technically you can paste a bunch of text from a Word document as a floating object in a box, it doesn't integrate all that naturally into an Excel worksheet. However, other types of objects may make more sense. One example is an image format that Excel doesn't support. You could paste this image as an embedded object by copying it from the appropriate drawing program. The trick is Excel's Home→Clipboard→Paste→Paste Special command, which works just like it does in all other Office programs, including Word.

Figure 27-5 shows an example of an object transferred from Visio, a program that helps create diagrams.

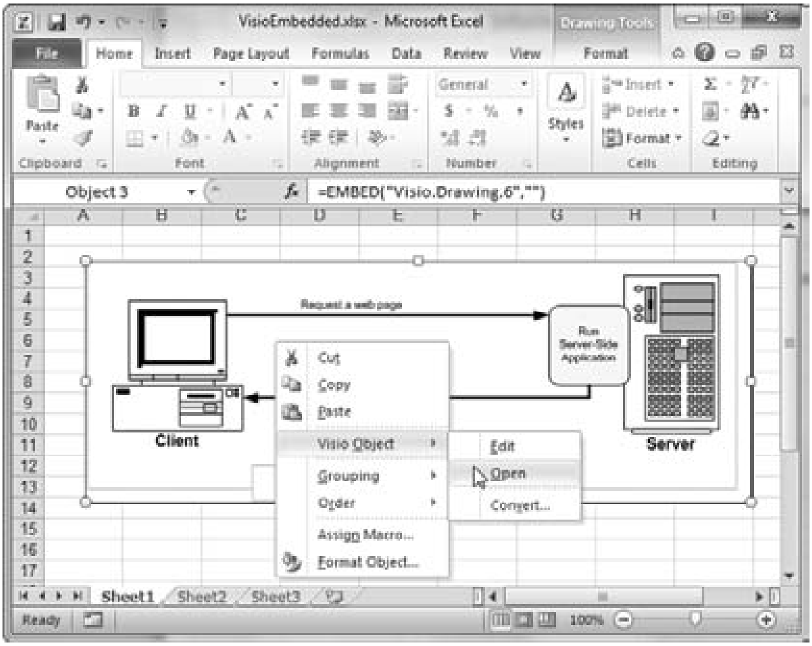

Figure 27-5: *Visio is a leading diagramming tool. Thanks to linking and embedding, you can integrate Visio objects into your worksheets as easily as charts or pictures. When you select Home→Paste→Paste Special in Excel and there's a Visio drawing on the clipboard, Excel lets you convert the diagram to a graphic, or paste it as a linked Visio object. If you choose the latter (as shown here), you're able to edit the diagram with Visio at any time.*

Transferring Data

With embedding and linking, two or more programs work together to create a compound document. However, even though the different objects appear side-by-side, different programs actually own them. When you embed an Excel table in a Word document, you can't spruce up any text in your table by using Word's built-in grammar checker.

Furthermore, sometimes the program you're embedding an object *into* can just as easily create the object you've inserted. Consider, for example, a typical Word document, which supports all the same formatting options as Excel, and can organize information into a grid perfectly well by using its table-creation features.

In such cases, you could transfer raw *data*, rather than objects. So instead of embedding a worksheet object inside Word, you could simply copy Excel's formatted worksheet data, and then move it into Word. You do lose the ability to update the information with Excel, but you gain a few benefits as well:

- You can edit the data directly in Word without needing access to Excel.

- You can edit the data quickly and more conveniently. This ability is particularly important if you want to format the data to match the rest of your document.
- The file is smaller than it would be if you used an embedded object.
- You avoid accidentally modifying information if you change the source worksheet (as you would if you used a linked object).

For these reasons, it's worth carefully considering whether you should copy a full-fledged worksheet object, or just transfer the information you want to use.

Exporting Tables of Data

The secret to pasting worksheet data into another program like Word is the same Paste Special command you use to create embedded and linked objects. The following steps walk you through the process.

1. **Select a range of cells from your worksheet, and then choose Home→Clipboard→Copy.**

   For best results, try to avoid selecting empty rows and columns.

2. **Switch to the target document, and then use the Paste Special command.**

   The Paste Special dialog box appears. This dialog box shows what you'd see if you were copying a selection of cells from Excel into Word. When you copy a selection of cells, the Paste Special dialog box gives you a slew of choices for how you want to import your data from Excel. As shown in Figure 27-6, the target program recognizes that your clipboard data includes ordinary text content. In most programs, you have the option to insert formatted or unformatted versions of the text, a linked or embedded object, or even a picture of the text that you've selected!

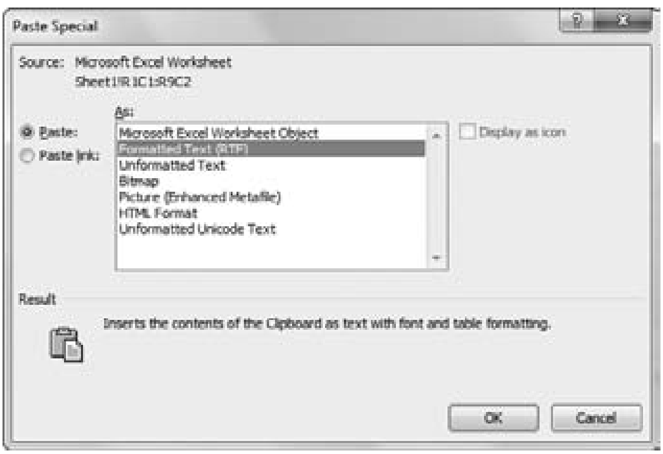

Figure 27-6:
In most cases, the best approach is to avoid embedding and linking altogether, and just choose formatted or unformatted text. Keep in mind that if you embed a worksheet object, you wind up embedding all the worksheet's data, even though you see, in the embedded object box, only the portion of data that you selected.

3. **In the Paste Special dialog box, select the type of object you want to insert.**

   You can choose any supported format from the list. The format choices depend on the type of object you're inserting. For example, Excel lets you paste the selected cells as an Excel object, or convert it to unformatted text, HTML text, or even a low-resolution picture.

   You can also choose how Excel synchronizes the pasted data with the source. Choose "Paste link" if you want to create a linked object that Excel refreshes based on changes you make to the source information. Choose "Paste" if you want to create an embedded object that you can modify independently.

4. **Click OK.**

   The Excel information appears in the format you chose.

Importing Tables of Data

There's also no reason that you can't take text out of another program, and then paste it into Excel. This tactic works best if the information is already arranged in a grid-like structure (for example, a Word table), but Excel can insert the data no matter how it's organized. Figure 27-7 shows an example.

Instead of using Home→Clipboard→Paste Special to paste content into Excel, you can use Home→Clipboard→Paste as a shortcut. When you do, Excel inserts formatted text, and a smart tag icon appears next to the newly inserted content. You can click this icon, and then choose either Keep Source Formatting (which applies whatever formatting was used in the original program) or Match Destination Formatting (which ignores the source formatting and keeps whatever formatting was already applied to the cells).

Importing Text Files

Sometimes you want your Excel spreadsheet to use data that's stored in an ordinary text file. In this case, you can import the information into Excel, but you need to go through a special conversion process. This conversion process scans the text and splits it into separate rows and columns, so that Excel can insert it into a worksheet's cells. Excel provides an intelligent tool, the Text Import Wizard, which helps you perform text file conversions.

To import a text file, follow these steps:

1. **From within Excel, select File→Open.**

   The Open dialog box appears.

2. **From the "Files of type" list at the bottom of the window, select Text Files.**

   Excel knows you're looking for files that have the extension .prn, .txt, or .csv. If you have a text file that has a different file extension, choose the All Files type instead.

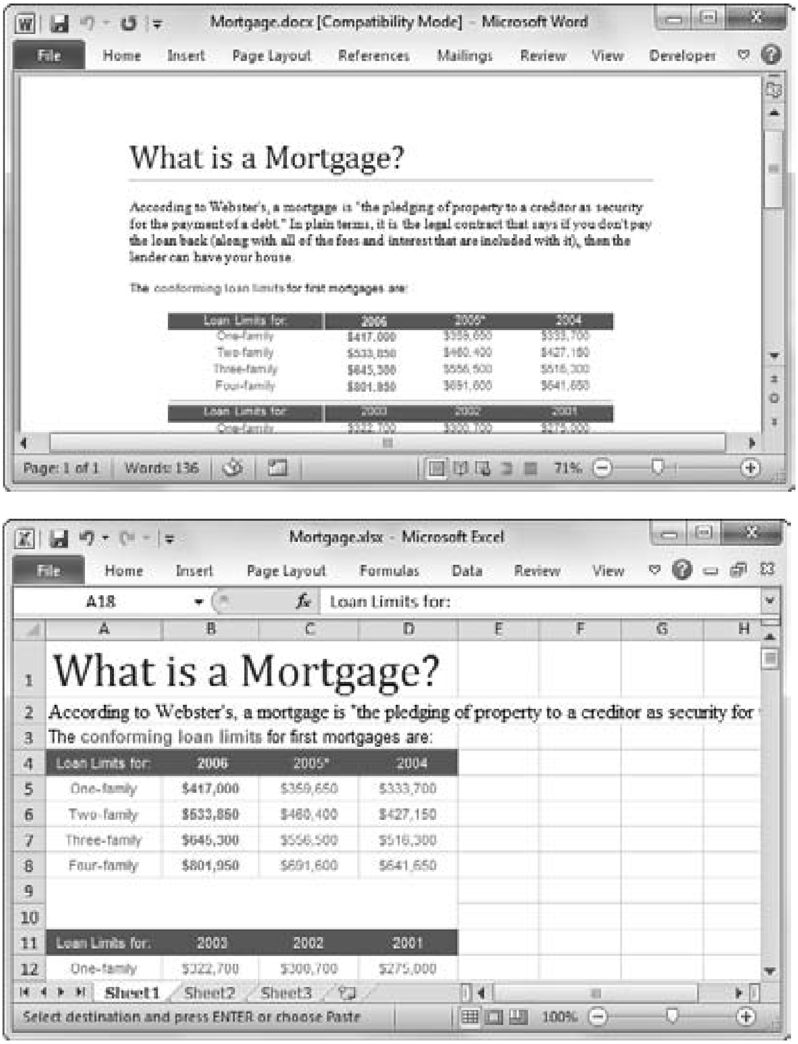

Figure 27-7: *These examples show some Word content (top) and how it appears when it's copied into Excel (bottom). As you can see, you can transfer large paragraphs, but the result isn't too pleasant. The entire paragraph is crammed into one cell, and it's hardly readable. On the other hand, tables are much easier to copy and paste. Excel correctly distributes the rows and columns into cells on the worksheet, and it even merges cells in order to match the original document's formatting. This technique works just as well if you copy a table of information from a Web page, and then paste it into Excel.*

3. **Browse to the file you want to import, and then select it.**

   For a quick and easy test, use the log.txt file included with the downloadable content for this chapter on the Missing CD page at *www.missingmanuals.com/cds*. This text file contains a list of log entries that reflect system activity on a computer.

4. **Click OK to continue.**

   Excel starts the Text Import Wizard.

UP TO SPEED

Why Text Files Matter

With all the copying, pasting, and formatting features available in Windows, you may wonder why anyone would ever stoop to the lowly level of plain text. In fact, using plain text is more common than you may think, because text files are the lowest common denominator when it comes to data. Even if you have a program that runs on a dinosaur-era operating system like DOS, you can still transfer information, as long as your program includes a text export feature.

Even if you're using a Windows program, you can still find yourself in this situation if the program doesn't provide a way to select and copy objects. You'll often find yourself in this predicament with programs that aren't document-oriented (in other words, programs that aren't designed to help you create, edit, and save some sort of document).

Consider a management tool that lets you assign projects to employees in a company. This tool probably doesn't let you copy employee and project information directly, but it might include an export feature that dumps this information into a text file. Once the text file is ready, you can use Excel's Text Import wizard to bring it into your worksheet.

5. **In step 1 of the wizard (Figure 27-8), choose "Original data type." You have two options, "Delimited" or "Fixed width," depending on your text file's format.**

   Usually, the program that creates the text file indicates what type of format it's using, or gives you the chance to choose one of several supported formats. If you're unsure of the exported format, you may need to take a closer look at the file by first opening it in a program like Notepad, which should help you figure out what kind of format you're dealing with.

   Delimited means that there's a separator—usually a comma, a tab, or a semicolon—between each column. In a comma-delimited file, the following data would comprise three columns of information, one with the region name (Region 1 is the first name), followed by separate columns containing the month and the number at the end of the line:

   ```
   Region 1,January,43432
   Region 2343,March,839
   ```

   Fixed width means that the columns are separated using a series of spaces. All the columns are spaced evenly. The same data is shown in a fixed-width format here. The problem with fixed-width formats is that a value can't exceed the maximum space allocated to the column:

   ```
   Region  1     January  43432
   Region  2343  March    839
   ```

Note: Most programs that export data to text use delimited text files. When you're importing data into Excel, you'll often use the delimited option. To tell whether your text file's delimited, look for its telltale sign—a repeated character that's used to separate each piece of information. Any character can separate columns in a delimited file, but commas and tabs are the top choices.

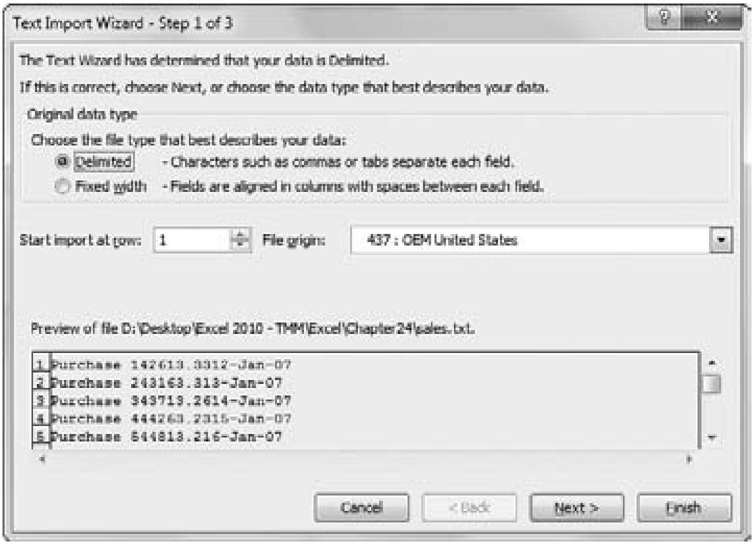

Figure 27-8: *In this example, a delimited text file containing a list of purchases is being imported. In the first step, the file appears as a stream of unstructured information.*

6. **Click Next, and then complete step 2 of the Text Import Wizard (Figure 27-9).**

   This step varies depending on whether you're importing a fixed width or a delimited file. If you're importing a delimited file, you need to specify, in the Delimiter box, the character that's used to separate columns. If you're importing a comma-delimited file (where a comma separates each column), then turn on the Comma checkbox. Once you make your change, Excel updates the preview, separating the data into columns using the delimiter you chose.

   If you're importing a fixed-width file, you need to indicate where each column starts. Click the preview window in the appropriate location, once for each column. As you click, Excel adds column dividers between each column.

7. **Click Next to move to the last step (Figure 27-10).**

   The last step lets you specify the data type for each column. Usually, the standard option (General) works perfectly well. If you choose General, Excel examines the content, and then changes it into a text, number, or date entry, depending on the type of values the file contains.

   If you want to override this automatic decision-making process, simply select a column in the "Data preview" section. Then, from the column in the "Column data format" box, choose the data type. You may take this extra step if you have a cell that contains numeric content, but you want to treat it as text.

8. **Click Finish to complete the wizard.**

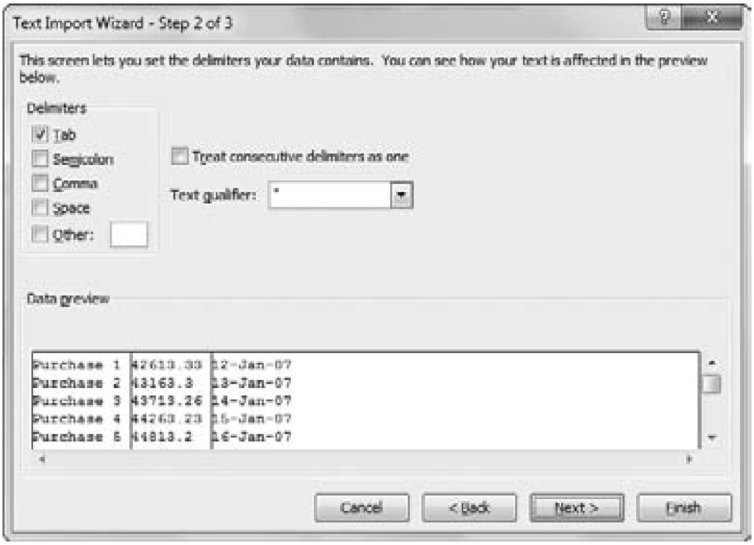

Figure 27-9: *Life gets better in step 2, which gives you the chance to specify the delimiter that separates columns (in this case, the tab character). Once Excel knows that the file is organized using tab characters, it can parse the file and display it properly in the "Data preview" window.*

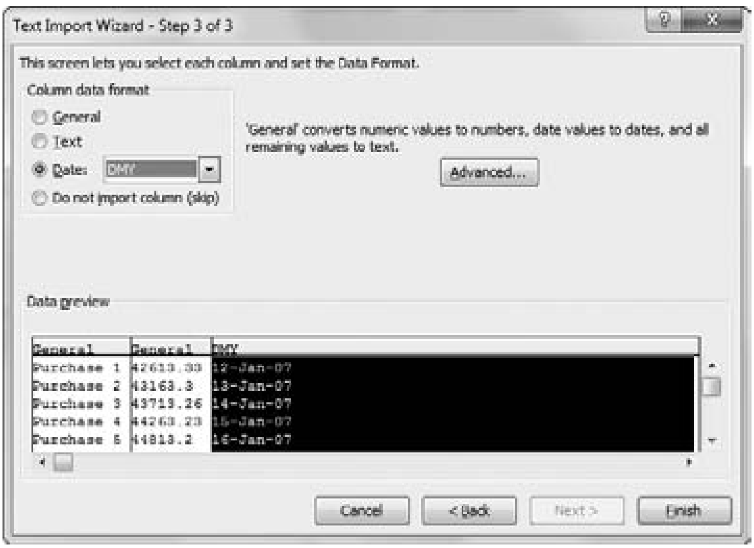

Figure 27-10: *The third step provides a last chance to change how Excel interprets the data in each column. Here, it's important to tell Excel the third column contains date information, not ordinary text.*

POWER USERS' CLINIC

Importing Text into an Existing File

Rather than using your text file to create a new workbook, you place it into an existing workbook. This process uses the same Text Import Wizard–the difference is that once the last step is finished, Excel asks you where you want to place the data. (You can use a specific location in the current worksheet, or create a new worksheet.) To launch the process, choose Data→Get External Data→From Text.

At first glance, you may wonder why Excel has two features that seem so similar. However, there's a subtle but important difference in the *way* you use these features. The File→Open command is designed for Excel fans who want to import external data, and then change it. On the other hand, the Data→Get External Data→From Text command is for those who just want to *analyze* the information that's in a text file. Here are two examples that make the difference a bit clearer:

- Your company has oodles of old data created by a Paleolithic piece of mainframe software. This data is stored in bare-bones text files. Now, your company is finally moving into the new millennium. It's up to you to take this data and change it into Excel workbooks, where you can maintain it more easily. In this situation, you need the File→Open command to make the jump from text files to more-capable Excel spreadsheets.
- It's your job to review the daily sales transactions and look for emerging patterns or warning signs. To perform this analysis, you use Excel's formula smarts and charting tools. But the data you need to analyze is emailed to you at the end of each day as a text file. It's up to you to pull this information out of the text file and look at it in Excel. In this situation, you need the Data→Get External Data→From Text command. That way, you set up the link once, and you can refresh your connection whenever you need to bring it the newly updated data from the latest version of the file (just choose Data→Manage Connections→Refresh All). Best of all, your linked formulas and charts update themselves automatically when the new data flows in.

Data comes in many forms, and Excel isn't limited to ordinary text files. In fact, the Data→Get External Data→From Text command is part of a broader set of features for creating *data connections*. These connections can link to outside text files, XML files, databases, or even pages on a website. Using the data connection feature, you can perform regular (and even automatic) refreshes so you always have the latest data on hand. For more information about Excel's data connection features, see Chapter 23.

CHAPTER
28

Automating Tasks with Macros

No one wants to repeat the same task over and over again. Excel gurus especially hate the drudgery of repetitious work like data entry and formatting cell after cell after cell. Fortunately, Excel offers serious time-saving tools called *macros*—miniature programs that automatically perform a series of steps within any workbook.

Excel macros are written in a full-blown programming language called *VBA* (short for Visual Basic for Applications). VBA is a scaled-down version of the incredibly popular Visual Basic programming language, and it's fined-tuned for Excel and other Office programs. Fortunately, you don't need to be a programmer to use VBA. In fact, you can create a simple macro using a special Excel tool called the *macro recorder*. The macro recorder works like a tape recorder, but instead of recording audio, it records keystrokes and mouse actions as you perform them.

In this chapter, you'll learn how to use the macro recorder, how macros work, and where Excel stores them. You'll also learn how to attach macros to shortcut keys, the Quick Access toolbar, and buttons in your worksheet. In the next chapter, your exploration goes one step further and plunges into the VBA language.

Macros 101

Macros can automate everything from simple tasks to complex operations, making your life immensely easier. Here are examples of tasks people commonly use macros for:

- Quickly inserting a group of cells you need often (like a company header for a report).

- Applying complex formatting to multiple rows or columns.
- Cleaning up text, rearranging cells, or transferring information from one place to another.
- Printing data in a specific format. This operation could involve switching to a custom view, selecting a part of a worksheet, and choosing specific print options (like paper size or the number of copies).
- Analyzing a selection of cells, perhaps by adding a set of functions or even by generating a specialized chart.

These examples are just the tip of the macro iceberg. When you create a macro, you can use the full range of Excel's features. It's like having a personal assistant, skilled in the ways of Excel, ready to carry out whatever instructions you provide.

Excel gives you two ways to create a macro:

- **Write the macro by hand, using pure VBA code**. This option isn't as difficult as you may think, but it isn't easy either. The problem isn't learning the *syntax*, or rules, of the VBA language (which are fairly basic); it's learning how to find all the Excel features you want to use. When you want to enter text in a cell, create a new worksheet, or print a selection, you need to find the right Excel *object* that can handle the job. Objects are programming tools that let you get at features you need. For example, if you want to write macro code that spell checks your spreadsheet, you need to know the name of the spell checker object.
- **Record the macro using the Excel macro recorder**. You turn the recorder on, and then go about your business, entering text, navigating a worksheet, and selecting choices from Excel's ribbon. While you work, Excel records each operation and translates it into the corresponding VBA code. When you're finished, you can stop the recorder, save the macro, and replay it to repeat all the actions you just performed.

Using the macro recorder is the simplest approach, particularly if you've never used a programming language before. However, writing macros by hand is much more powerful and flexible. It lets you create macros that can make decisions, repeat actions in a loop, and even prompt whoever's viewing your worksheet to enter some information.

In the next chapter, you'll take a closer look at the intricacies of VBA and start digging through the vast collection of Excel objects. In this chapter, you'll get right to work creating simple macros with the macro recorder.

Note: There's a middle ground between writing macros yourself and having Excel record your actions. You can record a macro with the macro recorder, and then tweak the VBA code with the Visual Basic editor in order to make it more powerful. At the start of the next chapter, you'll take a look at what a recorded macro looks like inside the Visual Basic editor.

Macro-Free and Macro-Enabled Workbooks

In Chapter 1, you took your first look at Excel's file formats (page 39). So far, you've been sticking with Excel's *.xlsx* file type, the format of choice. However, .xlsx files have a serious limitation when it comes to macro writing—they can't store macros. That means you need to either choose a different format for your workbook (one that does support macros), or store your macros elsewhere, in a workbook that does have macro-storing abilities.

Fortunately, Excel has two perfectly good file formats that support macros and all the same features as their macro-free counterparts:

- **.xlsm files** are macro-enabled workbooks. They're identical to .xlsx files, except they have the added ability to store macro code. Like .xlsx files, .xlsm files are stored using XML.
- **.xlsb files** are workbooks stored in binary format, which is more efficient for certain types of files (like very large spreadsheets). They also allow macro code.

Microsoft recommends that you use the .xlsm format. In some situations, extremely complex files may load faster if they're stored in the .xlsb format. To review the difference between Excel's XML-based formats (like .xlsm) and its binary formats (like .xlsb), see page 39.

Figure 28-1:
Excel gives macro-enabled workbooks (both .xlsm and .xlsb files) a different icon, with a superimposed exclamation mark. This icon lets you recognize a macro-enabled workbook–like the MacroWorkbook shown here on the Windows desktop–before you open it. But don't get too paranoid just yet. As you'll discover on page 803, Excel's macro security features make sure there's little to fear about malicious macro code.

The Macro Recorder

The macro recorder is easy to use, but keep a few key points in mind while you're recording so that you end up with a great collection of really useful tools:

- **Excel is watching you**. Excel captures every ribbon command or keyboard shortcut you use. In other words, don't do anything that you don't want recorded as part of the macro, unless you're willing to edit the VBA macro code after the fact. Also, try to avoid switching to another program while you're recording an Excel macro. Even though the macro recorder ignores anything you do outside of Excel, it's easy to confuse yourself and inadvertently add macro code you don't want by jumping back and forth.

FREQUENTLY ASKED QUESTION

A Reason to Like Macro-Free Workbooks

Why doesn't everyone use .xlsm files? Whose idea was it to create a file format that supports fewer features?

At first glance, it seems that Microsoft could have avoided a lot of confusion if it had just created a single file format and made sure that file format was able to store macro files. After all, why encourage people to use a file format that can't do everything?

The reason Microsoft created .xlsx files (and the reason they're standard when you don't need macro support) is security. In recent years, Microsoft's become increasingly paranoid, as have the thousands of companies using Excel in a business setting. Their nightmare is that one day the company accountant opens an innocent-looking Excel spreadsheet, and a malicious macro toasts the computer.

Of course, evil Excel macros are fairly rare, and you're guaranteed to stay safe if you avoid opening spreadsheets that come from mysterious sources (like a "Get Well Endowed" email message). But as more companies use Excel to deal with sensitive, mission-critical information, even a small risk is too high. As you'll learn a little later (page 803), Microsoft automatically disables macros when you open a macro-enabled workbook for the first time. It's up to you to re-enable them and accept the risk.

So what does all this have to do with the two file formats? These formats let network administrators quickly distinguish between macro-free files (.xlsx), which are always safe, and potentially risky macro-enabled files (.xslm). And thanks to the different file extensions, administrators can create spam-blocking rules that automatically remove macro-enabled workbooks that are attached to email messages. In other words, a little bit of extra complexity makes it easier to keep company systems safe.

- **You don't need to work fast**. The macro recorder doesn't record anything in between each action you perform. If you browse through the ribbon for 20 minutes before you eventually select Home→Clipboard→Paste, Excel just adds a single line of VBA code to your macro—the line it needs to invoke the Home→Clipboard→Paste command. As a result, macros tend to execute very quickly when you replay them—much more quickly, in fact, than when you perform the actions yourself.
- **Try to be generic**. The ideal macro is general enough that you can reuse it in a wide range of scenarios. If you make a macro so specific that you can use it only once, the macro won't get much use, and all your hard work will be wasted.

Before you jump into the macro-recording studio, it helps to understand the difference between recording modes and to know where Excel saves your macros. The next two sections cover these topics in more detail.

Relative and Absolute Recording

When you click a command in the ribbon, Excel's macro recorder knows exactly what you're doing. However, sometimes Excel needs a little guidance from you in order to decide how to interpret some of your actions. When you type in text, Excel can interpret your action in two different ways, depending on which of two recording modes you're in (page 801 tells you how to change modes):

- **Absolute reference mode**. In absolute reference mode, Excel stores the absolute references for the cells that you're modifying (for a refresher on absolute cell references, see page 254). When you play the macro again, the macro affects only these cells.
- **Relative reference mode**. In relative reference mode, Excel tracks how far you move from your starting position. That means when you play the macro again, Excel takes your current location into account.

You can most easily understand the difference between the two modes by following an example. Imagine you move to cell A1, and then start recording a macro. You then move two columns to the right to cell C1, type in the number 42, and then save the macro.

If you used absolute reference mode, here's the series of instructions that Excel stores in your macro:

1. **Move to cell C1.**
2. **Enter the number 42.**

On the other hand, if you used relative reference mode, Excel stores this list of instructions:

1. **Move two columns to the right (from wherever the active cell is).**
2. **Enter the number 42.**

You'll see the difference if you clear the worksheet, move to cell E10, and play the macro. If you used absolute reference mode when you recorded the macro, then the macro always returns to cell C1 to enter its information. If you used relative reference mode, Excel enters the number 42 two cells to the right of the current cell, in cell G10.

Note: Absolute reference mode works if your data always needs to be in the same position. This need arises most often when you have to add some sort of header information at the top of a worksheet. Relative reference mode works if you need to repeat a task in several places—like bolding, italicizing, and enlarging the font in a series of cells you choose. It's the most common approach.

Where Macros Live

Once you've crafted the perfect macro, you need to make sure Excel stores it so that you can find it later on. In Excel, every macro's attached to a workbook and saved in that workbook's .xlsm file. When you open a workbook that contains a set of macros, Excel makes them available instantly.

Excel beginners often assume that before you can use a macro in a particular workbook, you need to store the macro inside *that* workbook. In fact, macros have a much greater range. As soon as you open a workbook that contains macros, Excel makes those macros available to every other workbook that's currently open.

In other words, imagine you're editing a workbook named SalesReport.xlsx, and you open another workbook named MyMacroCollection.xlsm, which contains a few useful macros. You can use the macros contained in MyMacroCollection.xlsm with SalesReport.xlsx without a hitch. Once you close MyMacroCollection.xlsm, those macros are no longer available.

This design makes it easy to share and reuse macros across workbooks (and between different people). It's up to you whether you want to store macros in your most commonly used workbooks or create standalone workbooks that hold collections of useful macros.

When you record a macro, Excel gives you three slightly different storage options:

- **This Workbook**. If you choose this option, Excel stores your macro in the current workbook. Remember, you need to save this workbook as a macro-enabled .xlsm file or a binary .xlsb file, or you'll lose your macros.
- **New Workbook**. If you choose this option, Excel automatically creates a new workbook (which it opens in a separate window) and stores your macro there.
- **Personal Macro Workbook**. If you choose this option, Excel stores your macro in a special hidden workbook named Personal.xlsb. The Personal.xlsb workbook opens automatically whenever you start Excel (although it remains hidden), so macros in this workbook are always available no matter what workbook you're using.

Note: If you try to save a workbook that contains a macro using Excel's macro-free .xlsx file format, you receive an error message warning you that you'll lose your macros. Instead, choose File→Save As, and then, from the "Save as type" list, pick "Excel Macro Enabled Workbook (*.xlsm)."

Just because you place a macro in a specific location doesn't mean it needs to stay there. In fact, it's quite easy to copy a macro from one workbook to another, or even move it into and out of the personal macro workbook. All you have to do is fire up the Visual Basic editor and learn a little bit about how it organizes macro code. You'll get an introduction to the editor in the next chapter.

Tip: One useful place to put macros is in an Excel template. That way, every time you or anyone else creates a new workbook based on the template, you have immediate access to the macros. In order to put macros in a template, record them inside the template using the This Workbook option. See Chapter 16 for more information about templates.

FREQUENTLY ASKED QUESTION

Locating the Personal Macro Workbook

Where's the personal macro workbook hidden?

It makes sense to save your most useful macros in your personal macro workbook, so that they're always available, without requiring any extra steps. Excel stores the personal macro workbook in a file called Personal.xlsb, which lives in a folder named XLStart.

On a Windows Vista or Windows 7 computer, you can find the XLStart folder in a location like *C:\Users\UserName\AppData\Roaming\Microsoft\XLStart*, where UserName indicates the account name of whomever is currently logged on to Windows. On a Windows XP computer, you'll probably find it in a folder like *C:\Documents and Settings\UserName\Application Data\Microsoft\XLStart.* Keen eyes will notice that the XLStart folder is right next to the Templates folder (page 455). In other words, if your Templates folder is located at *C:\Users\billjones\AppData\Roaming\Microsoft\Templates,* your XLStart folder will be at *C:\Users\billjones\AppData\Roaming\Microsoft\XLStart.*

Each person with an account on your PC has a separate personal workbook, so if you log on as someone else, you won't have the same collection of macros. Also, keep in mind that Excel doesn't actually create the personal macro workbook until you add your first macro to it. So, unless you've recorded at least one macro for the personal macro workbook to hold, there's no point hunting for it or the XLStart folder.

But if you're in the habit of placing a lot of important macros in the Personal.xlsb file, you should search it out. That way, you can easily back it up so that you don't lose all your hard work when the next computer virus strikes.

Recording a Macro

Now that you've learned the basics of Excel macros, it's time to try creating one:

1. **Choose View→Macros→Macros→Record Macro.**

   The Record Macro dialog box appears (Figure 28-2).

Figure 28-2:
Here, Excel is about to start recording a macro named InsertHeader. Excel will store this macro in the current workbook along with an optional description. Since the "Shortcut key" box has been left blank, Excel won't assign a shortcut key to the macro.

Tip: For an even faster way to start recording a new macro, click the macro recording button in the Status bar. It looks like a worksheet window with a red circle (representing a record button) superimposed on top. If you don't see the button, right-click the Status bar and make sure the Macro Record item is checked.

2. **Type in a name for your macro.**

   The macro name can include letters, numbers, and the underscore character (_), but it can't include spaces or other special characters, and it must begin with a letter. Use a descriptive macro name like CompanyHeader or Create-SalesChart, because you'll identify the macro based on the name later on. (The macro name is also the name Excel uses for the subroutine where it records all the VBA instructions.)

3. **If you want to create a keyboard shortcut, choose one now.**

   Shortcut keys let you launch your macros in a hurry. The only shortcut keys you can use are Ctrl+*letter* combinations. You can use uppercase or lowercase letters, as in Ctrl+A (which means hold down the Ctrl key, and then press A) and Ctrl+Shift+A (which means hold down the Ctrl and Shift keys, and then press A).

4. **Choose a storage location from the "Store macro in" list.**

   This option determines where Excel saves the macro.

WORD TO THE WISE

The Dangers of Macro Shortcuts

Using shortcut keys can be dangerous because Excel doesn't warn you if you choose a shortcut key that already corresponds to another Excel task.

If this sort of conflict happens, Excel always uses the shortcut key for the *macro*. This habit can cause confusion if other people use your macros and you replace a common shortcut key. Imagine their surprise when they hit Ctrl+S to save a document and end up triggering a macro that turns all negative numbers positive.

Here are some common key combinations that you should *never* assign to macro shortcuts, because people use them too frequently:

- Ctrl+S (Save)
- Ctrl+P (Print)
- Ctrl+O (Open)
- Ctrl+N (New)
- Ctrl+X (Exit)
- Ctrl+Z (Undo)
- Ctrl+Y (Redo/Repeat)
- Ctrl+C (Copy)
- Ctrl+X (Cut)
- Ctrl+V (Paste)

To avoid problems, always use Ctrl+Shift+*letter* macro key combinations, because these combinations are much less common than the Ctrl+*letter* shortcut keys. And if you're in doubt, don't assign a shortcut key when you create a new, untested macro, especially if you think it may conflict with an Excel shortcut key. You can always assign one later if the macro proves extremely useful.

5. **In the Description text box, type a description for the macro.**

   You can use the description to help you identify macros later on. Excel also adds the description to the macro code as a series of comments.

6. **Click OK to start recording the macro.**

   At this point, macro recording begins. If the red record circle is visible on the Status bar, it changes to a blue stop square (Figure 28-3).

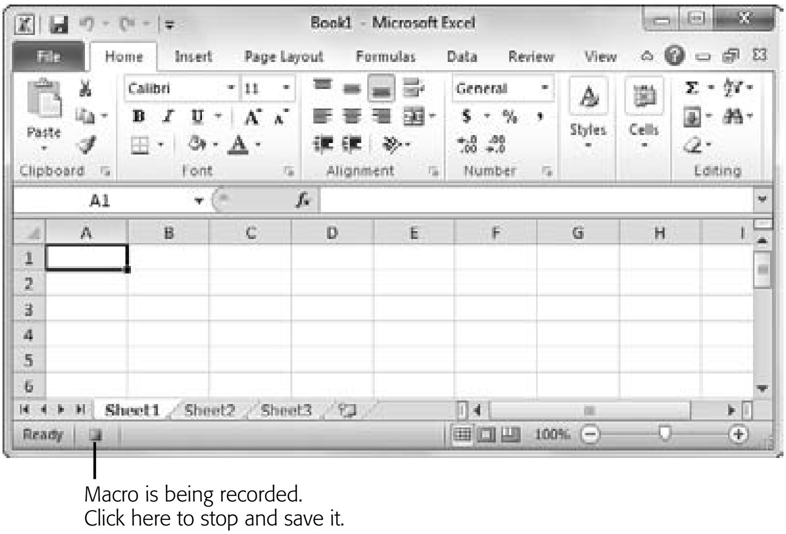

Figure 28-3: *The stop button (in the shape of a square) in a spreadsheet's lower-left corner is a visual indication that you're in record mode. If you don't see the macro stop button in the Status bar, then you can switch it on by right-clicking the Status bar, and then, in the popup menu, choosing Macro Recording.*

7. **Decide whether you want to use absolute or relative references. If you want to use relative references, make sure the View→Macros→Macros→Use Relative References button is selected. If you want to use absolute references, make sure it's *not* selected.**

   Remember, if you use absolute references, Excel always replays your actions on the same cells. If you use relative references, Excel executes your actions according to where you are at the time you launch the macro (see page 797 for a full description of the difference between the two modes).

Note: It's possible to create a macro that performs some actions in relative reference mode, and some in absolute reference mode. You just click the Relative Reference button on or off while you're recording the macro. However, this approach can lead to much confusion, so unless you're very comfortable with macros, you'll probably want to pick one system before you start recording, and then stick with it.

8. **Perform the actions you want Excel to record.**

   The macro recorder records all the actions you take, from worksheet editing and formatting to ribbon commands (like switching views, printing a document, creating a chart, and so on).

Note: The macro recorder stores code, not your specific actions. That means it doesn't matter whether you activate a feature using a shortcut key, the Quick Access toolbar, or the ribbon. In all cases, the macro code is the same.

9. **Choose View→Macros→Macros→Stop Recording to save the macro (or click the blue stop square in the Status bar).**

   Excel stops recording and quietly stores your macro. Excel doesn't show any confirmation message—the only way you can tell that the macro recorder has switched itself off is by looking at the ribbon or Status bar, where the stop button is replaced by the record button.

Playing a Macro

Once you've recorded a macro, it's even easier to play it back. Just follow these steps:

1. **Move to the location where you want to play the macro.**

   The location doesn't make a difference if your macro doesn't modify the worksheet, or if it uses absolute reference mode. On the other hand, if your macro uses relative reference mode and it modifies or formats cells, it's important to position yourself where you want to apply the changes.

   If you want to test your macro in a completely new workbook, select File→New to create a new file. Just remember to keep your original workbook open if it contains the macro you want to use.

2. **Choose View→Macros→Macros. Keyboard lovers can also use the handy Alt+F8 shortcut.**

   The Macro dialog box appears, with a list of all the macros in the personal workbook and any open workbooks (see Figure 28-4). You can filter this list (to show things like only macros in the current workbook or those in the personal macro workbook) by choosing a different option from the "Macros in" list.

Note: If you get an error message informing you that Excel has turned off the macros in your workbook (which it does automatically when you close and reopen it), you need to take a few extra steps. The next section has the scoop on macro security.

3. **Select the macro from the list, and then click Run.**

   Excel closes the Macro dialog box and runs the macro. Don't blink—99 percent of all macros replay so quickly that you can't see the changes being made. Instead, the modifications seem to appear all at once.

Caution: Once you play a macro, you're stuck with the changes. Sadly, the Undo feature can't reverse a macro, so make sure you save your worksheet before trying out an untested macro.

Figure 28-4: *The Macro dialog box shows all the macros in every open workbook, as well as those in the personal workbook. This example has two available macros. The first one, InsertHeader, is located in another workbook. Excel uses a name that starts with the name of the workbook file to indicate this fact (MacroCollection.xlsm!InsertHeader). The second macro, PrintCustom, is in the current workbook, and so the file name isn't present. The Macro dialog box also includes a Delete button that lets you remove a macro, and several additional buttons for editing and debugging macro code, which you'll learn about in the next chapter.*

Using the Macros dialog box, you can attach a shortcut key to an already created macro. Just select View→Macros→Macros, and then select the macro you want to modify. Click Options, and then type in (or change) the shortcut key. Once your macro has a shortcut key, you can trigger it by pressing the appropriate key combination, without making a return visit to the Macros dialog box.

Macro Security

Excel's macro language is surprisingly powerful. In fact, the VBA language packs enough power for expert gurus *and* hackers who want to design malicious worksheet viruses. Unfortunately, macros aren't limited to moving from cell to cell, entering information, formatting data, and so on. Instead, macros can use full-fledged VBA code (which you'll consider in the next chapter) to delete files and even lobotomize your operating system.

Note: Excel macro viruses are adept at spreading—they work by copying themselves from an infected workbook to other currently open workbooks. But only a few Excel viruses exist, and almost none of them are destructive. They may annoy you, but they aren't likely to trash your computer. You can find a catalog of macro viruses that affect Office programs at *http://tinyurl.com/lgmls*.

To keep your machine clean, the best solution is to avoid using macros in Excel spreadsheets that you don't trust. (And if you do get infected, antivirus software can help you out.) Happily, Excel's got your back—whenever you open a macro-enabled workbook, Excel automatically disables all the macros it contains. In fact, this is even true for the workbooks you create.

To see Excel's automatic security in action, try this test. Create a new workbook (or open an existing one), record a macro, save the workbook (as a .xlsm or .xlsb file), and close it. Now try opening it again. Excel automatically disables the macro you created and pops open a message bar that explains (sort of) what happened (Figure 28-5).

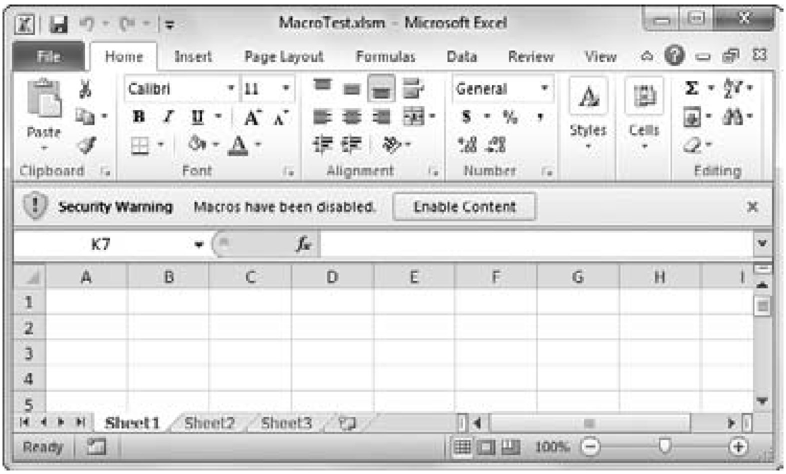

Figure 28-5: *The message bar informs you when Excel switches off your macros. To let macros run, just click the Enable Content button. Or, if you don't mind leaving your macros switched off for now, you can ignore the message bar or click the X on the right side to hide it. (You can always turn your macros back on later using the trust center, as described on page 806.)*

Trusted Documents

In the previous version of Excel, a security warning appears every single time you open a workbook with macros, data connections (page 658), or linked formulas (page 258), unless you configure Excel's security settings. Excel 2010 tries hard to be less annoying. If you click Enable Content, Excel remembers that you trust this workbook, and it won't ask you again the next time you open the file (even if you change the macro code). This feature is called *trusted documents*.

Here's how the trusted document system works. Every time you decide to trust a workbook, Excel stores the information about that document in the Windows registry. (The registry is a massive catalog of settings that every Windows computer uses.) The next time you open that document, Excel looks in the registry and discovers that you already trust it.

However, Excel *will* turn off your macros and show the message bar again if you rename the file, move it to another folder, or move it to another computer, because in these situations it can't be sure that it really is the same document. Similarly, Excel won't trust your document if you log into your computer as another user, and you can't fool Excel by replacing a trusted workbook with another file that has the same name. (This sneaky strategy fails because Excel stores other details about the file, like its size and creation date. If these don't match, Excel knows that something's up.)

Note: You may wonder why Excel doesn't wait until you try to run a macro before it warns you. Some macros can be set to run even if you don't explicitly use the View→Macros→Macros command. These macros run automatically in response to certain Excel events (like opening a document). To further complicate life, nefarious macros can sometimes hide or password-protect their code, making it difficult to find out that a problem even exists.

There's one exception to Excel's trusted document system, and it happens when you open a macro-enabled workbook from a from a network drive. When you click Enable Content on a network file, Excel turns on the macros but it doesn't automatically designate the workbook as a trusted document. Instead, it shows an additional warning message (Figure 28-6) that gives you the choice. If you're certain you want to make this workbook a trusted document, click Yes. If you want to turn on the macros for this session only, click No. All the macros will work, but Excel will switch them off again the next time you open the document.

Figure 28-6:
Network documents represent a different type of risk. Because they aren't stored on your computer, it's possible that someone else might modify the file after you decide to trust it. If you're worried about this scenario, click No to turn on the macros but refrain from designating the file as a trusted document. If you're not so paranoid, and you're opening the file from a personal or company network, you're probably safe to click Yes.

You can also give temporary trust to a macro-enabled workbook that's stored on your computer hard drive, but it's a lot less common and requires a little more work. The following section shows you how.

Temporary Trust

Usually, the trusted document system is exactly what you want. Potentially dangerous files are locked up until you check them out, and then click Enable Content. After that point, they're family, and Excel lets them run their macro code. However, it's possible that you might decide to temporarily trust a document—in other words, turn on its macros for one time only.

If this is what you want, don't click the Enable Content button. Instead, follow these steps:

1. **Click File→Info.**

   This switches you to backstage view.

2. **Click the big Enable Content button, and then choose Advanced Options.**

   A new dialog box appears, with two options: keep the macros disabled or let them run just this once.

3. **Choose "Enable content for this session," and then click OK.**

   Now the workbook can run its macros—at least until the next time you close and reopen it.

The Trust Center

If you want more control over Excel's macro security settings, you need to visit Excel's Trust Center. It gives you several options to make it even easier to work with files that contain macros:

- You can lower the Excel security settings to allow macros. This approach isn't a great idea, because it allows all code to run—even potentially malicious code in other files.
- You can tell Excel to trust the files in certain folders on your computer (or on other computers). This option is the most convenient way to go.
- You can tell Excel to trust workbooks created by a *trusted publisher*. This option's the most secure, but in order to set it up, you need to pay another company to get a security certificate. For that reason, big companies with money to burn are usually the only ones who opt for this route.

To get to the Trust Center, follow these steps:

1. **Choose File→Options.**

   The Excel Options window appears.

2. **In the list on the left, choose Trust Center.**

3. **Click the Trust Center Settings button.**

   The Trust Center window appears (Figure 28-7).

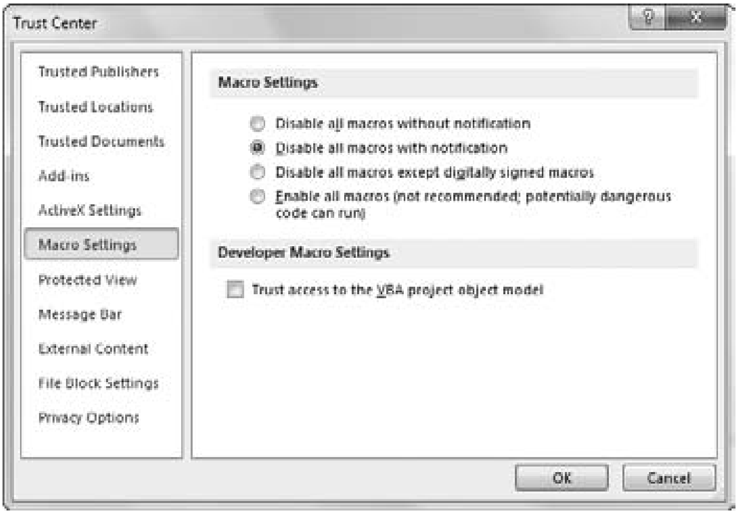

Figure 28-7:
The Trust Center window is divided into several sections. (You move from one section to another using the list box on the left.) The Macro Settings section lets you decide how Excel reacts to macros. You can choose to turn macros on or off, and you can choose whether Excel should notify you when it turns something off.

The Trust Center has several sections:

- **Trusted Publishers** lets you tell Excel to trust workbooks that are *digitally signed* by certain people. To use this feature, your company needs to buy a digital certificate from a company like VeriSign (*www.verisign.com*). Then, when you open a signed workbook, Excel contacts the company that issued the certificate and checks that it's valid. If it's valid, everything's kosher, the workbook is trusted, and Excel allows all macros. Digital certificates are outside the scope of this book, but you can learn about them at *www.verisign.com.au/repository/tutorial/digital/intro1.shtml*.

Note: If you dig around long enough, you'll discover that Microsoft has a tool (known as *MakeCert.exe*) for generating your own digital certificates. However, this tool is only for testing purposes, as the certificates it generates don't work on anyone else's computer. Beware–some Excel books and websites could lead you astray.

- **Trusted Locations** lets you pick out the places on your hard drive where you store *your* workbooks. That way, Excel knows to trust *your* workbook files but not anyone else's. You'll learn how to set up a trusted location in the following section.

- **Trusted Documents** lets you switch off the trusted document feature altogether, or switch it off for network documents only. You can also remove the current list of trusted documents (in other words, you can "untrust" everything you've trusted so far) by clicking the Clear button. However, Excel doesn't provide a list of the documents you trust.

- **Add-ins** lets you adjust whether Excel add-ins (mini programs that extend the features in Excel) should be supported even if a supported publisher didn't create them. Ordinarily, Excel allows all add-ins. (After all, if you don't trust a specific add-in, don't install it!) Only people in corporate environments, where they need to lock down Excel severely to prevent any chance of a problem, use this setting.
- **ActiveX Settings** lets you adjust how Excel treats ActiveX controls. ActiveX controls are graphical widgets (like super-fancy buttons) that developers create (or buy), and then plop into workbooks and other documents. People don't often use ActiveX controls with Excel.
- **Macro Settings** lets you configure how Excel deals with macros. You can make protection more rigorous (so that you don't let in any macros, unless they're from a trusted publisher), or less (so that all macros are allowed, no matter what they may do). By far the best choice is leaving this option at the standard setting: "Disable all macros with notification."
- **Protected View** configures the way Excel's protected view feature works. As discussed on page 57, Excel uses protected view to open unknown documents in a carefully limited window, ensuring that potentially dangerous files can't do any damage. Using these settings, you can choose whether Excel should use protected view automatically on workbooks you download from the Web, workbooks you receive in your email, workbooks you open from potentially risky folders (like the Temporary Internet Files folder), and so on.
- **Message Bar** lets you set whether Excel shows the message bar when it blocks macros in a suspicious workbook. It's best to keep this feature on, so you know when your macros have been put out of commission.
- **External Content** lets you configure how Excel deals with formulas that link to other workbooks (page 258) and data connections that pull information out of other sources, like databases (page 658). Ordinarily, Excel is a bit cautious with both these ingredients. When you open a workbook that uses these features, Excel turns off its links and data connections, and shows the security warning in the message bar. (You can then click Options, and then click Enable Content to tell Excel to switch these features back on.) If clicking the Enable Content button is sucking too much time out of your life, the best way to remedy this inconvenience is to use a trusted location.
- **File Block Settings** lets you specifically prevent Excel from opening or saving certain file types. It's up to you whether you want to prevent Excel from opening restricted file types altogether or simply force Excel to open them in protected view (page 57). The standard settings configure Excel to use protected view for Excel file formats that are older than Excel 95.
- **Privacy Options** lets you tweak a few options that aren't related to macros at all. You can choose whether Excel checks the Web for updated Help content, and whether it sends troubleshooting information to Microsoft when a problem

happens (so that Microsoft can spot bugs and learn how to improve Excel in the future). If you're paranoid about Internet spies, you may want to turn off some of these options. Most of the time, these settings are for conspiracy theorists only.

Setting Up a Trusted Location

Wouldn't it be nice to have a way to distinguish between your workbooks, which contain perfectly harmless code, and other workbooks, which may not be so nice? Excel lets you designate a specific folder on your hard drive as a trusted location. If you open a workbook stored in this location, it's automatically trusted. And if one of these workbooks contains macros, data connections, or links, Excel switches on all these features right away.

Note: Of course, it's still up to you to make sure that you place only *your* workbooks in the trusted location. If you put a potentially dangerous workbook in the trusted location, you don't have any protection when you open it.

Here's how you can set up a new trusted location:

1. **Open the Trust Center window.**

   If you're not there already, follow the steps on page 806.

2. **Select the Trusted Locations section.**

   You see a window that lists all your trusted locations (Figure 28-8). When you first install Excel, it creates a few trusted locations so it can store templates, add-ins, and other important files that it uses.

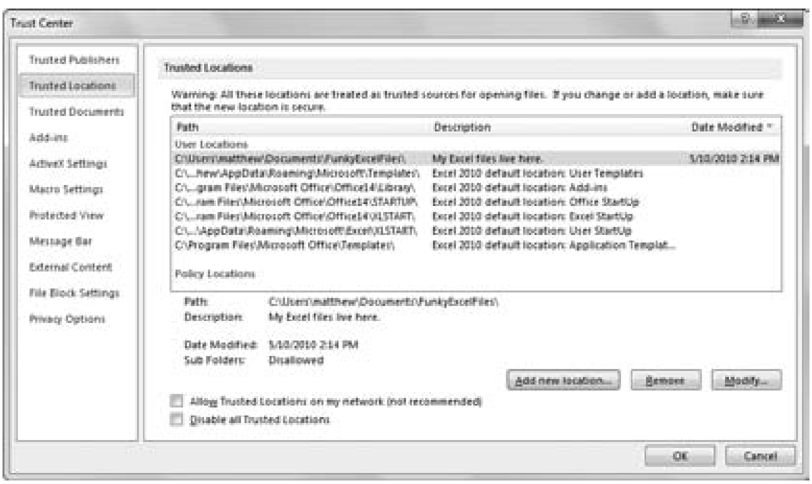

Figure 28-8: *In this example, there's only one user-added trusted location (which is selected here). It's the FunkyExcelFiles folder in the hard drive's My Documents section.*

3. **If you want to trust a specific folder on your company or home network, turn on the "Allow trusted locations on my network" checkbox before you go any further.**

   You're taking a bit more of a risk when you turn this setting on, because a network location is out of your control. A hacker could sneak a virus-laden workbook into that location without you noticing. However, if you're reasonably certain that the network's secure (and the other people who use the folder aren't likely to download workbooks from the Web and place them there), you probably don't need to worry.

4. **Click "Add new location."**

   Excel asks you to fill in a few pieces of information, as shown in Figure 28-9.

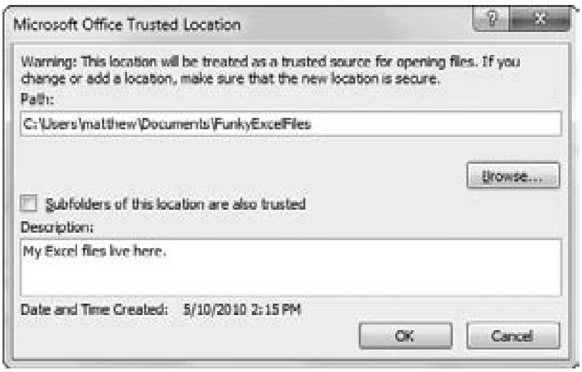

Figure 28-9:
To configure a trusted location, you need to specify the path (click Browse to hunt it down). You can also choose whether all this folder's subfolders should be trusted automatically, and you can fill in an optional description that appears in the list of trusted locations.

5. **Click OK to add the location to the list.**

   You can configure the location or remove it at any time by selecting it in the list, and then using the clear-as-a-bell Remove and Modify buttons shown in Figure 28-8.

Creating Practical Macros

If you're still wondering what macros can do for you, it helps to look at a few straightforward examples. In this section, you'll learn about some helpful macros that you can create for practice and profit.

POWER USERS' CLINIC

Authenticode Macro Signing

If your company frequently creates and distributes workbooks with macros, you may want to use an advanced option called Authenticode signing. With Authenticode, you sign all your macros with a digital signature. When someone opens the workbook, their computer checks the signature, and verifies that it's from a known, trusted source. If it's not from a trusted source, the Security Warning dialog box appears (if Excel's using the standard Medium security level), or Excel turns the macro off (if it's using High or Very High security).

Authenticode macro signing is problematic, because the person who opens the document needs a way to verify that the digital signature's trusted. You need to create a *digital certificate* for the person who writes the macro, and register this digital certificate with everyone who needs to use macros created by this person.

This process can be complicated, and large organizations that use macros heavily go through it. Authenticode signing is far outside the scope of this book, but you can download a document that explains this feature and the technology it uses from *http://office.microsoft.com/home.* Just search for *macro security whitepaper*.

Inserting a Header

A really simple macro to start out with is one that inserts some boilerplate information into a group of cells. Consider the header at the top of a company spreadsheet, which typically includes a company's name, the spreadsheet's author, the date the spreadsheet was created, and a title. If you're often adding this same information to many different worksheets, you can automate the work with a macro.

Here's what you need to do:

1. **Fire up the macro recorder (select View→Macros→Macros→Record Macro).**

   The Record Macro dialog box appears.

2. **Name your macro.**

   Choose a memorable name (like InsertHeader), and then click OK to get started with the actual macro recording.

3. **Make sure you're in absolute reference mode (the View→Macros→Macros→Use Relative References button should *not* be selected).**

   You want absolute reference mode because you always want Excel to place the header in the same few cells at the top of the worksheet.

4. **Move to cell A1, and then insert a generic title (like "Sales Report").**

   You can of course pick whatever title you want to use.

5. **Select cells A1 to C1, right-click the selection, and then choose Format Cells.**

   In the next step, you're going to create a little more breathing room for your title by increasing its size.

6. **Change the font to something large and dramatic.**

   In the Alignment tab, turn on the "Merge cells" checkbox to group these three cells into one larger cell that can accommodate the title. Click OK once you're done.

7. **Move to cell A2 and enter the text "Created". Move to cell B2, and then type the formula =*TODAY()*.**

   The TODAY() function inserts the current date. However, you want this cell to reflect the creation date, so that Excel doesn't update the date every time someone opens up the spreadsheet. To make sure the date can't change, you need to replace the formula with the calculated date value. Press F2 to put the cell in edit mode, and then F9 to replace the formula with the result. Press Enter to commit the change.

8. **Optionally, add any extra text you'd like to include (like a company slogan or a copyright message).**

   Feel free to tweak column sizes as well.

9. **Choose View→Macros→Macros→Stop Recording to save your macro.**

   To try out your macro (Figure 28-10), replay it in a different worksheet in the same workbook.

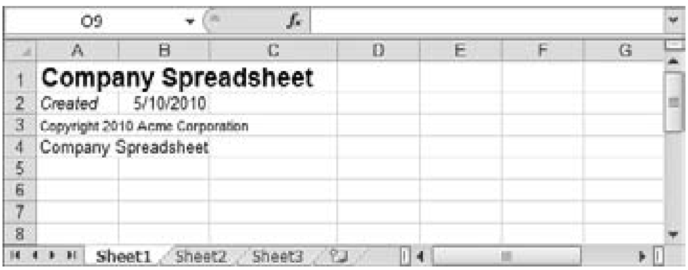

Figure 28-10: *The InsertHeader macro (which you can download from the Missing CD page at* www.missingmanuals.com/cds*) inserts a generic header at the top of a worksheet and automatically grabs the current date, and inserts it as text.*

Alternating Row Formatting

A custom macro is a useful tool for complex and unusual formatting tasks. Macros are particularly useful if you want to apply a pattern of formatting to a large block of cells. For example, if you want to format every third row with a different background color, a macro can come in quite handy. In this case, the easiest approach is to create a macro that formats just a few rows; you can then replay this macro multiple times to format a whole table of data.

Here's how to build a macro for applying alternating row formatting:

1. **Move to the cell where you want to start before you begin recording.**

   In this example, it makes sense to use relative reference mode. Therefore, it doesn't matter which cell you start from; Excel doesn't record the cell address as part of the macro.

2. **Choose View→Macros→Macros→Record Macro, and then choose a name you'll remember (like FormatRow).**

   You can also specify a shortcut key (like Ctrl+Shift+F). The FormatRow macro described here is much more practical with a shortcut key, because you'll want to use it multiple times in quick succession. When you're finished, click OK.

3. **Make sure you're in relative reference mode.**

   Click the View→Macros→Macros→Use Relative References button so that it's highlighted. You want this button selected since you're going to repeatedly apply this macro across your worksheet.

4. **Click the row number button at the left of the current row to select the entire row. Right-click the selection, and then choose Format Cells.**

   The Format Cells dialog box appears.

5. **Select the Patterns tab, choose a new color, and then click OK.**

6. **Press the down arrow key twice to move down two rows.**

   If you're in row 1, move down to the third row.

7. **Choose View→Macros→Macros→Stop Recording to save your macro.**

The neat thing about this macro is that you can replay it multiple times to format unlimited expanses of data. To try this trick out, start at the top of a table of data. Press the macro shortcut key (if you've created one) to format the first row and automatically move down two rows. Then press the macro shortcut key again to format the current row and move down again. (If you haven't used a shortcut, you'll have to use the more awkward approach of selecting View→Macros→Macros, and then selecting the macro from the list of available macros.) You can continue this process for as long as you want. The end result is that you'll have created a table where every second row is highlighted with a different background color.

A Combined Task

Really sophisticated macros don't just add data or format cells. Instead, they perform a whole *series* of operations. To get a better feeling for these types of macros, you can try out the following example, which figures out the average and median values for a list of numbers.

1. **Before you start recording the macro, create a list of numbers in an Excel spreadsheet. These numbers are the ones the macro will analyze.**

   If you don't want to type in your own numbers, you can open an existing worksheet, like the student grade worksheet from Chapter 9.

2. **Select the numbers.**

   Note that you select the numbers *before* you create the macro. That's because the macro you're creating is designed to analyze a selection of cells. Before you replay the macro, you'll need to select a range of cells you want to use.

3. **Choose View→Macros→Macros→Record Macro.**

   Choose a suitable name (like AnalyzeSelectedCells), and then click OK.

4. **Before you take any other steps, define a new name for the selected cells.**

   In the formula bar's Name box, type AnalyzedCells. (The name box is at the extreme left side of the formula bar, and it usually displays the reference for the active cell.)

5. **Press the right arrow key once or twice to move to a free column.**

6. **Enter the formula *=AVERAGE(AnalyzedCells)* to calculate the average of the named range you created in step 4.**

7. **In a cell underneath, enter the formula *=MEDIAN(AnalyzedCells)*.**

   This formula calculates the median value of the named range you created in step 4.

8. **Add labels next to the cells with the two formulas, if you want.**

9. **Chose View→Macros→Macros→Stop Recording to save your macro.**

To test this macro, select the group of cells you want to analyze before you play the macro. The neat thing about this macro is that you can use it to analyze any number of cells, just as long as you select them all. The only limitation is that you can have just one range with the same name in a worksheet, so you can't use this macro in more than one place on the same worksheet. If you wrote the same macro by hand in VBA code, you could circumvent this limitation.

Placing a Macro on the Quick Access Toolbar

Once you've created a useful macro, you may want to attach it to the Quick Access toolbar at the top of the Excel window so that it's conveniently available when you need it. You'll especially want easy access if you're creating a macro that you're frequently going to use (and it also helps if you've saved the macro in the personal macro workbook, since that ensures that you can use that macro from any workbook, as explained on page 798).

Note: If you attach a macro to the Quick Access toolbar and the macro isn't in your personal macro workbook, you could run into trouble. If you rename or move the workbook that contains the macro later on, the button won't work anymore.

Adding a Quick Access toolbar button that activates macros isn't difficult. Just follow these steps:

1. **If the macro isn't in the current workbook, start by opening the workbook that contains it.**

   Ideally, the macro is stored in your personal macro workbook. (If you want to copy a macro from another workbook into your personal workbook, jump ahead to page 824 in the next chapter.)

2. **Right-click the Quick Access toolbar, and then choose Customize Quick Access Toolbar.**

   The Excel Options dialog box appears, with the Customize section chosen (see Figure 28-11).

3. **In the "Choose commands from" drop-down list, choose Macros.**

   Underneath, a list appears with all the macros available in the personal macro workbook and any other currently open workbooks.

4. **Choose the macro you want in the list, and then click Add to place it in the Quick Access toolbar.**

   The item appears in the box on the right, which shows the current list of Quick Access commands.

5. **Choose the newly added macro (at the bottom of the list), and then click Modify.**

   The Modify Button dialog box appears.

6. **Choose a new icon and display name.**

   Initially, all macros use a hideously bizarre icon and a tongue-twisting display name that includes the name of the workbook where they're stored. You can do better.

   The icon's the miniature picture that appears in the Quick Access toolbar. Pick one of the ready-made icons, all of which look better than the stock macro icon.

   The display name is the pop-up text that appears when you hover over the icon in the Quick Access toolbar. Try using a more understandable name. Instead of SuperMacroWorkbook.xlsm!MySuperMacro, consider "My Super Macro" or "MySuperMacro from SuperMacroWorkbook" or "Formats Alternating Rows".

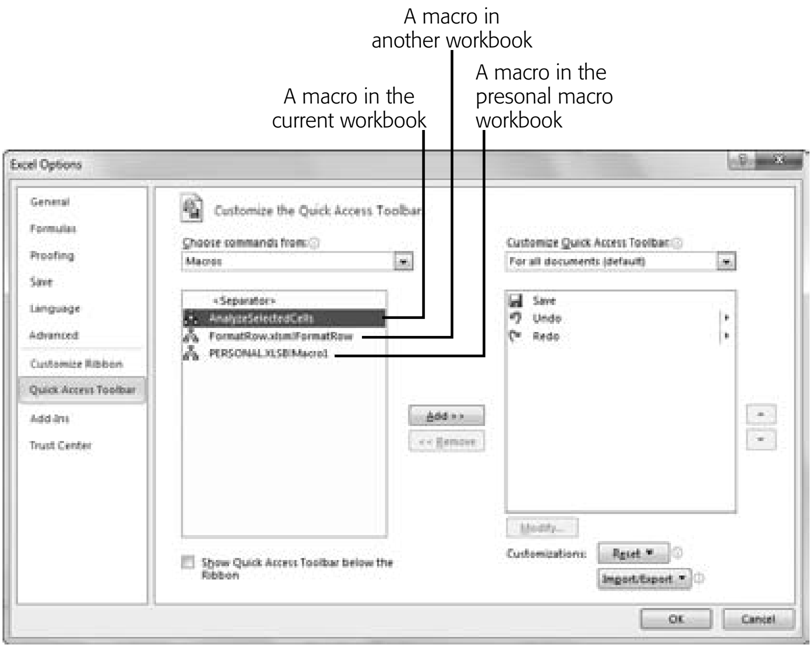

Figure 28-11: *Choose the category of commands you want to see by picking from the "Choose commands from" list. The right side of the dialog box shows the Quick Access toolbar's current contents. You can move items on or off the toolbar using Add and Remove, and you can change the order of items using the arrow buttons.*

7. **If you don't like your macro button's positioning, select it, and then use the up and down arrow buttons.**

   When you add a new item to the Quick Access toolbar, it heads straight to the end of the list (which means it appears at the toolbar's right edge). If you want, you can move it to a more prominent position.

8. **Click OK.**

   Your macro appears in the Quick Access toolbar. If you tire of it, right-click the button, and then choose Remove from Quick Access Toolbar.

Tip: When you customize Excel's Quick Access toolbar, Excel stores your customizations on your computer, but not in any workbook file. People who use your workbook files on other computers don't see your new buttons. They need to rely on the Macros dialog box or the macro shortcut key. You can get around this limitation by attaching a macro to a worksheet button instead, as described in the next section.

Attaching a Macro to a Button Inside a Worksheet

Only a few extremely useful macros will ever be worth space on the Quick Access toolbar. But if you create a macro that you use frequently with a specific workbook—

say, a macro that produces a special printout or performs a complex calculation—you may want it easily available all the time, but for that workbook only.

The solution? Use a *button control*. You can place this button anywhere you want on your worksheet and configure its text. When the person using the spreadsheet clicks the button, the linked macro runs automatically. This property is particularly useful if a lot of different people use the worksheet; not everyone can remember a specific macro name or shortcut key, but nobody has any trouble clicking a large inviting button.

To add a button to a worksheet, you need the Developer tab (Figure 28-11). The Developer tab doesn't appear unless you ask Excel to show it. To do so, choose File→Options, and then choose the Customize Ribbon section. On the right of the window, a large box lists all the tabs that are currently shown in the ribbon. Near the bottom, you'll see an unchecked item named Developer. To show the Developer tab, check this box, and then click OK.

The Developer tab has buttons for recording and playing macros (which duplicate the options from the View→Macros section of the ribbon), along with more advanced programming commands, most of which you'll never use unless you become a hard-core code jockey. Although you can record and play macros with Status bar buttons, it's a good idea to show the Developer tab before you continue any further. The Developer tab includes a few more options, like the ability to control whether the macro records relative or absolute references (page 797), and a way to manage Excel's macro security settings (page 803).

Once you've shown the Developer tab, you're ready to add a button to your worksheet and attach a macro to it. Just follow these steps:

1. **If the macro isn't in the current workbook, start by opening the workbook that contains it.**

   Ideally, the macro is stored in the same workbook where you're placing the button or the personal macro workbook. Otherwise, you're probably complicating your life unnecessarily.

2. **Choose Developer→Controls→Insert→Button (*not* the similarly named Command Button).**

   When you choose Developer→Controls→Insert, a list of different controls appears (Figure 28-12).

3. **Drag to "draw" the button onto your worksheet.**

   There's no restriction as to where you can place a button, or how large it can be. However, you don't want to obscure important data in the cells underneath. Once you finish drawing the button, the Assign Macro dialog box appears, with a list of available macros.

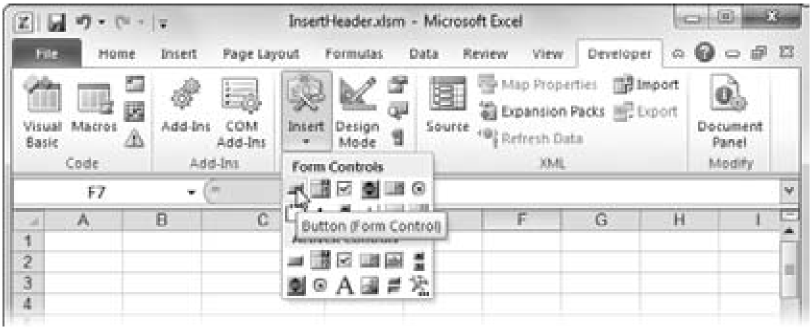

Figure 28-12: *If you don't spot the Button icon right away, just move your cursor over all the icons until the tooltip text Button appears. Then click it.*

4. **Select a macro, and then click OK.**

   Ideally, you should choose a macro that's stored in the current workbook or the personal workbook. Otherwise, the button won't work if you move or rename the required workbook.

5. **Right-click the button, and then choose Edit Text. Replace the standard text (Button 1) with something more descriptive (like Update Totals).**

6. **Click the worksheet.**

   Click anywhere on the worksheet to finish this procedure. At this point, the button is fully functional, and you can click it to run the linked macro (see Figure 28-13).

   If at any point you want to change the button text or move it somewhere else, start by right-clicking the button to select it. You can then drag it, resize it, delete it (by pressing the Delete key), or click to change the button text, without inadvertently triggering the linked macro.

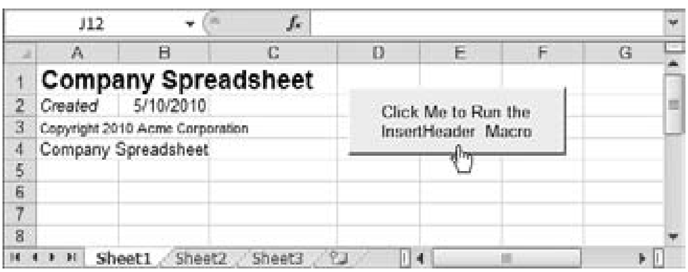

Figure 28-13: *This worksheet includes a custom button that runs a linked macro. When you move the mouse pointer over the button, it changes into a hand, indicating that you can click the button to unleash the macro.*

CHAPTER

29

Programming Spreadsheets with VBA

When you use Excel's macro recorder (covered in the previous chapter), you actually build a small program. Excel translates every action you take, from selecting a menu item to editing a cell, into a line of *VBA* code (short for Visual Basic for Applications), and inserts it into your new macro. The obvious benefit is that you can use the macro recorder without knowing the first thing about programming. However, this free pass has limitations.

You'll find that you can record only actions that you can perform yourself. If you want to create a macro that inserts a column with 100 identical cell values, you need to go through the drudgery of typing in each cell value so that the macro recorder knows what to capture. A more serious problem is that when you record a macro, you can respond only to the worksheet that's in front of you. If you want to make a more flexible macro that has the ability to examine a variety of cells and the intelligence to respond to different conditions, you need to tap into some of VBA's more advanced capabilities.

Sooner or later, every Excel guru delves into macro code. Often, your first encounter with macro code happens when you need to enhance an existing recorded macro to make it more nimble or to correct a problem. Depending on your preference, these encounters may be the only time you come face to face with VBA, or it could be the start of a new career path as an Excel programmer extraordinaire.

There's no way to explain a complete programming language in one chapter. Instead, this chapter walks you through the Visual Basic editor, explains some important code techniques, and shows a few useful macros. By the time you finish this chapter, you'll know enough about VBA to grab a useful snippet of macro code from an Excel website (and understand what it's supposed to do), or dive into an advanced book that's dedicated to VBA programming.

The Visual Basic Editor

Before you can modify a macro, you first need to find it. The tool you use to edit macros isn't actually part of Excel. Instead, it's a separate program called the Visual Basic editor.

The easiest way to get to the Visual Basic editor is to use the Developer tab. The only problem is that Excel likes to hide its high-powered coding features from ordinary people. To ask Excel to show the Developer tab, you need to choose File→Options, and then choose the Customize Ribbon section. On the right of the window, a large box lists all the tabs that are currently shown in the ribbon. Near the bottom, you'll see an unchecked item named Developer. Check this box, and then click OK to reveal the Developer tab. Now you can switch the Visual Basic editor window at any time by choosing Developer→Code→Visual Basic. When you do, Excel launches the standalone window shown in Figure 29-1.

Tip: Alternatively, if you've already created a macro, you can jump straight to its code using the familiar Macros dialog box. Just choose View→Macros→Macros, select the macro in the list, and then click Edit.

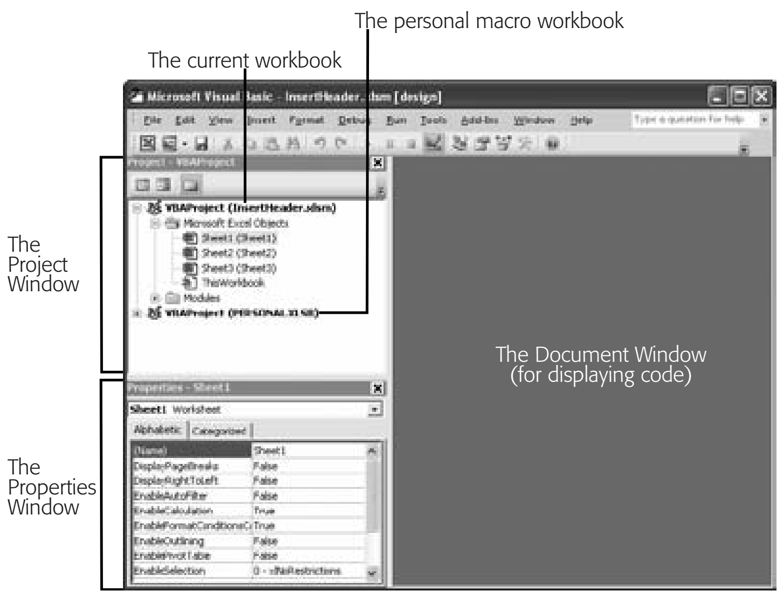

Figure 29-1: *The Visual Basic editor window is divided into three main regions: the Project window, the Properties window, and the Document window. The Document window is where your code appears (it starts off blank). The Project window lists all your projects (two are shown here: InsertHeader.xlsm and the personal macro workbook). The Properties window shows individual settings you can change for the currently selected project.*

The Project Window

The best way to orient yourself in the Visual Basic editor window is to start by looking at the Project window in the top-left corner. This window shows a directory tree representing all the currently open workbooks. Each open workbook becomes a separate VBA project, which means that it contains its own set of macro code.

Along with the workbooks you currently have open, you may also find a few workbooks that you don't recognize in the project window. Don't panic—these interlopers are completely kosher. They correspond to *hidden workbooks*—workbooks that Excel opens but doesn't show you or let you edit. Often, a hidden workbook loads when you activate an add-in like Solver (page 604). That's because all the code and specialized functions that the add-in provides are actually written using VBA *modules* (code files). A little later in this chapter, you'll see how you can use this technique to not only create macros, but also to define your own specialized functions.

Some of the unexpected workbooks that you may find in the project window include:

- **PERSONAL.XLSB**. This project is the personal macro workbook (page 798) that contains the macros you share between all workbooks. If you haven't recorded any macros to the personal macro workbook, you don't see this project because Excel hasn't created it.
- **SOLVER.XLAM**. This project contains the code for the Solver add-in, if you've turned it on (page 604).
- **SUMIF.XLAM**. This project contains the code for the Conditional Sum wizard add-in, if you've turned it on.

Note: Don't try to open the add-in projects. They're password-protected so that you can't modify any of the code that they contain (or even look at it). However, you can change the macros you create, as well as your personal macro workbook.

Each workbook project in the Project window contains one or two folders. The first folder, Microsoft Excel Objects, is always present and contains a separate entry for each worksheet in the workbook. You use these entries to attach code that reacts to a specific worksheet event. You could create a task that springs into action every time the person using the workbook selects this worksheet. The Microsoft Excel Objects folder also contains a single workbook object—called ThisWorkbook—where you can respond to workbook actions (like saving, opening, or printing the workbook).

Note: Creating code that reacts to events is an advanced Excel technique. For example, you could use an Excel event to run some code every time your workbook is opened. This book doesn't cover events.

Modules is the second folder in a workbook project. It appears only if you've created a macro for this particular workbook. Figure 29-2 details the action inside one particular Project window.

Note: As you open new workbooks, they automatically appear in the Visual Basic editor's Project window. Similarly, as you close workbooks, they disappear from the Project window.

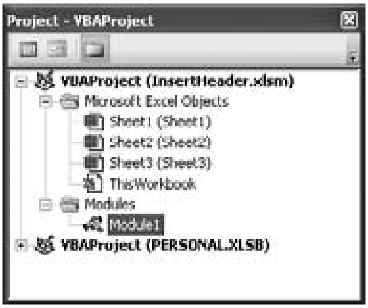

Figure 29-2:
This window has two VBA projects. The first, InsertHeader.xlsm, is an Excel workbook file with three worksheets and one module for macro code. The last, PERSONAL.XLSB, is the personal macro workbook, which may also contain worksheets and macro modules (although that branch of the tree is collapsed, so these objects aren't shown).

Modules and Macros

Modules contain macro code. Ordinarily, Excel creates a new module the first time you record a macro, and names it Module1. Excel then places every macro you record into that module. If you want, you can separate macros into different modules for better organization. However, the number of modules you create, or the module you use for a given macro, has no effect on a macro's function.

Every module contains one or more VBA *subroutines*. Each subroutine is a named unit of code that performs a distinct task. In the VBA language, subroutines start with the word Sub followed by the name, like all VBA commands. They end with the statement End Sub. Here's an example:

```
Sub MyMacro
    ' Your macro code goes here.
End Sub
```

This small snippet of VBA code shows two important principles. First, it shows you how to start and end any subroutine (by using the statement's Sub and End Sub). This code also shows you how to create a *comment*. Comments are special statements that Excel ignores completely; they're notes to yourself (like explaining in plain English what the following or preceding line of code actually does). To create a comment, you just place an apostrophe (') at the beginning of the line.

Tip: Master programmers always leave comments in their code (or at least they feel terribly guilty when they don't). Comments are the best way to clarify what you want your code to do, so you can remember your intentions when you review it a few months later.

In Excel, each macro is a separate subroutine. (In some cases, you may want to break a complex macro down into more than one subroutine, but the macro recorder doesn't do this for you.) When the macro recorder goes to work recording a new macro, it generates a new subroutine using the name of the macro that you assigned. It also adds any description you entered when you were creating the macro. Then, it places all the code it generates into the subroutine.

Here's the basic skeleton for the InsertHeader macro created in the last chapter:

```
Sub InsertHeader()
    '
    ' InsertHeader Macro
    ' Macro recorded 3/6/2010 by Matthew MacDonald
    ' (Code goes here.)
End Sub
```

To take a look at the subroutines in a module, double-click the module in the Project window. The Module opens in a new window, as shown in Figure 29-3. You can scroll through this window to see all the macro procedures it contains.

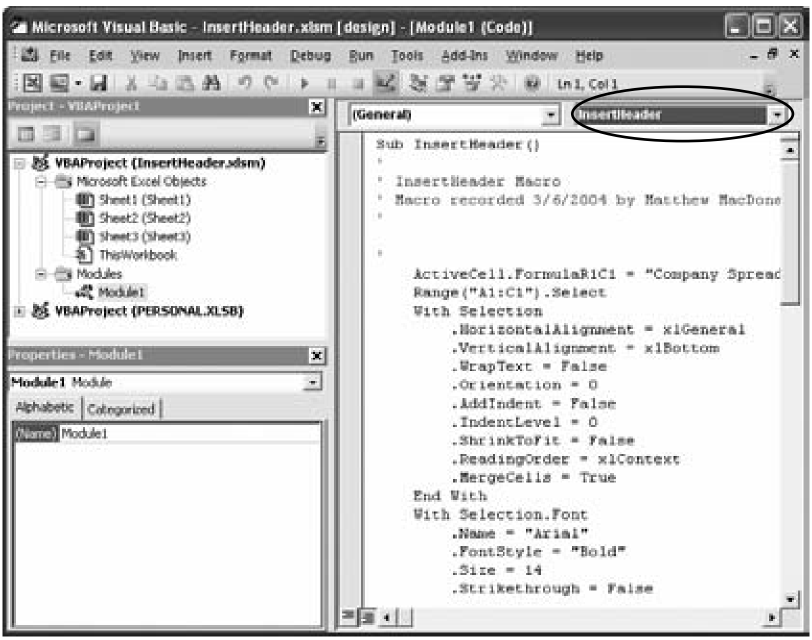

Figure 29-3: *In this example, the Visual Basic editor shows a portion of the code for the InsertHeader macro. If you scrolled down through the window, you'd see additional macros in separate modules if this workbook contained any. Between each subroutine is a solid line (not shown). You can also jump directly to the subroutine that interests you by selecting the subroutine's name from the list box in the top-right corner of the code window (circled).*

Finding and Moving Macros

With your new knowledge of VBA, you can transfer a macro from one workbook to another. Just follow these steps:

1. **Open both workbooks in Excel.**

   Before you start, you need to open the workbook that contains the macro you want to move or copy (that's the *source workbook*), and the workbook where you want to move the macro to (the *target workbook*).

2. **Switch to the Visual Basic editor window.**

   If it's not already open, use the View→Macros→Macros→View Macros command, pick a macro, and then click Edit.

3. **Using the Project window, expand the source workbook until you see the module where the macro's stored (this module's usually named Module1). Double-click this module.**

   The code window appears with the macro code for the source workbook.

4. **Scroll through the code window until you find the macro you want to copy or move. Select it all, from the starting Sub line to the concluding End Sub statement. Press Ctrl+C to copy, or Ctrl+X to cut the macro code.**

   If you use Ctrl+X, you remove the macro code from the source workbook, and the macro's no longer available in that workbook.

5. **Using the Project window, expand the target workbook. If the target workbook doesn't already contain a macro module, then you need to insert one. To do so, right-click the project name (like "VBA Project (NewWorkbook.xlsm)"), and then choose Insert→Module.**

   When you add a new module, it appears in the project window with a generic name (like Module1).

6. **Double-click the new module.**

   A blank code window appears.

7. **Click inside the code window, and then press Ctrl+V to paste in your macro code.**

   This action places the macro code into the new workbook. From this point on, you can use the macro whenever this workbook is open.

Debugging a Macro

The Visual Basic editor isn't just for editing. You can also use it to *run* any macro. Just open the module that has the macro you want, scroll through the code window, and then click to place the cursor somewhere inside the macro. Now, on the Visual Basic toolbar, click the Run button (which looks like the play button on a VCR control). If you have a long-running macro, you can click the pause button to temporarily halt the code, or the stop button to abort it altogether.

Note: You may assume that Excel performs the macro on the workbook that *contains* the macro code, but this scenario isn't necessarily true. Instead, Excel runs the macro on the *active workbook*, which is the one that you looked at *most recently*. In other words, if you have two workbooks open and you use the Windows task bar to switch to the second workbook, and then back to the Visual Basic editor, Excel runs the macro on the second workbook.

Macro debugging is another neat trick you can perform from the Visual Basic editor. This feature lets you walk through your macro code and run it one line at a time. After each line executes, you can take a look at what's happened so far in the workbook. Macro debugging is a great tool for finding mistakes or diagnosing bizarre behavior. It's also a good way to learn what each statement in a macro actually does.

To debug a macro, follow these steps:

1. **Find the macro you want to debug, and then click anywhere inside the code.**

   Excel's debugging tools aren't limited to misbehaving macros. You can use them to watch any macro in action, just for the fun of it.

2. **From the Visual Basic editor menu, select Debug→Step Into, or press F8.**

   Excel highlights the first line of the macro (the Sub statement) in yellow, and places an arrow next to it. The arrow indicates that this line is ready to run.

3. **Press F8 to run the highlighted line.**

   Once Excel runs this line, it moves to the next line (skipping any comments), and highlights it in yellow. Then it moves you to the first actual line of code in the macro, as shown in Figure 29-4.

4. **Press F8 to run the highlighted line.**

   Excel runs the highlighted line, moves one line down to the next line, highlights it, and waits for your command. At this point, you can switch back to the workbook to see if the first line of code produced any visible effect. Be careful not to change anything in your worksheet, however. If, for instance, you clear the current selection or move to a different place in the worksheet, then you could throw the macro off completely.

   When you're finished looking at the worksheet, switch back to the Visual Basic editor.

5. **Return to step 4, and keep repeating it for each line in the macro.**

   At any point, you can stop stepping through your code and run everything that's left by clicking the play button, or you can cancel the remainder of the macro commands by clicking the stop button. When you reach the last line, the macro ends.

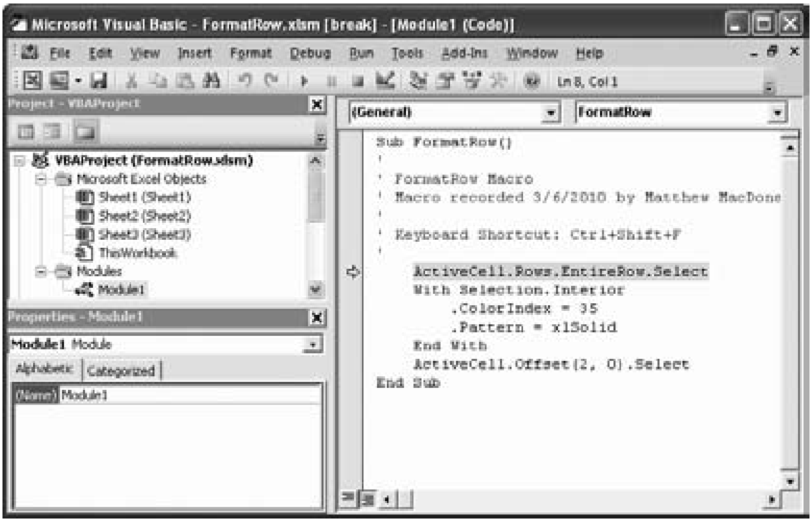

Figure 29-4: *When you press F8, Excel runs the line that's highlighted in yellow and marked with an arrow. Then it moves to the next statement and pauses again.*

Tip: Excel gives you a way to jump right to a macro that interests you. From the Excel window, choose View→Macros→Macros→View Macros to show the list of available macros. Select a macro, and then click Step Into to open it in the Visual Basic editor and begin debugging it, running one line at a time.

Understanding Macro Code

Now that you've learned how macros are organized, it's time to dive in and tackle the actual macro code. In the following sections, you'll take your first look at VBA code and learn how Excel uses a special programming tool called *objects*. (Objects are programming tools that you use to trigger Excel features. You'll learn more about them on page 828.)

The Anatomy of a Macro

A good place to start learning about the guts of a macro is the FormatRow macro shown in the previous chapter, which automatically highlights a row with a light green background, and then moves down two rows.

Following is the complete VBA code for the FormatRow macro. To make it easier to analyze, each line of code has been numbered, and the code has been simplified a little from what Excel generated automatically:

```
Sub FormatRow()
    '
    ' FormatRow Macro
    ' Macro recorded 3/6/2010 by Matthew MacDonald
```

```
5      '
6      ' Keyboard Shortcut: Ctrl+Shift+F
7      '
8      ActiveCell.Rows.EntireRow.Select
9      Selection.Interior.ColorIndex = 35
10     Selection.Interior.Pattern = xlSolid
11     ActiveCell.Offset(2, 0).Select
12 End Sub
```

Line 1 starts the macro and defines its name. Lines 2–7 are simply comments. They appear in green writing in the editor and don't actually do anything (other than convey information to the person reading the code).

The action gets started with line 8. Here, the code accesses a special object called ActiveCell, and uses it to select the entire current row. Next, line 9 changes the background color of the selected cells (to light green), and line 10 sets the type of background fill (to solid). Both of these lines use the Selection object.

Finally, line 11 returns to the ActiveCell object, and uses its Offset command to jump down two rows from the current cell. Line 12 marks the end of the macro code routine.

POWER USERS' CLINIC

Adding New Modules

Excel lets you use the macros that are stored in every module of your workbook project. That means it doesn't matter how many modules you use, or even what name each module has. Excel can find and use your macro code just as easily no matter what the arrangement.

But you could choose to use multiple modules in the same workbook project to help organize your code. This approach makes sense if you have a lot of macros in one place, and you want to organize them by function so they're easier to edit and maintain in the future. You could, for instance, create a MyMacroCollection.xlsm workbook that has three macro modules: QuickCalculations, Formatting, and Printing.

You can easily add new modules. All you need to do is right-click the name of your project in the Project window, and then choose Insert→Module. Once you've created a new module, you can follow the steps shown earlier to transfer a macro from one module to another.

Of course, if you create multiple modules with generic names (like Module1, Module2, and so on), you'll have a hard time keeping track of what's what. So, once you've created your module, change the name to something more descriptive by selecting the module in the Project window and looking at the Properties window underneath it. In Figure 29-4, the Properties window shows one entry, called (Name). You can click here to replace the current name with something more specific. Module names have all the same restrictions as macro names, which means they have to start with a letter, and use only letters, numbers, and the underscore.

If you end up with a module you don't want, it's easy enough to remove it. Just right-click the module in the Project window, and then choose Remove. (At this point, the Visual Basic editor asks if you want to export your macro code to a separate file so that you can import it back into some other workbook. If you just want to get rid of your module for good, choose No.)

Altogether, this macro doesn't do much, but the code is quite dense, and fairly difficult to read on first sight. The problem isn't the VBA language—in fact, the only language-specific details in this example are the Sub and End Sub statements (not to mention all those odd periods, which are explained in the next section).

The real complexity comes from understanding all the different objects that are available. To write a macro like this one, you need to know that there's an ActiveCell object that lets you select rows and a Selection object that lets you adjust their formatting. These details (and many more that aren't shown in this macro) make up Excel's *object model*. If you want to perform any task in a macro, from printing a document to saving a worksheet, you need to first figure out which object can do your bidding.

Objects 101

In many programming languages, including VBA, everything revolves around objects. So what exactly is an object?

In the programming world, an object is nothing more than a convenient way to group together some related functionality. In the FormatRow macro, two objects are front and center: one named ActiveCell, and one named Selection. The ActiveCell object bundles together everything you may want to do with the current cell, including editing, selecting, and moving from cell to cell. The Selection object offers other features for modifying a group of selected cells, including ways to change their borders, background colors, and font.

Programmers embraced objects long ago because they're a great way to organize code (not to mention a great way to share and reuse it). Objects also make it easier for developers to discover new features. For example, once you learn about ActiveCell, you immediately know what object you need to use for any task related to the current cell.

You can use objects in a number of different ways. Altogether, you interact with objects in three ways:

- **Properties**. Properties are pieces of information about an object. You change properties to modify the object or how it behaves. The FormatRow macro uses the ColorIndex property to change the background color of a row.
- **Methods**. Methods are actions you can perform with an object. The FormatRow macro uses the Select method to select the current row.
- **Events**. Events are notifications that an object sends out that you, as the macro programmer, can respond to. This chapter doesn't look at events, but you can use them to react to certain actions that someone using your worksheet takes (like saving a workbook).

In the next section, you'll take a closer look at how to use properties and methods.

Using Properties and Methods

So how do you change properties or use methods? The answer is the lowly period. Imagine you have a Car object that provides a StartIgnition method. In this case, you use the following syntax to start your engine:

```
Car.StartIgnition
```

The same technique works with properties, but it tends to look a little different. With properties, you typically want to perform one of two actions. Either you want to retrieve information about the property, or you want to change the property. To change the property value, you use the equal sign (=). The following line of code changes the number of people in a car by modifying the Passengers property:

```
Car.Passengers = 2
```

Note: Think of the equal sign as an arrow pointing to the left. It takes whatever information's on the right side (in this case, the number 2) and stuffs it in whatever receptacle is on the left side (in this case, the Passengers property of the Car object).

Surprisingly, this information is just about everything you need to know about objects. Of course, if you look back at the FormatRow example, you'll notice that none of the lines look quite as simple as the previous two Car object examples. That's because you need to use multiple objects at once to accomplish many Excel tasks. Excel is so feature-laden that its developers needed to divide and subdivide its features into dozens of different objects to make them more manageable.

Consider the following statement (line 9 from the FormatRow macro):

```
Selection.Interior.ColorIndex = 35
```

In this example, two objects are at work. The Selection object contains another object named Interior. Technically, Interior is a property of the Selection object. Unfortunately, the Selection object doesn't give you any way to change the background color of the selected cells. To do that, you need to use the ColorIndex property, which is a part of the Interior object. That's why this statement has two periods. The first one accesses the Interior object, and the second period accesses the ColorIndex property.

To make life even more interesting, consider line 8, which has *three* periods:

```
ActiveCell.Rows.EntireRow.Select
```

In this case, the ActiveCell object has a property called Rows, which is also an object. The Rows object has a property named EntireRow, which is yet another object. The EntireRow object is the one you want—it provides the Select method that highlights the current row.

In short, understanding the basics of objects really isn't that difficult, but you could spend days wandering around Excel's family of objects trying to find the ones you need. Excel includes objects that represent worksheets, workbooks, cells, ranges, selections, charts, the Excel window, and more. The next section tells you how to get started.

UNDER THE HOOD

The With Statement

Once you find the right object, you'll probably need to use several of its properties or methods. To save the effort of typing in the object name each time, you can use a *With block*. A With block starts off by identifying an object that you want to use. The following statements (up until the final End With) don't need to include the object name. Instead, they can skip straight to the period, and use it to start the line.

The FormatHeader macro uses the following statements to set the formatting of the current selection:

```
Selection.Interior.ColorIndex = 35
Selection.Interior.Pattern = xlSolid
```

You can rewrite this statement using the With statement as follows:

```
With Selection.Interior
    .ColorIndex = 35
    .Pattern = xlSolid
End With
```

Either way, the result is the same. But it helps to be familiar with the With statement, because the macro recorder uses it frequently. In fact, if you look at the original FormatHeader code that the macro recorder generated, you'll find that it uses a With block.

Hunting for Objects

Finding the objects you need isn't trivial, and if you do more than a little bit of macro code editing, you'll want to invest in a dedicated reference book. However, the Visual Basic editor tries to help you out while you're typing code with a feature called IntelliSense. Every time you press the period after typing in a valid object name, a pop-up menu appears with a list of all the properties and methods for that object (see Figure 29-5). If you don't know the exact name of the property or method you want to use, you can scroll through the list, and then select it. Of course, you still need to know *how* to use the property or method, but this feature helps you get started.

Exploring the VBA Language

Now that you've learned a bit about the Visual Basic editor and how Excel uses objects, it's time to jump right in and get some firsthand experience with VBA. First, you'll see how custom-programmed macros can perform some basic editing and formatting. After that, you'll learn how you can use simple macro code to solve problems that you can't tackle with the macro recorder alone. In these sections, you'll see how to manipulate data more flexibly, make decisions, and repeat actions.

Before you get started, here are a few helpful tips:

- For easier macro coding, try to arrange your windows so you can see the Excel window and the Visual Basic editor window at the same time, side-by-side.
- To play a macro, move to the position in the worksheet where the macro should execute, and then switch to the Visual Basic editor. In the Visual Basic editor, move to the appropriate subroutine, and then click the play button.

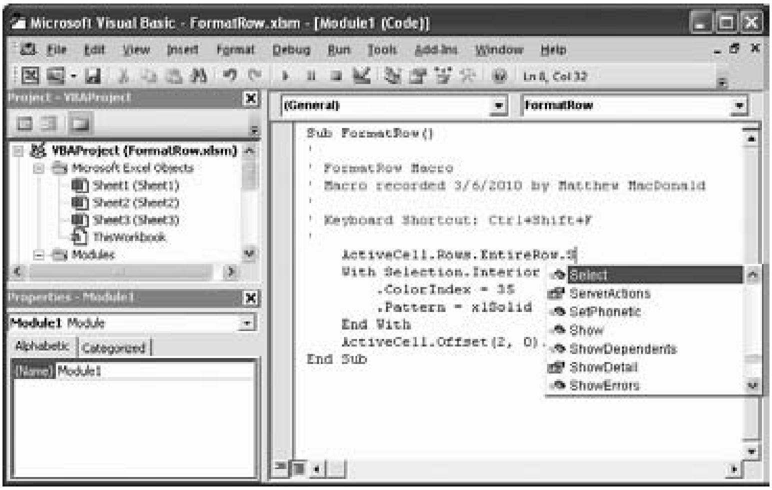

Figure 29-5: *When you press the period key, a Visual Basic feature called IntelliSense shows you the available properties and methods for the current object. Properties are represented with a hand and page icon, while methods are shown with what appears to be a flying green eraser.*

- Make sure you type in every command and object name correctly. Minor differences cause mysterious errors.
- If Excel finds some invalid code when running a macro, it stops executing it and highlights the problem line in yellow. At this point, you can correct the line and press play to keep going, or stop and give up altogether (many programmers find long walks on the beach or guzzling a jug of Mountain Dew a helpful way to clear their heads).

Entering Text in the Current Cell

A few examples can go a long way to showing you how a typical macro works. First of all, check out the following macro subroutine, which represents one of the simplest possible macros you can create. It inserts the text "Hello World!" in the current cell (replacing whatever content may already be there):

```
Sub MyMacro
    ActiveCell.Value = "Hello World!"
End Sub
```

With a little more effort, you can *edit* the current cell value instead of replacing it with new content. Imagine you want to take the current text value, and add the message "INVALID: " before the text. To accomplish this feat, you can use the following macro code:

```
Sub MyMacro
    ActiveCell.Value = "INVALID: " & ActiveCell.Value
End Sub
```

In this macro, Excel begins by joining together two pieces of text: the word "INVALID:" and whatever content's in the current cell. It then takes that combined piece of text and places it back into the cell. In this macro, the ampersand symbol (&) is key. It's a *concatenation operator*, which means it glues together different pieces of text.

You can use a similar approach to adjust a cell that has a number in it. In this case, you can use all the ordinary numeric operators, like +, –, /, *, and ^. Here's an example that multiplies the current cell value by 2, subtracts 1, and then enters the new value:

```
Sub MyMacro
    ActiveCell.Value = (ActiveCell.Value * 2) - 1
End Sub
```

Tip: When using arithmetic operators, make sure the current cell contains a valid number. Otherwise, your code fails with a cryptic "type mismatch" error, which is a reminder that Excel can't perform numeric calculations with text.

Interacting with Other Cells

The ActiveCell object is the starting point for everything you want to do with the current cell. However, it doesn't let you change the content of *other* cells. If you want to do that, you need to access these cells by using the Offset property.

The property looks a little more complicated than other properties because it needs two pieces of information: a row offset and a column offset (in that order). The row offset tells Excel how many rows down you want to move. The column offset tells Excel how many columns to the right you want to move. If you want to move up or left, you need to use a negative number.

The following macro places the phrase *Top cell* in the current cell, and then places the phrase *Bottom cell* in the cell that's immediately underneath it.

```
Sub MyMacro
    ' Change the top cell.
    ActiveCell.Value = "Top cell"

    ' Change the bottom cell.
    ActiveCell.Offset(1,0) = "Bottom cell"
End Sub
```

You need to note one important factor about this code. Although it uses the Offset() property to change the cell that's under the current cell, it doesn't actually *move* to that cell. Instead, when the macro ends, you're still positioned in the top cell.

If you actually want to move to the new cell, you need to use the Activate or Select method, as shown here:

```
Sub MyMacro
    ' Change the top cell.
    ActiveCell.Value = "Top cell"

    ' Move down one cell.
    ActiveCell.Offset(1,0).Select
```

```
    ' Now this changes the bottom cell.
    ActiveCell.Value = "Bottom cell"
End Sub
```

Tip: If you look at the code that Excel creates when you record a macro, you'll see that it often uses the Select method to move around the worksheet. However, when you write your own macros from scratch, it's much simpler–and a far better design–to stay in one place. You can still change cells throughout the worksheet, but you rarely need to actually go there.

Editing Specific Cells

Using Value and Offset, you can romp around your worksheet changing cells as you please. However, you probably remember from the last chapter that you can edit cells in two ways—using relative or absolute references. The example macros you've seen so far use relative references, which means they start working in the current position in the worksheet. However, in some situations you want to move to a specific cell. In order to do this in macro code, use the Range object.

The basic technique is easy. You supply the cell address (like A2) as an argument to the Range object. This points Excel to the cell you want to manipulate. Then, you can do anything you want, including modifying the cell's value, changing its formatting, and so on, all without leaving your current position.

Here's an example that shows this technique:

```
Sub MyMacro
    ' Change cell A1.
    Range("A1").Value = "This is A1"
End Sub
```

Interestingly, you can even modify multiple cells at once using a range reference (like A1:A2). In this case, if you set the value, that value appears in every selected cell.

```
Sub MyMacro
    ' Insert the text "Hello" in ten cells
    Range("A1:A10").Value = "Hello"
End Sub
```

For a little more excitement, take a look at the next macro. It starts by creating a new worksheet for your workbook, and then it fills in several cells in that new worksheet.

```
Sub MyMacro
    ' Create the worksheet using the Add method.
    ActiveWorkbook.Worksheets.Add

    ' Enter several cell values.
    Range("A1").Value = "Company Report"
    Range("A2").Value = "Generated by an Excel macro"

    ' Get the name of the person who owns this copy of Excel
    ' using the UserName property of the Application object.
    Range("A3").Value = "Generated for " & Application.UserName
End Sub
```

Formatting Cells

Conceptually, using macros to format cells is just as easy as using them to edit text. The difference is that you need to think about many more properties, because you can format a cell in dozens of different ways.

Before you can apply any formatting, you need to start off by finding the cells you want to change. If you know the exact reference, you can use the Range object (as shown in the previous section). If you know where the cells are relative to your current cell, you can use the ActiveCell.Offset property (as shown on page 832).

And once you have the object that represents your range of cells, you can use a wide range of properties to format it, including HorizontalAlignment, VerticalAlignment, and MergeCells. You can also use some properties that are themselves objects, and so provide a slew of sub-properties, including Interior (which lets you set fills and patterns) and Font (which lets you configure the typeface and font size).

Here's an example that shows some of the code from the InsertHeader macro:

```
Sub InsertHeader
    With Range("A1:C1")
        ' Note that the alignment properties take special constant values.
        .HorizontalAlignment = xlGeneral
        .VerticalAlignment = xlBottom
        .MergeCells = True
    End With

    ' Change the font of the selected cells.
    With Range("A1:C1").Font
        .Name = "Arial"
        .FontStyle = "Bold"
        .Size = 14
    End With

    ' (Other code omitted.)
End Sub
```

This code grabs a range of three cells (A1 to C1) and changes their alignment and font.

Using Variables

Every programming language includes the concept of *variables*, which are temporary storage containers where you can keep track of important information. In an Excel macro, you can use variables to get around problems that you just can't avoid with the macro recorder.

Imagine you want to swap the content in two cells. On the surface, this operation seems fairly straightforward. All you need to do is copy the text in one cell, place it in the other, and insert the other cell's text in the first cell. Unfortunately, once you paste the new cell content into the second cell, you end up overwriting the content you want to put in the first cell. The easiest way around this problem is to use a variable to keep track of the information you need.

To create a variable in VBA, use the oddly named *Dim* keyword (short for *dimension*, which is programmer jargon for "define a new variable"). After the word Dim, you enter the name of the variable.

Here's how you'd create a variable named CellContent:

```
Dim CellContent
```

Once you've created the variable, you're free to put information in it and take information out. To perform both these operations, use the familiar equal sign, just as you would with properties.

Here's an example that stores some text in a variable:

```
CellContent = "Test text"
```

The following macro puts it all together. It uses a variable to swap the content of two cells.

```
Sub SwapTextWithCellOnRight()
    ' Create the variable you need.
    Dim CellContent

    ' Store the content that's in the current cell.
    CellContent = ActiveCell.Value

    ' Copy the value from the cell on the right
    ' into the current cell.
    ActiveCell.Value = ActiveCell.Offset(0, 1).Value

    ' Copy the value from the variable into the
    ' cell on the right.
    ActiveCell.Offset(0, 1).Value = CellContent
End Sub
```

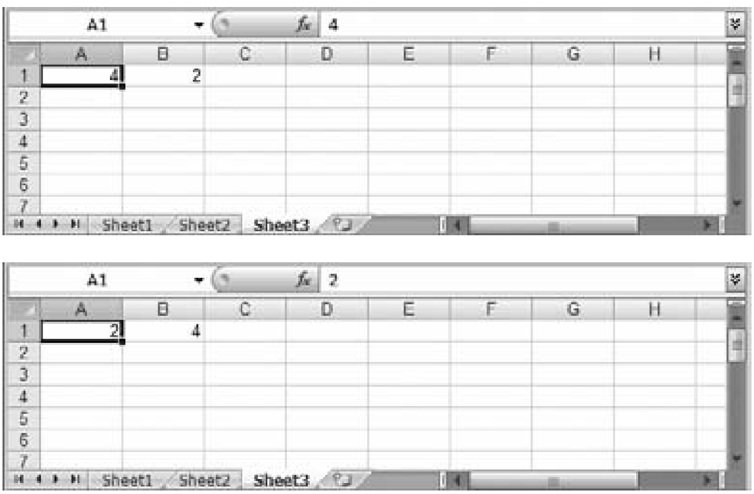

Figure 29-6: *Top: To try out the SwapTextWithCellOnRight macro (included with the downloads for this book), type number 4 in cell A1 and the number 2 in cell A2. Now move to cell A1, and then run the macro.*

Bottom: The macro moves the number 2 into cell A1 and the number 4 into cell A2. It isn't rocket science–but it's almost impossible to do without using variables.

POWER USER'S CLINIC

Variable Data Types

The variable examples in this chapter are a bit lazy. They don't specifically tell Excel what type of information the variable will hold. This approach works perfectly well, but real programmers frown on it, because it lets certain types of errors slip into your code undetected. For example, you might inadvertently store the wrong type of data in a variable.

A better approach is to add the As keyword after your variable name, followed by the variable's *data type*. The only problem is that in order to do this, you need to know what data types VBA supports. Some common data types include String (for storing text information, which you have to enclose in quotation marks), Integer (for storing positive and negative whole numbers), Single (for storing fractional numbers) and Date (for storing dates). Here's an example:

```
Dim MyText As String
MyText = "This is string data."

Dim MyNumber As Integer
MyNumber = 42
```

Simple data types are easy enough, but the Excel objects also have their own data types too. For example, to define the CellContent variable in the previous example, which represents the content in a cell, you need to use the Range data type:

```
Dim CellContent As Range
```

As you can imagine, you need a fair bit of experience before you know when to use the right data type. To learn more, get yourself a good VBA reference, as described on page 842.

Making Decisions

Conditional logic is another programming staple, and it's code that runs only if a certain condition is true. There's no limit to the number of ways you can use conditional logic. You could want to perform a different calculation based on the value of a cell, apply different formatting based on the number of cells, or create a different printout depending on the date. All these operations and many more are possible when using conditional logic.

All conditional logic starts with a *condition*, a simple expression that can turn out to be true or false (programmers call this process *evaluating* to true or false). Your code can then make a decision to execute different logic depending on the outcome of the condition. To build a condition, you need to compare a variable or property using a logical operator like = (equal to), < (less than), > (greater than), and <> (not equal to). ActiveCell.Value = 10 is a condition. It can be true (if the current cell contains the number 10), or false (if the current cell contains something else).

On its own, a condition can't do anything. However, when used in conjunction with other code, it can become tremendously powerful. Once you've created a suitable condition, you can put it inside a special structure called the *If block*. The If block evaluates a condition, and then runs a section of code if the condition is true. If the condition isn't true, Excel completely ignores the code.

Here's a macro that looks at the current cell value. If that value exceeds 100, then Excel changes it to 100. If the cell value is less than 100, nothing happens, and the current value remains:

```
Sub MyMacro
    If ActiveCell.Value > 100 Then
        ' This value is too big. Change it to the maximum of 100.
        ActiveCell.Value = 100
    End If
End Sub
```

Note that the If block always starts with If and ends with End If. Everything else is conditional and runs only if the condition is true.

An If block can also evaluate several different conditions. Here's an example that considers the current value of a cell. Depending on the cell's value, the If block uses a different calculation to arrive at the sales commission, which it places in another cell:

```
Sub MyMacro
    If ActiveCell.Value > 1000 Then
        ' Use the 5% commission rate.
        ActiveCell.Offset(0,1).Value = ActiveCell.Value * 0.05
    ElseIf ActiveCell.Value > 500
        ' Use the 2.5% commission rate.
        ActiveCell.Offset(0,1).Value = ActiveCell.Value * 0.025
    Else
        ' Give a basic $5 comission.
        ActiveCell.Offset(0,1).Value = 5
    End If
End Sub
```

Here, only one segment of code runs. Excel works its way through the If block, testing each condition until one matches. If the cell value is greater than 1,000, it runs the first conditional block of code, and then jumps down to the closing End If statement. It then continues with any other code that's in the macro. If the cell value is less than (or equal to) 1,000 but greater than 500, the first condition is false, and Excel tries the second one, which is true. If no condition matches, Excel runs the code in the final Else clause.

These examples scratch only the surface of what careful conditional logic can do. You can use *And* and *Or* keywords to combine conditions, put one conditional block inside another, and much more. To learn more about these approaches, you may want to consult the Web or a dedicated book about VBA programming. Those ready to dive right in to the nitty-gritty, can try *Excel 2010 Power Programming with VBA*, by John Walkenbach (Wiley).

Repeating Actions with a Loop

Computers work particularly well when you need to automate a tedious task. While you may tire out after typing in your 100th cell value, an Excel macro has no such weakness, and can perform thousands of operations without pausing.

The *loop* is one of the best tools for repeating operations. A loop is another type of block, one that repeats itself over and over again. Here's an example:

```
Do
    ActiveCell.Select
    Selection.Interior.ColorIndex = 35
    Selection.Interior.Pattern = xlSolid
    ActiveCell.Offset(1,0).Select
Loop
```

When Excel reaches the final Loop statement at the bottom of this loop, it automatically jumps back to the beginning and repeats your code. However, there's one problem—this process continues infinitely! That means if you make the mistake of running this macro, your worksheet is locked up indefinitely (until you press the emergency-stop key combination, Ctrl+Break).

To avoid this situation, you should build all loops with an *exit condition*. This condition signals when the loop should end. Here's a rewritten version of the same loop that stops as soon as it finds an empty cell:

```
Do Until ActiveCell.Value = ""
    ActiveCell.Select
    Selection.Interior.ColorIndex = 35
    Selection.Interior.Pattern = xlSolid
    ActiveCell.Offset(1,0).Select
Loop
```

This technique is quite powerful. Consider the following macro, which uses a loop to format all the rows in a table. It gives each row an alternating color and stops when there are no values left:

```
Sub FormatAllCellsInColumn
    Do Until ActiveCell.Value = ""
        ' Format the first row.
        ActiveCell.Rows.EntireRow.Select
        Selection.Interior.ColorIndex = 35
        Selection.Interior.Pattern = xlSolid

        ' Move down two rows.
        ActiveCell.Offset(2,0).Select
    Loop
End Sub
```

This macro is really an enhanced version of the FormatRow macro shown in the previous chapter. Unlike FormatRow, you need to run this macro only once, and it takes care of all the rows in your worksheet that contain data.

Excel actually has different types of loops. The For Each loop is another useful loop, which repeats itself once for every item in a collection of objects. For Each loops come in handy if you need to process all the cells in the current selection.

Imagine you want to fix up many cells that have a jumble of upper-and lowercase letters. As you learned on page 319, the Excel PROPER() function can do the trick and convert a string like "hElLo THERE" to a respectable "Hello There." The downside is that you need to write a separate formula for each cell you want to change. A much better solution is to use the PROPER() function from inside a macro to perform the change automatically.

The following macro does the trick. It accesses the PROPER() function through the Application.WorksheetFunction object:

```
Sub FixText()
    ActiveCell.Value = Application.WorksheetFunction.Proper(ActiveSheet.Value)
End Sub
```

This useful macro quickly cleans up the current cell. However, if you select *multiple* cells and run the macro again, you're likely to be disappointed. The FixText macro changes only the current cell. It ignores all the other selected cells.

If you want to take these other cells into account, you need to create a loop using For Each. That's because the For Each block lets you scan through all the selected cells, and then run a series of code statements once for each cell.

Here's the revised macro, which cleans up every selected cell:

```
Sub FixTextInAllCells()
    ' This variable represents the cell you want to change.
    Dim Cell

    ' Change all the cells in the current selection.
    For Each Cell In Selection
        ' This code repeats once for each cell in the selection.
        Cell.Value = Application.WorksheetFunction.Proper(Cell.Value)
    Next
End Sub
```

This code works even if you don't explicitly select any cells. That's because Excel treats the current cell as though it's selected. In this situation, the For Each loop performs just one pass.

Figure 29-7 shows the result.

Creating Custom Functions

So far, you've seen how you can use code to create powerful macros that take control of Excel. But you have another option for plugging your logic into Excel: You can create custom functions (known in programmer-ese as *user-defined functions*). Essentially, a custom function accepts some information (through arguments), performs a calculation, and then provides a result. Once you've created your custom function, you can use it in a formula in a cell, in exactly the same way that you use Excel's built-in functions.

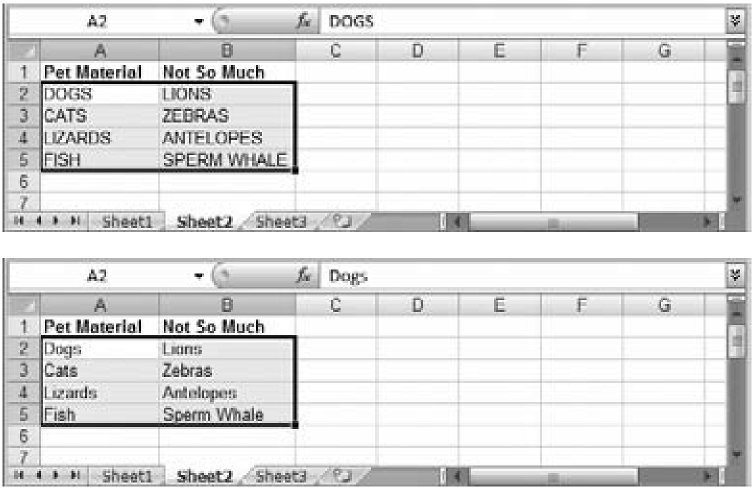

Figure 29-7: *Loops are one of the handiest tools for fixing up large tables in a hurry. Here, a table crammed with text that's in all capitals (top) is transformed into something much more palatable (bottom). The best part? You need to select the cells and run the FixTextInAllCells macro only once.*

You create custom functions in the same place that you create macros—in modules. In fact, any number of macros and functions can exist side-by-side in a module. The difference is that macros start with the word *Sub*. Custom functions start with the word *Function*.

Here's an example of one of simplest possible custom functions:

```
Function GetMyFavoriteColor()
    GetMyFavoriteColor = "Magenta"
End Function
```

Every function needs to provide a result, which is what appears in the cell when you use the function. To set the result, you use the name of the function, followed by an equal sign and the value, as shown here:

```
GetMyFavoriteColor = "Magenta"
```

The function provides the text "Magenta" as its result. If the function provided a number instead of a piece of text, you wouldn't use the quotation marks.

To use the function in your worksheet, just create a formula that uses the function. Figure 29-8 shows you how.

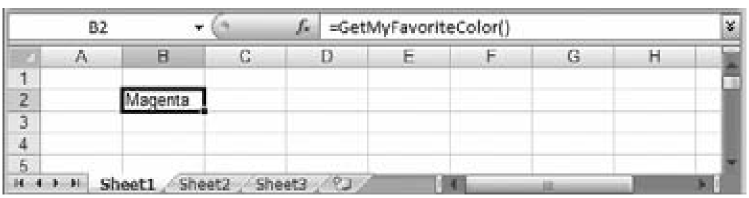

Figure 29-8: *In this example, the custom function is named GetMyFavoriteColor(). To use this homemade function in your worksheet, just create a formula that uses the function, like =GetMyFavoriteColor().*

The GetMyFavoriteColor() function is particularly simple because it doesn't use any arguments. But there's no reason you can't get a little fancier. Consider the following custom function, which takes two arguments—length and width—and calculates the total area by multiplying them together:

```
Function Area(Length, Width)
    Area = Length * Width
End Function
```

The two arguments are defined in the parentheses after the function name. You can add as many as you want, as long as you separate each one with a comma.

Now, to use this function in a cell in your worksheet, you can use this formula:

```
=Area(100,50)
```

This formula uses literal values, but you can easily go one step further and use cell references instead.

Neither the GetMyFavoriteColor() nor Area() function shows you anything you couldn't already achieve on your own. The following example shows a much more practical function that does something you *can't* otherwise achieve in Excel—*banker's rounding*. As discussed on page 267, Excel's rounding can contribute to biases when adding numbers. To avoid this problem, accountants sometimes use other types of rounding. Banker's rounding is one example—it rounds 0.5 up sometimes and down other times, depending on whether the number before it is even or odd. 1.5 always rounds up to 2, while 2.5 always rounds down to 2. With ordinary rounding, 0.5 is always rounded up, and that can slightly inflate long columns of numbers.

Implementing banker's rounding by hand is tricky, and it requires some conditional logic. The VBA language, however, gives you a powerful shortcut. It provides a built-in Round() function that always uses banker's rounding, unlike Excel's ROUND() function.

Here's the complete function you need:

```
Function BankerRound(NumberToRound)
    BankerRound = Round(NumberToRound)
End Function
```

To test it out, try the following two formulas. The first one produces a result of 0, while the second one has a result of 2:

```
=BankerRound(0.5)
=BankerRound(1.5)
```

POWER USERS' CLINIC

Power Programming

If this taste of Excel programming has whet your appetite, there's no need to stop now. The next logical step is to read a book that's all about Excel programming. Two good choices are *Excel 2010 Power Programming with VBA* (Wiley) and *Programming Excel with VBA and .NET* (O'Reilly). If possible, look for an edition that's updated for Excel 2010, but don't worry if you're stuck with an older copy. Other than a few additions to cover new Excel features, the VBA language really hasn't changed over the years.

If you already have solid programming skills, you might be interested to know that the VBA macro language isn't the only game in town. For example, just about any Windows programming language can communicate with Excel—and take control—using COM. (COM, or Component Object Model, is a Microsoft technology for sharing useful objects between different programs.) For example, a Visual Basic developer can use this technique to build a standalone program that "drives" Excel—creating worksheets, inserting data, and performing calculations. The best part is that you still use the same set of Excel objects, methods, and properties that you do in a macro, but you place your code in a standalone program that has much greater power. This approach gives you the ability to create large-scale solutions, like programs that generate Excel documents or websites that read Excel data.

COM has "oldie but goodie" status. It's been around for much more than a decade, and while it still works, cutting-edge developers have newer and slicker tools. One of these is Visual Studio Tools for Office (known to programming geeks as VSTO), a tool that's included with many versions of Microsoft's Visual Studio development tool. VSTO is the most powerful way to program Excel, but it's also the most complicated. It makes sense if you're already skilled with a .NET programming language like C# of VB .NET. To learn more, you can visit Microsoft's Office Developer Center at *http://msdn.microsoft.com/en-us/office/aa905411.aspx.*

APPENDIX

A

Customizing the Ribbon

When Microsoft introduced the ribbon in Office 2007, they clamped down on customization. Quite simply, the designers of Office were concerned that overly creative Office fans would replace the standard arrangement of buttons with a jumble of personal favorites. Their worst fear was that Excel customizers would transform Excel so completely that no one else would be able to use Excel on their computers, and the instructions in books like this one would be useless. To prevent this crisis, Microsoft made it extremely difficult to customize the ribbon. In fact, ribbon customization was only available to programming gearheads who were willing to work with the intimidating RibbonX standard.

Excel 2010 isn't nearly as paranoid. It lets anyone rename tabs, hide them, add groups, and remove them. It even lets you create an entirely new tab with your own button selections. When used carefully, this feature gives you a great way to speed up your work and put your favorite commands in a central spot. In more reckless hands, it's a great way to confuse friends, family, coworkers, and even yourself.

Along with its surprisingly powerful ribbon customization ability, Excel 2010 also lets users customize the Quick Access toolbar—the sequence of tiny buttons that sits just about the ribbon and its tabs. If you're not so ambitious that you want to create your own custom tab, the Quick Access toolbar is a convenient place to stick your absolute favorite buttons. This is the task you'll consider first.

Adding Your Favorites to the QAT

You've already seen the Quick Access toolbar (known to Excel nerds as the QAT). It's the micro-size toolbar that sits above the ribbon. The Quick Access toolbar has only icons, but you can hover over a button to get the full command text.

When you first start out with Excel, the Quick Access toolbar is a lonely place, with buttons for quickly saving your workbook and undoing or redoing the last action. However, Excel gives you complete control over this space, including the ability to add new buttons. You can most quickly add stuff by clicking the down-pointing arrow at the far right side. Figure A-1 shows how it works.

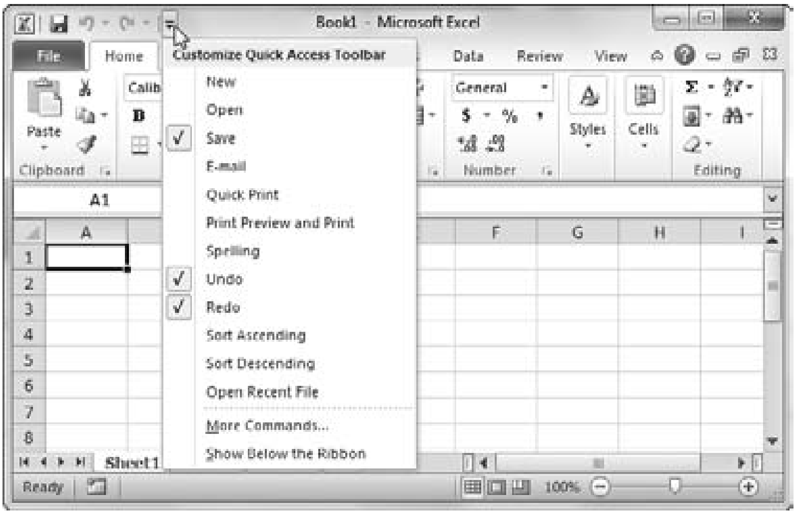

***Figure A-1:** When you click the drop-down arrow on the Quick Access toolbar, Excel shows a list of often-used commands that you can add just by clicking them. These commands include some for creating a new workbook, opening an existing workbook, sending it to the printer with no questions asked, and firing up the spell checker. But to see all your possibilities (or to rearrange the commands you've added), you need to choose More Commands.*

Tip: If you don't like the Quick Access toolbar's placement, Excel gives you one other option. Click the drop-down arrow, and then choose "Show Below the Ribbon" to move your toolbar under the ribbon so your mouse has less distance to travel.

You might add buttons to the Quick Access toolbar for two reasons:

- **To make it easier to get to a command you use frequently**. If it's in the Quick Access toolbar, you don't need to memorize a keyboard shortcut or change the current ribbon tab.
- **To get to a command that the ribbon doesn't provide**. Excel has a small set of less popular commands that it lets you use, but that it doesn't keep in the ribbon. Many of these commands are holdovers from previous versions of Excel. If you have a long-lost favorite Excel feature that's missing, it just might be available via the Quick Access toolbar's extra buttons.

Keyboard lovers can also trigger the commands in the Quick Access toolbar with lightning speed thanks to Excel's KeyTips feature (page 5). When you press the Alt key, Excel displays a number superimposed over every command in the Quick Access toolbar (starting at 1 and going up from there). You can then press the number to trigger the command. So in the Quick Access toolbar shown in Figure A-1, Alt+1 saves the workbook, Alt+2 opens the Undo list, and so on.

You've already seen how to customize the Quick Access toolbar in a few places in this book. You used it to get to the Compare and Merge Workbooks feature in Chapter 25 and to get quick access to macros in Chapter 28.

Tip: If you want to add a command that duplicates something that's already in the ribbon, here's a shortcut. Find the command in the ribbon, right-click it, and then choose Add to Quick Access Toolbar.

Adding Buttons

To add a button to the Quick Access toolbar, you follow these steps:

1. **Click the drop-down arrow on the Quick Access toolbar, and then choose More Commands.**

   The Excel Options window opens and positions you at the Quick Access Toolbar section where you need to be (Figure A-2).

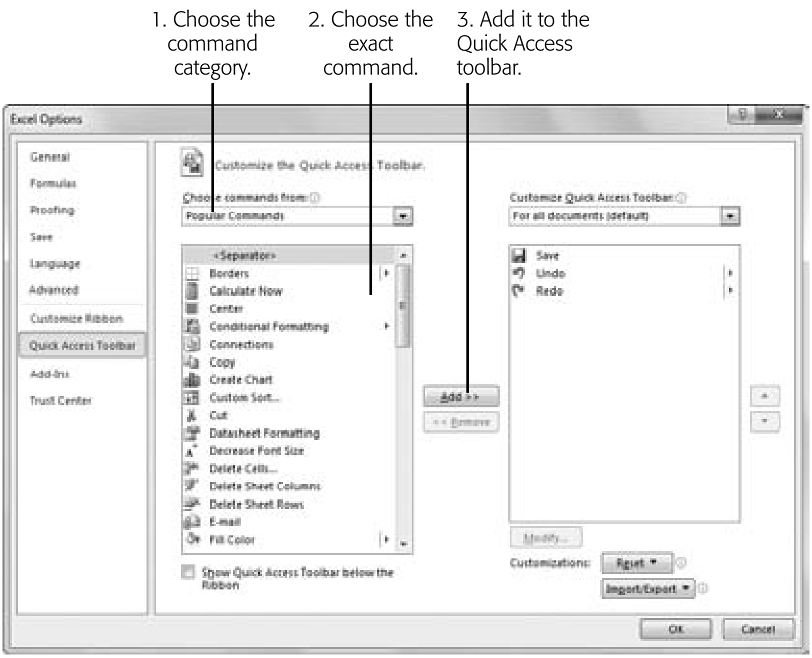

Figure A-2: *The Quick Access Toolbar section of the Excel Options window has two areas. The list on the left lets you choose the command you want to add. The list on the right shows the commands that currently appear in the Quick Access toolbar.*

2. **Choose a category from the "Choose commands from" list.**

   The library of commands that you can add to the Quick Access toolbar is enormous. To make it easier to find what you want, it makes sense to choose a specific category. Many of the categories overlap—Excel simply provides them to make finding what you want easier. Here are the top choices:

   - **Popular Commands** gives you a short list of commands that Excel jockeys love. If you're trying to get quick access to a commonly used feature, you just might find it here.
   - **Commands Not in the Ribbon** provides all the leftovers—commands that Microsoft didn't consider useful enough to include in the ribbon. This list holds some commands that are superseded or partially duplicated by other commands, commands that are included in other dialog boxes, and commands that were used in previous versions of Excel and put out to pasture in this release.
   - **All Commands** includes the full list of choices. As with the other categories, it's ordered alphabetically.
   - **Macros** shows all the macros in the currently open workbooks, including the personal macro workbook. Be careful about adding a macro that's not from the personal macro workbook—if you move the file, Excel can't run the macro when you click the button.

   Under these categories are several additional categories that correspond to the File menu and various tabs in the ribbon. For example, you can choose the Insert tab to see all the commands that appear in the ribbon's Insert tab. At the top of the tab list is the File tab, which has the commands from the File menu and backstage view. At the bottom of the list are the "Tools" tabs—tabs that appear only when certain objects are selected, like charts, pictures, or images.

3. **Once you've chosen the category you want, pick the command from the list below, and then click Add.**

   The command moves from the list on the left to the list on the right, placing it on the Quick Access toolbar (Figure A-3).

4. **You can repeat this process (starting at step 2) to add more commands.**

   Optionally, you can rearrange the order of items in the Quick Access toolbar. Just pick a command, and then use the up and down arrow buttons to move it. The topmost commands in the list are displayed to the left on the Quick Access toolbar.

Tip: If you've customized the heck out of your Quick Access toolbar and want to go back to a simpler way of life, then click the Reset button.

Figure A-3: *In this example, the Calculator command is being added to the Quick Access toolbar, so you can pop open the handy Windows calculator whenever you need it, without leaving Excel.*

5. **When you're finished, click OK to return to Excel with the revamped Quick Access toolbar.**

   Adding to the Quick Access toolbar isn't a lifetime commitment. To get rid of a command you don't want anymore, right-click it, and then choose Remove from Quick Access Toolbar.

Note: You might notice the tempting Modify button, which lets you change a command's name and picture. Unfortunately, it works only for macro commands.

Customizing Specific Workbooks

Do you have a button or two that you're using incessantly, but just for a specific workbook? In this situation, it may not make sense to customize the Quick Access toolbar in the normal way. If you do, you'll get your extra buttons in *every* workbook you use, including those in which the extras aren't useful.

Excel has a great feature to help you out in this situation. You can customize the Quick Access toolbar for an individual workbook. That way, whenever you open that workbook, the buttons you need appear in the Quick Access toolbar. When you close it (or open another workbook in a separate window), the buttons disappear.

Note: Customizing individual workbooks has advantages and disadvantages. The disadvantage is that you need to perform this task separately for every workbook, which can take a lot of time. The advantage is that your customizations are recorded right in your workbook file. As a result, they stick around if you open the workbook on someone else's computer (or if you log onto your computer as a different user).

To customize the toolbar for a single workbook, you follow the same steps that you used in the previous section. Start by clicking the Quick Access toolbar's drop-down arrow, and then choosing More Commands. However, before you add any commands, change the selection in the Customize Quick Access Toolbar list, which appears just above the list of commands in the Quick Access toolbar. Instead of using "For all documents (default)," choose your workbook's name (as in "For SecretSanta.xlsx"). This list starts off empty. Then, follow the normal steps to add buttons.

When Excel displays the Quick Access toolbar, it combines the standard buttons (as configured in the previous section) with any buttons you've defined for the current workbook (Figure A-4).

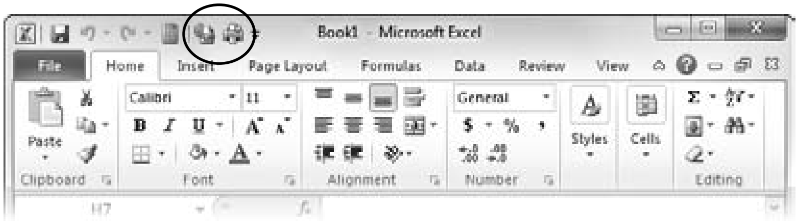

***Figure A-4:** The workbook-specific buttons (circled) always appear after standard buttons. You can define the same button as a standard button and a workbook-specific button, in which case it appears twice.*

Customizing the Quick Access toolbar for a specific workbook is a handy trick. It makes great sense with macros (Chapter 28) because it lets you create a single workbook that has a useful set of macros and handy buttons for running them. It's also great with templates (Chapter 16). If you customize the Quick Access toolbar for a specific template, every workbook that you create using that template gets the buttons you've added.

Personalizing the Ribbon

If you don't want the complexity that's involved in rethinking the ribbon, customizing the Quick Access toolbar makes perfect sense. But if you *are* hungering to make some more radical changes, and you think you could become more productive with a ribbon that's tailored to the tasks you perform most often, then ribbon customization is the best choice.

Ribbon customization takes place in the same Excel Options window you used to change the Quick Access toolbar. The easiest way to get to there is to right-click the ribbon, and then choose Customize Ribbon. Excel opens the Excel Options window and sends you straight to the Customize Ribbon section (Figure A-5).

NOSTALGIA CORNER

What Happens to Excel 2003 Toolbars?

Using the customization features in Excel 2003 (or before), you might have designed your own specialized toolbars (which you could then attach to a workbook file). The bad news is that Excel no longer provides this feature, so you can't enhance or modify your existing toolbars. The good news is that you can still access your custom toolbar, just in a slightly different way.

All the buttons from custom toolbars and custom menus end up on a separate tab in the Excel ribbon, named Add-Ins. This tab appears only if there's a custom toolbar or custom menu in the workbook. When it appears, it ends up at the end of the series of standard ribbon tabs. This arrangement may not be the prettiest, but it makes sure that old solutions keep working in Excel 2010.

If you want to keep your custom commands closer at hand, you can add them to the Quick Access toolbar. You need to look for three items (in the All Commands list): Custom Toolbars, Menu Commands, and Toolbar Commands. When you add these to the Quick Access toolbar, they appear as drop-down menus that, when clicked, show all the custom commands in your current workbook.

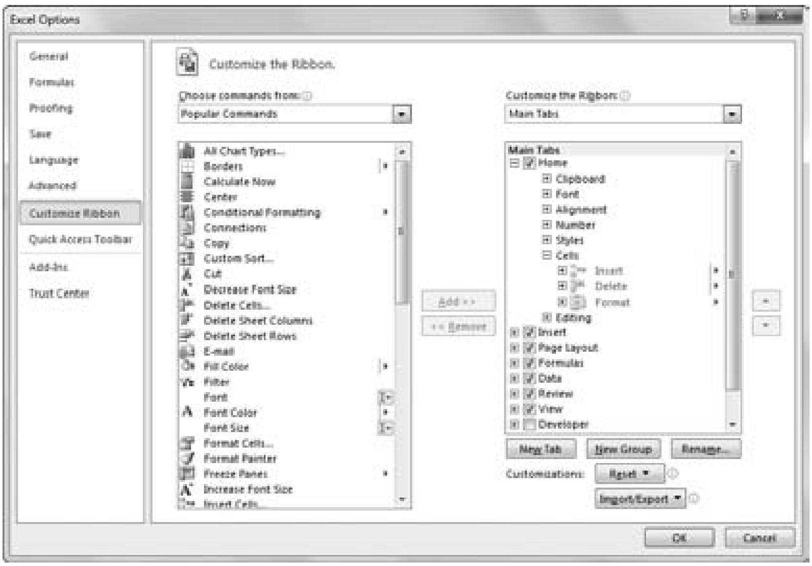

Figure A-5: *Customizing the ribbon is a lot like customizing the Quick Access toolbar. The most obvious difference is that there's a lot more information on right side of the window, because this list includes all the ribbon tabs. To see the groups in each tab and the commands in each group, click the tiny plus (+) icon next to the appropriate item. For example, here the Home tab is expanded to show all its groups, and the Cells group is expanded to show the buttons it contains.*

There's a lot you can do to fine-tune your ribbon. The next sections show you how.

Changing Existing Tabs

Tired of Excel's standard tabs? Surprisingly, you can do quite a bit to change them. Of course, *can* doesn't mean *should*, and if you're overly ambitious, you may just end up hiding something you need and making Excel much more difficult to use.

Here are the changes you can make to an existing tab:

- **Hide a tab.** If you're absolutely positive you don't need any of the commands in a tab, or if you're a system administrator trying to prevent your users from messing up spreadsheets with the wrong features, you can remove the tab from the ribbon. To do so, clear the checkbox next to the tab name. To restore it at a later time, add the checkmark back.
- **Remove a group.** A similar, but more fine-grained strategy is to remove a section of a tab. This leaves more room for the rest of the buttons (which, presumably, are ones that are more useful to you). To remove a group, click the plus (+) icon to expand the tab that has it, select the group, and then click the Remove button (which appears between the two lists).
- **Change the order of tabs.** If you use one tab far more than another, you might decide to change the tab order. Click to select the tab you want to move, and then click the up arrow button (to move the tab toward the left side, or start of the ribbon) or the down arrow (to move the tab toward the right side, or end of the ribbon). You can find both arrow buttons just to the right of the list.
- **Change the order of groups in a tab.** This works in the same way as changing the tab order, but it lets you rearrange the sections inside a tab. Click the plus (+) icon to expand the appropriate tab, select the group you want to move, and then use the up and down arrows to shuffle its position. But be warned: though this seems like a small change, it can seriously throw off other Excel fans, because people tend to remember the general position of the commands they use.

Note: If you keep pressing the up arrow when you reach the top of the tab's group list (or the down arrow when you reach the bottom), Excel moves the group to the next adjacent tab.

- **Rename a tab or group.** Select either a tab or a single group. (You can't rename individual buttons.) Then click the Rename button under the list. A small window pops up where you can type in the new name, and then click OK.
- **Add a new group to an existing tab.** First, expand the tab where you want to place your group. Select an existing group, and then click the New Group button. Your new group is placed immediately after the group you selected, but you can use the arrow buttons to move it. Now give your group a good name (click Rename), and then fill it up. To do that, select the command you want in the list on the left, and then click Add, just as you did when adding commands to the Quick Access toolbar.

Note: The ribbon changes you make are linked to your Windows user account. If someone else logs on to your computer with a different user name and runs Excel, that person gets the standard ribbon.

And here's the much smaller list of changes Excel doesn't let you make:

- **Delete or rearrange commands in a standard group.** You can take a standard group or leave it, but you can't change what's inside. Of course, nothing stops you from removing a standard group and creating a custom group that has just the commands you want.
- **Add custom commands to a standard group.** If you want to add new commands, you need to place them in a custom group in a standard tab or, even better, a custom group in a custom tab.

If you find you've gone too far, Excel lets you return the ribbon to its original state. Just click the Reset button, and then choose "Reset all customizations." Alternatively, you can reset a single tab: just select it in the list, click the Reset button, and then choose "Reset only selected Ribbon tab."

UP TO SPEED

Ribbon Tweaking: Too Much of a Good Thing?

Before you begin a wild bout of customization, it's worth asking which customizations are really worth the effort. Here's the rundown:

- Hiding tabs and removing groups makes sense if you're trying to simplify Excel. However, it risks losing features you'll need later, which is usually far more inconvenient.
- Changing tab and group names doesn't make you more efficient, but it could confuse other people, so it's usually a bad idea.
- Rearranging groups is a worthwhile strategy if you want to move the features you never use to the far right end of the tab, so they won't distract you.
- Adding new groups is a great way to get important features at your fingertips, but it probably makes even more sense to put your custom-picked buttons on a brand new tab you've added from scratch, as you'll do in the next section.

Any ribbon customization is a tradeoff between personalization and consistency. You already know that a revamped ribbon can confuse other Excel fans when they use your computer. But you might not have realized the more insidious reverse effect–namely, you'll be embarrassingly slow when you switch over to a normal Excel installation on a colleague's computer or the local copy shop. For that reason, it's best to practice a bit of restraint and follow these ribbon-customization guidelines:

- Customize only when you're absolutely sure it will make your life easier and more convenient.
- If you want to add more than two or three new commands to the ribbon, consider putting them all into a new tab, so they're clearly separate from the standard buttons.
- If you want to customize something just so you can express your own personal design aesthetic, stick to your computer's background wallpaper and desktop icons.

Creating Your Own Tab

The safest way to customize the ribbon is to put your custom buttons on a completely separate tab. This lets you keep the rest of the Excel user interface in its normal state (and ensures that you'll be able to use the instructions in this book). It's also oddly satisfying to have a ribbon tab all to yourself, to fill with your favorite shortcuts, as shown in Figure A-6.

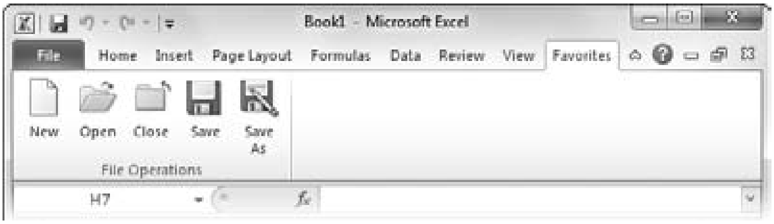

Figure A-6: *You can fit quite a bit more into a custom ribbon tab than you can fit in the Quick Access toolbar. Here, the custom tab (named Favorites) has a few items from the File menu, like Open, Save, and Close.*

You probably already know how to create and fill a custom tab. But if you're unsure, just follow this sequence of steps:

1. **In the Excel Options window, select the last tab in the list, and then click the New Tab button.**

   Excel adds a new tab named "New Tab" at the end of the ribbon. (You can move it by selecting and using the arrow buttons.) Inside the new tab, Excel adds a single new group named "New Group."

2. **Change the tab and group names.**

   First, select the new tab, click Rename, and then fill in a new name. Then, select the new group, click Rename, and give it a new name. For example, in Figure A-6 the tab is named Favorites and the first group is named File.

3. **Now add your favorite commands.**

   The Customize Ribbon section of the Excel Options dialog box works the same as the Quick Access Toolbar section. You pick commands in the list on the left, and then click Add to transfer them over to the currently selected custom group.

   For example, to add the Open command shown in Figure A-5, start by picking "File Tab" from the "Choose commands from" list. Select Open, and then click the Add button.

Note: The command list includes many variations of each command. For example, along with Save, you'll find Save As, Save As Other Format, Save as Another File Type, Publish as PDF or XPS, and so on. Although you can do just about everything with Save As, these variations are slightly more convenient for different saving tasks. You may want to explore a little and try the different variations before you settle on the exact commands you want.

4. **Optionally, add more groups.**

   To make a complete tab, you'll probably want several groups of commands. To create each new group, click Add Group. You can use the arrow buttons to rearrange them appropriately.

5. **Once you're finished, consider saving your ribbon for posterity.**

   The following section shows how to do this.

Saving and Reusing Your Custom Ribbon

Dedicated Excel customizers can spend a lot of effort getting exactly the right arrangement of tabs and buttons. After doing all that work, it's natural to wonder if you can reuse your custom ribbon on another computer. Fortunately, Excel has an export feature that preserves all your hard ribbon-customization work. It copies your ribbon settings to a special file (with the clunky file extension *.exportedUI*), which you can then apply to a copy of Excel that's installed somewhere else.

Here's how to use Excel's ribbon export feature:

1. **Start by perfecting your ribbon in the Excel Options window.**

2. **When you're finished, click the Import/Export button, and then choose "Export all customizations."**

   A File Save dialog box appears.

3. **Browse to the folder where you want to place your file, and then give it a suitable name, like DavesCustomRibbon. Click Save.**

   Excel adds the file extension .exportedUI, so if you type in DavesCustomRibbon, Excel actually saves a file named DavesCustomRibbon.exportedUI.

4. **Now go to the other computer where you want to replicate your ribbon, and take your file with you. In the Excel Options window, click the Import/Export button, and then choose "Import customization file."**

   A File Open dialog box appears.

5. **Browse to your file, select it, and then click Open.**

   Before it imports your ribbon, Excel gives you a last-minute confirmation warning, because the settings that are stored in your file will overwrite any custom ribbon settings that are already applied.

6. **Click Yes to seal the deal and import your custom ribbon.**

   Presto—Excel replaces the standard-issue ribbon with your personalized version.

FREQUENTLY ASKED QUESTIONS

Troublesome Tabs

Help! There's a tab in my Excel ribbon that I can't remove.

Sometimes, a tab may appear in the Excel ribbon but not in the Excel Options window. In this situation, the culprit is almost always an add-in–a program that plugs into Excel, extending it with new features and (sometimes) a ribbon tab that just won't go. The only way to remove this sort of tab is to turn off the add-in.

To switch off an add-in, you must go to the Add-Ins section of the Excel Options window, choose the appropriate add-in type from the Manage list (if you're not sure what it is, try both Excel Add-ins and COM Add-ins), and click the Go button. You'll then see a list of add-ins that you can disable or remove.

Index

Symbols

3-D charts, formatting, 528–531
3-D column charts, 483
: (colon) as range operator, 237
, (comma) as union operator, 237
& (concatenation operator), 832
= (equal sign) in formulas, 228
 # (number sign), column sizing and, 105

A

ABS() function, 281
absolute cell references, 254–255
absolute reference mode (macros), 796–797
accelerated depreciation functions, 310–311
Access databases, connecting to, 659–661
Accessibility Checker, 718
Accounting format, 144–145
Add-Ins tab (Ribbon), 849
ADDRESS() function, 351–353
aligning shapes, 563
alignment/orientation of cells, 159–162
alternating row formatting (macro example), 812–813
anchoring charts, 468
AND() function, 364
apostrophe in text identification, 67–68
area charts, 486
arguments, defined, 235
arithmetic operators, 229–230
arithmetic rounding, 267
array formulas, 277, 253
arrow notations, 13
arrows in worksheets, 394
aspect ratios (images), 541
auditing tools for formula errors, 387–388
Authenticode signing (macros), 811
AutoComplete feature, 72–73
AutoCorrect feature, 73–74
AutoFill feature, 75–77
AutoFit feature, 78–79
automatic filtering (tables), 412
automatic formatting of cells, 233
automatic subtotaling, 441–444
Auto Outline feature, 441
AutoRecover feature, 53–55, 82–85
AutoSelect feature, 99, 102
Auto-Sum button, 269
AVERAGE() function, 274–275
AVERAGEIF() function, 368–370
average/median values, finding, 274–275

B

background colors (charts), 509–510
backgrounds, adding to worksheets, 172
backing up files with AutoRecover, 53–55
backstage view
- downloading templates from, 448–451
- overview, 8–9, 16–17
- switching to, 36–37

banding (table layouts), 632
banker's rounding, 267, 841
bar charts, 484
bin constraints, 608
body fonts, 181
borders
- and lines (charts), 514–516
- and fills, formatting, 167–171
- for pictures, 550–551

brightness, adjusting (pictures), 543
browsers, uploading files to SkyDrive with, 753–755
bubble charts, 492–493

buttons
adding to Quick Access toolbar, 845–847
adding to worksheets, 817
button controls, 817
buttons with menus (Ribbon), 14

C

calculations
adding calculated fields to pivot tables, 640–642
fine-tuning in pivot tables, 639–643
table. *See* table calculations
tool (status bar), 98
capitalization, changing, 319–320
category axis (charts)
defined, 474, 498
scaling, 476–478
CEILING() function, 265–266
cell ranges
protecting with passwords, 709–711
protecting with user account permissions, 711–714
using with functions, 237–238
cell references
absolute, 254–255
basics of, 231–232
data validation with, 700–702
defined, 229–230
descriptive names for, 371–372
formatting cells containing, 232–233
partially fixed, 256
relative, 252
shifting, 349–350
stored as text, 351–353
using with functions, 236
cells
accessing with Offset property (macro), 832–833
adding only those specified, 367–370
alignment of, 159–162
automatic formatting of, 233
cell addressing systems, 27
CELL() function, 358–359
cell offsets, 255
changing data validation in multiple, 696
changing location of, 262
checking values inside, 356–358
copying/cutting parts of, 113
counting only those specified, 365–367
dependents of, 387
editing specific (macro), 833
entering text in (macro), 831–832
finding formatted, 133–134
formatting. *See* formatting cells
inserting copied or cut, 114–115
locked and hidden, 703–704, 707
merging, 162
moving around, 102–103
referencing between workbooks, 258–262
referencing in multiple worksheets, 258
retrieving values from, 346–347
selecting. *See* selecting cells
certificate authorities, 718
change tracking
accepting and rejecting changes, 737–739
change history log, 733–734
change history report, 735–737
features affected by, 732
highlighting changes, 734–735
overview, 729–730
switching on, 731–732
changing cells (Solver), 605
charts
adding/editing chart data, 470
anchoring, 468
area charts, 486
avoiding values in different scales, 474
bar charts, 484
bubble charts, 492–493
changing type of, 470
charting tables, 471
Chart Tools Ribbon tabs, 465
column charts, 474–475, 482–484
creating standalone charts, 468–469
creating with Ribbon toolbar, 463–465
data series, changing order of, 480
date/time scale data, 476–478
donut charts, 491
drawing graphic objects in, 565
embedded, 463, 471–472
exporting from Excel, 778–780
Insert Chart dialog box, 465
line charts, 474, 484–485
moving and resizing, 466–467
noncontiguous chart ranges, 478–480
overview, 462–463
pie charts, 486
pivot charts, 652–654
plotting blank values, 481–482
printing, 471–472
radar charts, 493
scatter charts, 478
standalone charts, 463, 472
stock charts, 489–490
surface charts, 490
types of, 481
with multiple series of numbers, 473–474
X-axis data, controlling, 473–474
XY Scatter charts, 488–489

charts, formatting
- 3-D charts, 528–531
- adding legends, 500–501
- adding titles, 499–500
- applying shadows, 515
- background colors, 509–510
- borders and lines, 514–516
- chart area, defined, 498
- chart layouts, 497–498
- chart styles, 495–497
- combination charts, 534–536
- controlling scaling, 520–523
- data labels, adding to series, 501–504
- data point labels, adding individual, 504–505
- data series/data points, 516–518
- data tables, adding, 506–507
- elements, overview of, 497–498
- elements, selecting, 507–509
- fancy fills, 510–514
- Format Data Labels dialog box, 502–503
- overlay charts, 532–533
- pie charts, 530–532
- saving formats with templates, 518–519
- scientific error bars, adding, 526–527
- trendlines, adding, 523–525

circular references (formulas), 243
CLEAN() function, 320
clip art library, 551–554
clipboard
- copying multiple items with, 110–112
- pane, 780

clustered column charts, 482
code
- for formatting dates and times, 153–154
- for formatting numbers, 155–156
- formula error codes, 241–242
- header and footer, 220
- for SUBTOTAL() function, 423–424
- VBA coding tips, 830

collaboration, workbook. *See* workbook collaboration
colors
- adjusting in pictures, 544
- background (chart elements), 509–510
- color scales, 578–579
- customizing, 162–167
- sorting tables by, 410

column charts
- example, 474–475
- gaps/widths/overlays in, 532–533
- subtypes of, 482–484

columns
- 3-D, changing shape of, 529–530
- adding/deleting in tables, 404
- adding titles to worksheets, 22–24
- column headers (tables), 398
- column names feature (tables), 421–422
- controlling widths of, 79
- deleting, 115–116
- freezing, 201–204
- grouping, 431–432
- hiding, 204
- inserting new, 113–114
- placing data bars in separate columns, 574–575

combination charts, 534–536
combinations and permutations (functions), 285–286
COMBIN() function, 234, 285–286
comma-delimited files, 789
comments
- creating new, 725–726
- editing, 727–728
- overview, 724
- printing, 729–730
- reviewing, 728
- showing/hiding, 726–727
- VBA, 822

Compatibility Checker, 41–43, 718
compatibility functions, 239–240
compatibility mode, 43
compound documents, 777, 783
compressing pictures, 545–548
concatenation operator (&), 315–316, 832
conditional formatting
- basics, 187
- highlighting specific values with, 187–190
- in older Excel versions, 572
- overview, 186
- rules, editing, 572–573
- Rules Manager, 193
- rules, using multiple, 192–193

conditional functions, 363–365
conditional logic in VBA programming, 836–837
conditional validation rules, 701
configuring Solver, 617–619
connections
- connecting shapes, 564–565
- database connections, 659
- reusing database connections, 668–670
- to SQL Server databases, 665–668

consolidating table data, 438–440
constants, naming, 375–376
constraint operators, 608
constraints (Solver), 606
contextual tabs, 14–15
continuous range cell selection, 95–97
contrast, adjusting (pictures), 544

copy-cut-paste operations
copying formulas, 252–254, 257
copying multiple items with clipboard, 110–112
copying/pasting pictures, 543
drag-and-drop for, 104–105
inserting copied/cut cells, 114–115
on parts of cells, 113
paste options, 107–109
paste preview feature, 106–107
Paste Special options, 109
Paste Values, 108–109
simple example, 102–103
COUNT()/COUNTA()/COUNTBLANK() functions, 269–270
COUNTIF()COUNTIFS() functions, 365–370
counting
items in lists (Math functions), 269–270
and summing using multiple criteria, 368–370
text string characters, 317–319
Create PivotTable dialog box, 626
criteria parameter (database functions), 426
cropping pictures, 547–548
Ctrl+arrow key combinations, 28
CUMIPMT()/CUMPRINC() functions, 304
cumulative interest payments, 304
cumulative principal, 304
Currency format, 144
custom formatting strings
creating, 150–152
date and time format strings, 153–155
number format strings, 155–158
string (formatting) codes, 152
text format strings, 158–159
customizing
AutoFill lists, 76–77
custom list dialog box, 409
dictionaries (spell checking), 89–90
filtering, 396
functions, 839–840
gradient fills (charts), 512
slicer styles, 649
specific workbooks, 847–848
styles, 176–178
templates, 454–457
views, 206–207
workspaces, 207
customizing Ribbon toolbar
changing existing tabs, 850–852
creating personal tabs, 852–853
overview, 848–849
saving/reusing customizations, 853–854
tips for, 851

D

data
adding to worksheets, 24–25, 63–65
changing order of data series (charts), 480
comparing different types of, 368
data connections, 659
data types, variable, 836
editing in worksheets, 25–26
entering in grouped worksheets, 127
grouping. *See* grouping data
hiding, 204–205
importing/exporting, 776
importing/exporting tables of, 786–787
querying, 658
selecting. *See* selecting data
tables, adding to charts, 506–507
visualizing. *See* data visualization
databases
connecting to Access databases, 659–661
connecting to SQL Server databases, 665–668
connections, reusing, 668–670
database functions for tables, 425–427
data source security, 664–665
refreshing data from, 661–663
Research pane, 690
types of, 658
Unlink command, 663
updating data connections automatically, 668
vs. worksheets, 402
Web queries and. *See* Web queries
data labels (charts)
adding to individual points, 504–505
adding to series, 501–504
data points (charts)
defined, 499
formatting, 516–518
data series (charts)
defined, 499
formatting, 516–518
data validation
defined, 456
error alerts, 698–700
with formulas and cell references, 700–702
input message, 696–698
with lists, 702–703
overview, 693–694
settings, 695–696
data visualization
color scales, 578–579
data bars, 570–572
data bars, editing, 573–577
editing conditional formatting rules, 572–573
icon sets, 579–582
sparklines. *See* sparklines

dates and times
calculating days in future, 338
calculating differences between dates, 332–334
calculating future dates, 335
calculating number of days between dates, 335–336
calculating percentage of year between dates, 336–337
calculating week for specific date, 338
in calculations, 327–329
converting into serial numbers, 333–334
counting number of business days, 337
DATEDIF() function, 332–334
DATE() function, 327–329
date leap year bug, 331
date/time scale data (charts), 476–478
DATEVALUE() function, 333–334
DAY() function, 329–330
DAYS360() function, 335–336
filtering dates, 413
finding last day of month, 337
format strings for, 153–155
formatting, 146–148
HOUR() function, 330
identifying in Excel, 68–69
inserting current date and time, 327
manipulating, 324–325
MINUTE() function, 330
MONTH() function, 329–330
regional dating, 70–71
SECOND() function, 330
TIME() function, 327–329
using with ordinary functions, 325–326
WEEKDAY() function, 330–331
YEAR() function, 329–330
DAY() function, 329–330
DAYS()360 function, 335–336
debugging macros, 824–825
decision variables, 605
declining balance depreciation method, 310–311
Defer Layout Update setting, 637
deleting
chart elements, 508
columns, 115–116
database connections, 670
duplicate rows automatically, 417–418
records from tables, 402
rows, 115–116
delimited files, 789
dependents, cell, 387
depreciation, calculating, 309–313
Developer tab, 817, 820
dialog box launchers, 15–16, 140
dictionaries, custom, 89–90
dif constraints, 608
digital certificates, 807
digital signatures, 718
digitally signed workbooks, 807–808
Dim keyword, 835–836
direct precedents (formulas), 391
discount rates, 306
display units (charts), 522
distorting images, 541
Distribute Horizontally/Vertically commands (shapes), 563
documents
compound, 777, 783
creating hyperlinks to, 91–93
document elements (XML), 675
Document Inspector, 718–719
document prolog (XML), 673
document properties, 719–721
trusted, 260
XML, mapping to worksheets, 677–680
DOLLAR() function, 321–323
donut charts, 491
double declining balance depreciation method, 310
double lookup, performing, 347–348
downloading templates, 448–453
drawing
Draw Border feature, 169–171
graphic objects in charts, 565
shapes, 554–557
duplicate rows, 416–418
dynamic calculations (tables), 419–420
dynamic invoicer tutorial, 359–361
dynamic links, creating, 355–356

E

EDATE() function, 335
editing
chart data, 470
collaborative in Excel Web App, 768–769
color scales, 578–579
comments, 727–728
conditional formatting rules, 572–573
data bars, 573–577
data in worksheets, 25–26
embedded objects, 782–783
formulas with point-and-click, 246
icon sets, 581–582
linked objects, 781
shapes, 558
specific cells (macro), 833
styles, 178
tables, 402–404
themes, 183–185
Web queries, 690–691

Edit Links dialog box, 261
EFFECT() function, 300
effects
defined, 180
for pictures, 550–551
themes and, 183–184
elements, XML, 673–675
e-mail, sending workbooks via, 723–724
embedded charts
basics, 463
printing, 471–472
embedding and linking objects
editing embedded objects, 782–783
editing linked objects, 781
overview, 776–778
Encrypt With Password option, 717
end-user databases, 658
Engineering functions, 292
EOMONTH() function, 337
equal sign (=) in formulas, 228
equations, complex (goal-seeking), 603–605
error alerts (data validation), 698–700
error bars, scientific (charts), 526–527
errors in formulas
error-checking feature, 392–394
formula auditing tools, 387–388
overview, 240–243
step-by-step evaluation, 387–388
tracing precedents and dependents, 389–392
ERROR.TYPE() function, 357–358
Evaluate Formula dialog box, 389–390
EVEN() function, 268–269
Evolutionary method (Solver), 617
example files, downloading, 17
Excel 2010
backstage view, 8–9
Excel 2010 Power Programming with VBA (Wiley), 842
exporting charts out of, 778–780
file formats, 39–41
identifying dates and times in, 68–69
identifying numbers in, 67–68
identifying text in, 66–68
importing objects into, 784–785
menu shortcuts in, 8
new features in, 10–11
Options window, 37–38
practical applications of, 2–3
preparing spreadsheets for Excel 2007, 41–42
using as pocket calculator, 233
version number, 2
Web query feature. *See* Web queries
XML mapping feature, 676–681
Excel Web App
browsers and, 750
collaborative editing in, 768–769
overview, 764–765
partially supported features in, 773–774
returning workbooks to desktop Excel, 770–771
saving files in, 768–769
SkyDrive and. *See* SkyDrive
supported features in, 765–766
unsupported features in, 771–772
uploading files, 751–752
vs. workbook sharing, 747
Web workflow, 753
exit conditions (loops), 838
exporting
charts from Excel, 778–780
Ribbon customizations, 853
XML, 681
extrapolation (trendlines), 525

F

FACT() function, 280
FALSE() function, 364
fax provider plug-in, 724
fields
defined, 396
field parameter (database functions), 426
rearranging in pivot tables, 633
files
formats for, 39–41
options for opening/saving, 61. *See also* opening files
saving, 37–39
uploading to Web. *See* Excel Web App
viewing in SkyDrive, 760–762
workspace, 209
XML, 675–676
fills
and borders, formatting, 167–171
fancy (charts), 510–514
filtering
custom (tables), 396
dates, 413
names, 377–378
numbers, 414
pivot tables. *See* filtering pivot tables
tables, 405–406
tables with list of values, 411–413
text, 415–416
filtering pivot tables
group filtering, 650–652
report filtering, 643–645
with slicers, 645–649

financial functions
depreciation, calculating, 309–313
figuring interest rates for future values, 305–306
future values, calculating, 297–300
internal rate of return, calculating, 306–308
loan payoffs/investment targets, calculating, 304
net present value, 306–308
number of loan payments, calculating, 302–304
overview, 295–296
specialized functions, 312–313
Find and Replace feature
advanced searches, 130–132
basics, 128–130
Find All feature, 129–130
FIND() function, 317–319
finding formatted cells, 133–134
finding/replacing values, 134–135
tips for applying, 136
using wildcards, 132
finding
macros, 824–825
numbers frequently occurring in lists, 279–280
FIXED() function, 321–323
fixed width columns, 789
FLOOR() function, 265–266
fonts
formatting, 162–167
in themes, 181–182
footers and headers, 220–221
For Each loops, 838–839
FORECAST() function, 289–291
Format Axis dialog box, 522
formatting
applying to individual characters, 165
borders and fills, 167–171
cells automatically, 233
cells containing cell references, 232–233
cells with macros, 834–835
charts. *See* charts, formatting
conditional. *See* conditional formatting
file formats, 39–41
Format Painter, 173–174
pivot tables, 632–633
styles, 174–179
tables, 399–402
themes, 180–186
formatting cells
alignment and orientation, 159–162
cell values, 138–140
custom formats, 150–157
dates and times, 146–148
finding formatted cells, 133–134
fonts and colors, 162–167
grouped worksheets, 126
numbers, 141–147
overview, 137–138
special numbers, 149
with macros, 834–835
formulas
applying names to existing, 380–381
array formulas, 253
cell references and, 231–232
circular references and, 243
conditions in, 363–365
copying, 252–254, 257
creating basic, 227–229
data validation with, 700–702
effects of changing cell locations, 262
errors in creating, 240–243
formatting cells containing cell references, 232–233
Formula AutoComplete, 234
formula bar, 31–32
Formulas tab, 247–249
functions. *See* functions
linked, evaluating, 389
logical operators, 243–245
naming, 375–376
order of operations, 229–231
point-and-click method for creating, 245–246
showing and printing, 245
solving errors in. *See* errors in formulas
storing, 255
sub-expressions in, 387–389
Fraction format, 144–145
freezing columns/rows, 201–204
FREQUENCY() function, 277–279
Full Screen view, 198
functions
basics of, 234–235
compatibility functions, 239–240
creating custom, 839–840
defined, 229
Engineering functions, 292
financial. *See* financial functions
Function Arguments dialog box, 250–251
Function Library, 247–248
information functions. *See* information functions
Insert Function feature, 249–252
learning new, 235
mathematical. *See* Math and Trig functions
reference functions. *See* lookup and reference functions

functions (*continued*)
that return arrays, 253
specialized functions, 236
statistical functions, 289–292
using cell ranges with, 237–238
using cell references with, 236
using in formulas, 234–236
Future Value (FV)
calculating, 297–300
defined, 295
FV() function, 297–300

G

gap width (column charts), 533–534
GCD() function, 284–285
General format (numbers), 142
genetic algorithms, 617
Goal Seek tool
basics, 599–602
complex equations and, 603–605
Go To feature
basics, 28–29
selecting cells with, 101
gradient fills (charts), 511–512
graphics. *See also* **clip art;** *See also* **drawing;** *See also* **pictures**
drawing objects in charts, 565
SmartArt feature, 565–568
vector graphics, 551
greatest/least common denominator functions, 284–286
GRG Nonlinear method (Solver), 617–618
gridlines, defined, 498–499
grids, moving around, 26–28
grouping
group filtering of pivot tables, 650–652
pivot tables vs., 625
slices in pie charts, 531–532
worksheets, 125–127
grouping data
automatic subtotaling, 441–444
Auto Outline feature, 441
collapsing/expanding multiple data groups, 436
combining data from multiple tables, 438–440
creating groups, 430–433
nesting groups within groups, 434–437
overview, 429–430
summarizing data, 436–437
tables vs. grouping, 430

H

handles, resizing, 540–541
headers and footers
creating, 220–221
inserting headers (macro example), 811–812
heading fonts, 181
hidden and locked cells, 703–704
hidden workbooks, 821–822
hiding
comments, 726–727
data, 204–205
icons, 583–584
and showing details in pivot tables, 637–638
worksheets, 121
highlighting changes (change tracking), 734–735
highlighting duplicate rows, 416–417
HLOOKUP() function, 340–341, 344–345
horizontal alignment of cells, 159–160
horizontal axis, defined (charts), 498
horizontal lookup function, 340–341, 344–345
horizontal splits (windows), 198–200
HOUR() function, 330
HTML files, saving spreadsheets as, 47–50
hyperlinks
adding to spreadsheets, 90–94
HYPERLINK() function, 355–356
to locations within worksheets, 93–94
to Web pages or documents, 91–93

I

icon sets, 579–582
IF() statements, nesting multiple, 365–366
importing/exporting
data, 776
XML, 681
exporting tables of data, 786–787
importing objects into Excel, 784–785
importing text files, 787–791
importing text into existing files, 792
INDEX() function, 346–347
INDIRECT() function, 351–353
information functions
checking values inside cells, 356–358
defined, 356
finding data/error type of values, 357–358
INFO() function, 358–359
obtaining cell information, 358–359
obtaining computer/operating system info, 358–359
overview, 339–340
Information Rights Management (IRM), 718, 708

input message (data validation), 696–698
inserting
copied/cut cells, 114–115
Insert Chart dialog box, 465
Insert Function feature, 249–252
Insert Sheet command, 119
Insert Sheet Rows command, 470
new columns, 113–114
new rows, 114
installed templates, 454
int constraints, 608
integer division, 281
interest payments, cumulative, 304
interest rates
monthly vs. annual, 300
to achieve future values, 305–306
intersection operators (spaces), 237
intersection, range, 381
INT() function, 266–267
investment targets, calculating time for, 304
invoices, generating from product catalog (tutorial), 359–361
IPMT() function, 302–304
IRR() function, 306–308
IS functions, 356–358

K

keyboards
keyboard accelerators, 5–6
selecting cells with, 100–102
keyboard shortcuts
for moving around grid, 26–27
overview, 8
for selecting cells, 100–101
keys, sorting, 406

L

LARGE() function, 272–273
last-in wins updating, 769–770
layering in worksheets, 562
laying out
charts, 497–498
pivot tables, 629–631
LCM() function, 284–285
leap year bug, 331
LEFT() function (text), 316–317
left lookup, performing, 346–347
legends (charts)
adding to charts, 500–501
defined, 498
LEN() function, 317–319
linear forecast trends (charts), 525
line charts, 474–475, 484–485
lines and borders (charts), 514–516
linking. *See also* **hyperlinks**
cells between workbooks, 258–262
dynamic links, creating, 355–356
linked formulas, evaluating, 389
objects. *See* embedding and linking objects
worksheet links to pictures, 540
lists
data validation with, 702–703
mapping in XML documents, 683–686
literals, defined, 157
live preview
defined, 766
turning off, 180
loan payoffs, calculating time for, 304
local templates, 455
locked and hidden cells, 703–704, 707
logarithmic scales, 522
logical operators (formulas), 243–245
logic, conditional, 836–837
logs, change history, 733–734
lookup and reference functions
advanced text lookup, 343
cell references stored as text, 351–353
creating dynamic links, 355–356
dynamic invoice tutorial, 359–361
finding position of items in ranges, 345–346
HLOOKUP() function, 340–341, 344–345
miscellaneous, 350–351
overview, 339–340
performing double lookup, 347–348
performing left lookup, 346–347
retrieving values from cells, 346–347
shifting cell references, 349–350
transposing rows and columns, 353–354
VLOOKUP() function, 340–343
loop statements (VBA), 838–840
LOWER() function, 319–320

M

macros. *See also* **VBA programming**
alternating row formatting (example), 812–813
attaching to buttons inside worksheets, 816–818
basics of, 793–794
code analysis, 826–827
combined tasks (example), 813
debugging, 824–825
defined, 33
finding/moving, 824–825
formatting cells with, 834–835
inserting headers (example), 811–812
macro code, 57–58
macro-free/macro-enabled workbooks, 795–797

macros (*continued*)
- macro recorder, 793–794, 795–801
- modules and, 822–823, 827
- personal macro workbook, 799
- placing on Quick Access toolbar, 814–816
- playing, 802–803
- shortcut dangers, 800
- stored location of, 797–798
- using variables in, 834–836

macro security
- Authenticode signing, 811
- overview, 803–804
- setting up trusted locations, 809–811
- temporary trust of documents, 806–807
- Trust Center, 806–808
- trusted documents, 804–805

MakeCert.exe tool, 807
manual page breaks, 222–223
mapping, XML, 676–680, 682–685
margins, adjusting (printing), 213–217
Mark As Final option, 717
markers, sparkline, 590–591
MATCH() function, 345–346
Math and Trig functions
- absolute values of numbers, 281
- advanced math functions, 286–289
- advanced ranking functions, 275–277
- Auto-Sum button, 269
- combinations and permutations, 285–286
- counting items in lists, 269–270
- determining signs of numbers, 281
- FACT() function, 280
- finding average or median values, 274–275
- finding maximum/minimum values, 271
- finding numbers in lists, 279–280
- generating random numbers, 282–284
- greatest/least common denominators, 284–285
- higher division functions, 281
- placing numbers in grouped ranges, 277–279
- POWER() function, 280–281
- PRODUCT() function, 280
- ranking numbers, 272–273
- rounding numbers, 263–268
- SQRT() function, 280
- summing up numbers, 268–269
- trigonometry and advanced math, 286–289

maximum/minimum values, finding, 271
MAX()/MIN() functions, 271
median values, finding, 274–275
merging
- cells, 162
- multiple workbook revisions, 739–741
- scenarios, 600

methods and properties (VBA), 829–830
mht web archive format, 51
MID() function (text), 316–317
MINUTE() function, 330
MissingManuals.com, 17
MODE() function, 279–280
MOD() function, 281
modules (macros), 822–823, 827
modulus, defined (math functions), 281
MONTH() function, 329–330
MROUND() function, 265–266
multi-layered pivot tables, 635–636
multiple sort criteria (tables), 407–409

N

named ranges
- automatically creating, 378–380
- creating and using, 372–374
- creating smarter, 373–374
- defined, 371
- managing, 376–378

naming
- applying names to existing formulas, 380–381
- formulas/constants, 375–376
- generating lists of names, 378
- Name Manager, 376–377
- table names, 421–422
- worksheets, 121–122

navigating through comments, 728
nesting
- data groups within groups, 434–437
- multiple IF() statements, 365–366
- XML elements, 673–674

net present value, calculating (investments), 306–308
NETWORKDAYS() function, 337
New Formatting Rule dialog box, 190
New Name dialog box, 375–376
noncontiguous cell selections, 99–100
noncontiguous chart ranges, 478–480
non-printing characters, removing, 320
NOT() function, 364
NOW() function, 327
NPER() function, 304
NPV() function, 306–308
Number of Payment Periods (NPER), 296
numbers
- absolute values of, 281
- converting to text, 321–323
- decimal precision of, 145
- filtering, 414
- formatting, 141–147
- functions for rounding, 263–266
- generating random, 282–284
- identifying in Excel, 67–68

placing in grouped ranges, 277–279
ranking, 272
using positive and negative (financial functions), 296

O

objective cells (Solver), 605
object model, 828
objects, basics of, 828–829
ODD() function, 268–269
offsets
cell, 255
OFFSET() function, 349–350
Offset property, 832–833
one-variable data tables, 382–384
opening files
menu options for, 58–59
multiple spreadsheets at once, 60–61
overview, 55–56
preferences for opening/saving files, 61
protected view, 57–58
recent documents, 56–57
operators
arithmetic, 229
logical, 243–245
Options window, 37–38
order of operations (formulas), 229–231
OR() function, 364
outlining
automatic subtotaling and, 441–444
Auto Outline feature, 441
overlay charts, 532–533

P

page breaks, 222–225
Page Layout view, 217–220
Page Setup dialog box, 215–217, 729
pagination (printing), 222–225
panes, multiple window, 198
partially fixed cell references, 256
passwords
Encrypt With Password option, 717
saving spreadsheets with, 50–52
workbook protection and, 709
Paste Special function, 543, 781
pasting cells. *See* **copy-cut-paste operations**
payments
calculating number of, 302–304
Payment (PMT), defined, 296
PDF files, saving spreadsheets as, 44–46
Percentage format, 144
percentile calculation functions, 275–277
permissions, user account, 711–714
PERMUT() function, 285–286
personal macro workbook, 799
pictures
borders/effects/styles, 550–551
clip art library, 551–554
compressing, 545–548
copying/pasting, 543
cropping/shaping, 547–548
inserting in worksheets, 538–540
picture fills (charts), 513–514
positioning/resizing, 540–542
touching up, 543–545
pie charts
exploding slices in, 530–531
grouping slices in, 531–532
subtypes of, 486
pinning documents, 56
pivot tables
advantages of (example), 622–624
building, 624–625
discovering trends and relationships in, 633–634
filtering. *See* filtering pivot tables
fine tuning calculations in, 639–643
formatting, 632–633
hiding/showing details in, 637–638
laying out, 629–631
multi-layered, 635–636
overview, 621
pivot charts, 652–654
preparing, 626–628
rearranging, 633
regions (fields) of, 628–629
Scenario PivotTable report, 598
slow refreshing of, 637
vs. grouping and subtotals, 625
playing macros, 802–803
plotting
blank values (charts), 481–482
defined, 462
plot area, defined, 498
PMT() function, 302–304
pocket calculator, using Excel as, 233
point-and-click formula creation, 245–246
points, connection (shapes), 564–565
polynomial trends (charts), 525
positioning images, 542
positive/negative numbers (financial functions), 296
Power and Exponential trends (charts), 525
POWER() function, 280–281
PPMT() function, 302–304
precedents
and dependents, tracing, 389–392
defined (formulas), 387

present values
calculating (finance), 300–302
Present Value (PV), defined, 295, 306
principal, cumulative, 304
printing
charts, 471–472
comments, 729–730
Excel files, procedure for, 211–217
formulas, 245
headers and footers, 220–221
overview, 210
Page Break Preview, 224–226
page breaks, 222–223
Page Layout view, 217–220
pagination, 222–225
parts of spreadsheets, 217
previewing, 217–220
scaling, 223
PRODUCT() function, 280
programming with VBA. *See* **VBA programming**
prolog, document (XML), 673
PROPER() function, 319–320, 839
properties and methods (VBA), 829–830
protection (security)
of cell ranges with passwords, 709–711
of online workbooks, 753
protected view, 57–58
Protect Structure and Windows dialog box, 708
Protect Workbook button, 716–721
of workbooks, 707–708
of worksheets. *See* worksheet protection
PV() function, 300–302

Q

quadratic function, 603
querying data, 658
Quick Access toolbar (QAT)
adding buttons to, 845–847
adding favorites to, 843–845
overview, 10–11
placing macros on, 814–816
QUOTIENT() function, 281

R

R1C1 referencing, 27, 255
radar charts, 493
RAND()/RANDBETWEEN() functions, 282–284
ranges
of cells. *See* cell ranges
named. *See* named ranges
noncontiguous chart ranges, 478–480
range intersection, 381
Range object, 833
range operator (colon), 237
table_range parameter (database functions), 426
ranking
advanced functions, 275–277
numbers (Math functions), 272–273
RANK() functions, 272–274
rates
defined (financial functions), 296–297
RATE() function, 305–306
recalculation, controlling, 386–387
recolor feature (pictures), 544
recorder, macro, 793–794, 795–801
records
adding to tables, 402
defined, 396
deleting from tables, 402–403
Redo/Undo features, 79–81, 163
reference functions. *See* **lookup and reference functions**
refreshing database data, 661–663
regional dating, 70–71, 147
relative cell references, 252
relative reference mode (macros), 796–797
Remove Duplicates feature (rows), 418
Repeat function, 81
reports
change history report, 735–737
report filtering (pivot tables), 643–645
summary reports (scenarios), 598–600
Research pane, 690
Research window, 88
resizing
chart elements, 508
charts, 466–467
pictures, 540–542
reviewing comments, 728
Ribbon toolbar
applying styles with, 176
buttons with menus, 14
collapsing, 30–31
contextual tabs in, 14–15
creating charts with, 463–465
customizing. *See* customizing Ribbon toolbar
Formulas tab, 247–249
numerical formatting shortcuts in, 146
overview, 4–5
tabs overview, 29–31
using with keyboard, 5–6
RIGHT() function (text), 316–317
root element (XML), 675
rounding numbers
arithmetic rounding, 267
banker's rounding, 267, 841

functions for, 263–266
VBA Round() function, 841
rows
deleting, 115–116
duplicate, 416–418
freezing, 201–204
grouping, 433–434
hiding, 204
inserting new, 114
rules, conditional formatting, 188, 190–192

S

Safari Books Online, 17
saving
chart formats with templates, 518–519
files, 37–39
files in Excel Web App, 768–769
options for opening/saving files, 61
Ribbon toolbar customizations, 853–854
Save to Web option, 724
Solver models, 615–617
spreadsheets as HTML files, 47–50
spreadsheets as PDF files, 44–46
spreadsheets for Excel 2003, 43–44
spreadsheets in various formats, 44
spreadsheets with passwords, 50–52
view settings, 205–207
scaling
charts, 520–523
data bars, 575–578
for printing, 223
scatter charts, 478
scenarios
creating, 594–598
managing with Scenario Manager, 597
merging/transferring, 600
overview, 594
storing Solver results in, 611
summary reports, 598–600
schemas, XML, 675–676, 682
scientific error bars (charts), 526–527
Scientific format, 146
SEARCH() function, 317–319
searching. *See* **Find and Replace feature**
SECOND() function, 330
security. *See also* **data validation;** *See also* **worksheet protection**
data connections and, 664–665
macros and. *See* macro security
selecting
data with AutoSelect, 99
parts of tables, 404–405
Selection and Visibility pane, 561–562
shapes, 560–563
selecting cells
continuous range selection, 95–97
noncontiguous selections, 99–100
with Go To feature, 101
with keyboard, 100–102
serial numbers, 324
server-side databases, 658
shapes
adding text to, 559–560
connecting, 564–565
drawing, 554–557
editing, 558
lining up, 563
selecting/arranging, 560–563
shaping pictures, 547–548
SharePoint, uploading to, 724, 759–760
sharing
templates, 458
workbooks. *See* workbook sharing
shortcuts
keyboard. *See* keyboard shortcuts
macro, 800
shrinking text, 162
SIGN() function, 281
simple operators, 229–230
Simplex LP method (Solver), 617
SkyDrive
creating workbooks in, 763
overview, 751–752
using browsers to upload files to, 753–755
using Excel to upload workbooks to, 756–758
viewing files in, 760–762
slicers, pivot table, 645–649
slices, grouping in pie charts, 531–532
SLN() function, 310
SMALL() function, 272–273
SmartArt diagrams, 565–568
smart tags, 109
smooth nonlinear functions, 617
Solver feature
advanced problems, 612–614
configuring, 617–619
defining a problem in (example), 606–611
overview, 604
restoring Solver models, 616
saving Solver models, 615–617
Solver model components, 605
Solver Parameters dialog box, 609–610
solvsamp.xls workbook (examples), 615
storing results in scenarios, 611

sorting tables
- applying simple order, 406
- by color, 410
- with custom lists, 409
- defined, 405
- with multiple criteria, 407–409
- sorting keys, 406

spaces
- as intersection operators, 237
- removing, 320
- usage in Excel, 229

sparklines
- basics of, 583–584
- creating, 585–588
- defined, 569
- markers, 590–591
- options for changing axes, 588–591

special characters (font symbols), 164–166

Special number format, 149

spell checking, 87–91

splitting windows, 198–200

spreadsheets. *See also* **worksheets**
- adding hyperlinks to, 90–94
- opening multiple at once, 60–61
- preparing for Excel 2007, 41–42
- printing parts of, 217
- programming with VBA. *See* VBA programming
- saving as HTML files, 47–50
- saving as PDF files, 44–46
- saving for Excel 2003, 43–44
- saving in various formats, 44
- saving with passwords, 50–52

SQL Server databases, connecting to, 665–668

SQRT() function, 280

stacked column charts, 482

stacked line charts, 485–486

standalone charts
- creating, 463, 468–469
- printing, 472

statistical functions, 289–292

status bar
- calculation tool in, 98
- overview, 32–35

stickiness (tables), 405

stock charts, 489–490

Stop If True column setting, 193

straight-line depreciation method, 309

string formatting codes, 152

string literals, 316

structured tables, 395

styles
- chart styles, 495–497
- formatting, 174–179
- for pictures, 550–551
- table styles, 399–400

sub-expressions (formulas), 387–389

subroutines (VBA), 822–823

SUBSTITUTE() function, 320

SUBTOTAL() function, 423–424, 437–438

subtotaling
- automatic, 441–444
- pivot tables vs. subtotals, 625

SUM() function, 238–239, 268–269

SUMIF()/SUMIFS() functions, 367–370

summary pivot tables (scenarios), 598

summary reports (scenarios), 598–600

summary tables, 622–624

summing/counting using multiple criteria, 368–370

sum-of-years-digits depreciation method, 310

surface charts, 490

symbols, font, 164–166

T

tables
- basics of, 396–397
- charting, 471
- creating, 396–398
- of data, importing/exporting, 786–787
- editing, 402–404
- filtering, 405–406
- filtering with list of values, 411–413
- formatting, 399–402
- grouping vs., 430
- pivot. *See* pivot tables
- selecting parts of, 404–405
- sorting. *See* sorting tables
- styles, 399–400
- summary. *See* summary tables

table calculations
- basics, 418–419
- column names feature, 421–422
- database functions, 425–427
- dynamic calculations, 419–420
- SUBTOTAL() function, 423–424
- table names and, 421–422
- total rows, 422

tabs, Ribbon, 29–31

target cells (Solver), 605

templates
- basics, 446
- creating, 454
- creating new workbooks from, 447–450

custom, 454–457
downloading from backstage view, 448–451
downloading from Office Online, 450–452
locating Templates folder, 455
overview, 445
saving chart formatting with, 518–519
searching for, 451
sharing with others, 458–460
types of, 454
workgroup, 455–456

temporary trust of documents (macros), 806–807

text
adding to shapes, 559–560
entering in current cell (macro), 831–832
files, importing/exporting, 787–791
filtering, 415–416
identifying in Excel, 66–68
importing into existing files, 792
shrinking, 162
Text format for numbers, 146–147
text format strings, 158–159

text functions
capitalization, changing, 319–320
concatenation operator, 315–316
converting text to numbers, 321–323
copying text strings, 316–317
counting string characters, 317–319
joining text strings, 316
miscellaneous, 323–324
overview, 315–316
removing spaces/non-printing characters, 320
replacing character sequences, 320
TEXT() function, 321–323

texture fills (charts), 511–512

themes (formatting)
basics, 180–184
colors in, 182–183
modifying, 183–185
shared with Office programs, 186

thesaurus feature, 88
tick marks (scaling), 521
times. *See also* **dates and times**
format strings for, 153–155
formatting, 146–148
identifying in Excel, 68–69
TIMEVALUE() function, 333–334

titles (charts)
adding, 499–500
defined, 498

TODAY() function, 327
total rows (tables), 422
touching up pictures, 543–545
tracing precedents and dependents (formulas), 389–392
tracking changes. *See* **change tracking**
translation button, 88
Transparent Color (pictures), 544–545
TRANSPOSE() function, 353–354
TREND() function, 289–291
trendlines, adding to charts, 523–525
trigonometry functions, 286–289
TRIM() function, 320
troubleshooting links between workbooks, 261–262
TRUE() function, 364
TRUNC() function, 266–267
Trust Center (macros), 806–808
trusted documents, 260, 804–805
trusted locations, 809–811
two-dimensional pivot tables, 631
two-dimensional summary tables, 624
two-variable data tables, 384–385
TYPE() function, 357–358

U

unchanging cells (Solver), 605
Undo/Redo features, 79–81
union operator (comma), 237

uploading
files to SkyDrive with browsers, 753–755
workbooks to SharePoint servers, 759–760
workbooks to SkyDrive with Excel, 756–758
workbooks to Web, 724

UPPER() function, 319–320
user accounts permissions, 711–714
user-defined functions, 839
user names in Excel, 716

V

values
cell, formatting, 138–140
filtering tables with list of, 411–413
finding/replacing, 134–135
numerical, formatting, 143
value axis (charts), 474, 498
VALUE() function, 321–323

variables
using in macros, 834–836
variable data tables, 382–385

VBA programming. *See also* **macros**
accessing cells with Offset property (macro), 832–833
coding tips, 830
conditional logic in, 836–837
custom functions, creating, 839–840
editing specific cells (macro), 833
entering text in current cell (macro), 831–832
formatting cells with macros, 834–835
loop statements, 838–840

VBA programming (*continued*)
macro code analysis, 826–827
objects, basics of, 828–829
Programming Excel with VBA and .NET (O'Reilly), 842
properties and methods, 829–830
variables in macros, 834–836
VBA, defined, 793
Visual Basic editor. *See* Visual Basic editor
With statements, 830
vector graphics, 551
venture investment, 306
VeriSign, 807–808
vertical alignment of cells, 160–161
vertical axis, defined (charts), 498
vertical lookup function, 340–343
vertical splits (windows), 198–200
viewing
files in SkyDrive, 760–762
multiple workbooks, 207–209
viewing worksheets
custom views, creating, 206–207
freezing columns/rows, 201–204
Full Screen view, 198
hiding data, 204–205
Page Layout view, 217–220
saving view settings, 205–207
viewing distant parts at once, 198–200
zoom feature, 196–198
Visio diagramming tool, 785
Visual Basic editor
debugging macros, 824–825
finding/moving macros, 824–825
modules and macros, 822–823
overview, 820
Project window, 820–822
visual fidelity, 764
visualizing data. *See* **data visualization**
Visual Studio Tools for Office (VSTO), 842
VLOOKUP() function, 340–343
volatile functions, defined, 282

W

Web Application, Excel. *See* Excel Web App
web archive format, 51
Web pages, creating hyperlinks to, 91–93
Web queries
creating, 687–690
editing, 690
limitations of, 687–688
overview, 686–687
websites, for further information
Verisign, 807
XML and XML schemas, 676
WEEKDAY() function, 330–331
WEEKNUM() function, 338
Windows Live ID, 751
Windows (Microsoft), sharing information within, 775–776
windows, splitting, 198–200
With statements, 830
workbooks
creating from templates, 447–450
creating in SkyDrive, 763
customizing specific, 847–848
digitally signed, 807–808
hidden, 821–822
macro-free/macro-enabled, 795–797
moving worksheets between, 123–124
overview, 118–119
protecting, 707–708
referencing cells between, 258–262
refreshing data from databases, 661–663
returning to desktop Excel from Web, 770–771
starting new files, 21
transferring styles between, 179–180
uploading to SharePoint servers, 759–760
uploading to SkyDrive using Excel, 756–758
viewing multiple, 207–209
workbook collaboration
change tracking. *See* change tracking
comments. *See* comments
in Excel Web App, 768
merging multiple workbook revisions, 739–741
overview, 715–716
sending by e-mail, 723–724
uploading files to Web, 724
workbook sharing. *See* workbook sharing
workbook sharing
limitations of, 745–747
multiple users not using, 742–743
overview, 741–742
turning on, 743–745
vs. Excel Web Application, 747
WORKDAY() function, 338
workflow, Web, 753
workgroup templates, 455
worksheets. *See also* **spreadsheets**
adding backgrounds to, 172
adding column titles to, 22–24
adding data to, 24–25
adding different data types to, 63–65
adding/removing, 119–120
attaching macros to buttons inside, 816–818
creating basic, 19–21
creating hyperlinks to locations within, 93–94
editing data, 25–26
Go To feature, 28–29
grouping, 125–127

hiding, 121
inserting clip art into, 551
inserting pictures in, 538–540
mapping XML documents to, 677–680
moving around grid, 26–28
moving between workbooks, 123–124
naming and rearranging, 121–122
overview, 118–119
protection feature, 456
referencing cells in other, 258
starting new, 21
tabs, changing colors of, 123
viewing. *See* viewing worksheets
vs. databases, 402
worksheet protection
locked and hidden cells, 703–704
overview, 693–694
steps to implement, 705–707
workspaces
custom, 207
workspace files, 209

X

X-axis chart data, controlling, 473–474
XIRR() function, 309
xls and xlsx formats, 39–41
xlsb files, 795
xlsm files, 795
XLStart folder, 799
XML (Extensible Markup Language)
files and schemas, 675–676
importing/exporting, 681
mapping feature in Excel, 676–680, 683–686
mapping with schemas, 682
overview, 671–672
tags, 674
three rules of, 672–675
XNPV() function, 309
XY Scatter charts, 488–489

Y

YEARFRAC() function, 336–337
YEAR() function, 329–330

Z

zoom feature (viewing), 196–198