Excel
eBook Sampler

Visit **quepublishing.com/excel**

TABLE OF CONTENTS

SAVE 40%
Enter discount code **EXCEL**

Visit **quepublishing.com/excel**

que

Shop & Save

QUE RESOURCE CENTERS

Microsoft Excel Save 40%	Microsoft Office Save 40%
Microsoft Windows 8.1 Save 40%	Apple Save 35%
Samsung Save 35%	Minecraft Save 30%
Senior Savings Save 40%	Maker / DIY Save 35%
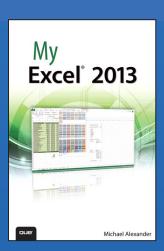
Free U.S. Shipping

Visit **quepublishing.com**

My Excel® 2013

Michael Alexander

CHAPTER 1:

Working with Excel's Ribbon Menus
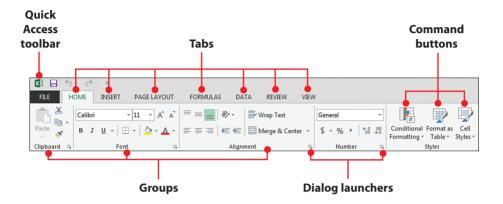
The friendly, full-color, task-based guide to Microsoft Excel 2013

- Organized and formatted for easy, quick reference by any user, regardless of experience

- Presents hands-on, task-focused coverage of the 100+ Excel 2013 tasks every user needs to master

- Step-by-step techniques show the easiest, fastest, most direct way to accomplish each task

- Carefully annotated, numbered full-color screenshots make sure you always know exactly what to do

- Written by world-class Excel trainer and author Michael Alexander

ABOUT THE AUTHOR

Michael Alexander has been named a Microsoft MVP for his ongoing contributions to the Excel community. A Microsoft Certified Application Developer (MCAD), he has authored several books on advanced business analysis with Excel and Access, including Easy Microsoft Excel 2010. He has 15+ years' experience developing Office solutions, and runs the free Excel and Access tutorial and tips site datapigtechnologies.com.

QUE®

Quick Access toolbar

Tabs

Command buttons

| FILE | HOME | INSERT | PAGE LAYOUT | FORMULAS | DATA | REVIEW | VIEW |

Calibri | 11 | General

Paste | B I U | Wrap Text | Merge & Center | $ % , | Conditional Formatting | Format as Table | Cell Styles

Clipboard | Font | Alignment | Number | Styles

Groups

Dialog launchers

In this chapter, you are introduced to Excel's Ribbon menu and learn the basics of working with workbooks and worksheets. Topics covered in this chapter include

➡ Familiarizing yourself with the Ribbon tabs
➡ Understanding contextual tabs
➡ Understanding workbooks and worksheets
➡ Customizing the Quick Access Toolbar

1

Working with Excel's Ribbon Menus

Like any other application, Excel has a basic workspace called the user interface. A user *interface* is the combination of screens, menus, and icons you use to interact with an application. In Excel, the user interface is primarily composed of the Ribbon menu, workbooks, and worksheets. The Ribbon is the name given to the row of tabs and buttons you see at the top of Excel. The Ribbon's tabs and buttons bring your favorite commands into the open by showing multiple commands grouped in specific categories.

The Ribbon is made up of five basic components: the Quick Access Toolbar, tabs, groups, command buttons, and dialog launchers.

- The **Quick Access Toolbar** is essentially a customizable toolbar to which you can add commands that you use most frequently.

- **Tabs** contain groups of commands that are loosely related to core tasks. Actually, it helps to think of each tab as a category.

- **Groups** contain sets of commands that fall under the umbrella of that tab's core task. Each group contains buttons, which you click to activate the command you want to use.
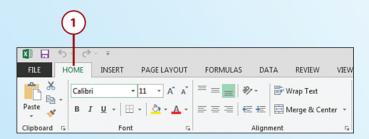
- **Dialog launchers** are activated by clicking the small arrow located in the lower-right corner of certain groups. Clicking any dialog launcher activates a dialog box containing all the commands available for a given group.

Familiarizing Yourself with the Ribbon Tabs

Each tab on the Ribbon contains groups of commands loosely related to a central task. Don't be alarmed by the number of commands on each tab. As you go through this book, you'll quickly become familiar with each of the common Excel commands. For now, take some time to become familiar with each of Excel's default tabs (the way they are set up before you customize it to fit your working needs).
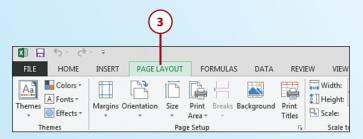
1. Click the Home tab. This tab contains commands for common actions such as formatting, copying, pasting, inserting, and deleting columns and rows.
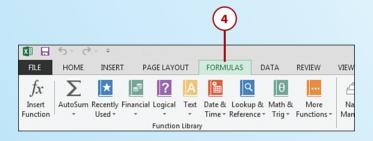
2. Click the Insert tab. This tab contains commands that enable you to insert objects such as charts and shapes into your spreadsheets.
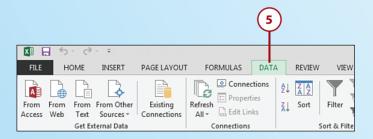
3. Click the Page Layout tab. This tab holds all the commands that enable you to determine how your spreadsheet looks, both onscreen and when printed. These commands control options such as theme colors, page margins, and print area.
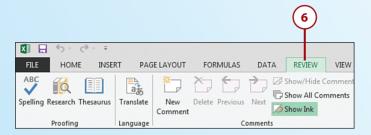
4. Click the Formulas tab. This tab holds all the commands that help define, control, and audit Excel formulas.
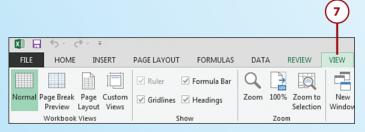
5. Click the Data tab. This tab features commands that enable you to connect to external data, as well as manage the data in your spreadsheet.
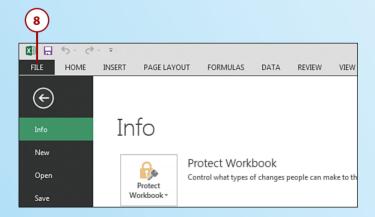
6. Click the Review tab. With commands such as Spell Check, Protect Sheet, Protect Workbook, and Track Changes, the theme of the Review tab is protecting data integrity in your spreadsheet.

7. Click the View tab. The commands on this tab are designed to help you control how you visually interact with your spreadsheet.

8. The File tab exposes the Backstage view, where you find commands to help you open existing Excel workbooks, create new workbooks, save workbooks, apply protection, and much more. Chapter 2, "Managing Workbooks and Worksheets," covers the File tab in detail.

Minimizing the Ribbon

If you feel the Ribbon takes up too much space at the top of your screen, you can minimize it—that is to say, you can hide the Ribbon and show only the tab names. To do this, simply right-click the Ribbon and select Unpin the Ribbon. At this point, the Ribbon becomes visible only if you click on a tab. To revert back to the default Ribbon view, right-click the Ribbon, and click the Unpin the Ribbon toggle.
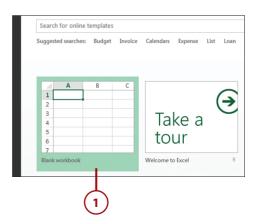
Understanding Workbooks and Worksheets

An Excel file, often referred to as a *workbook*, contains one or more spread-sheets, or *worksheets*. Each box in the worksheet area is referred to as a *cell*. Each cell has a *cell address*, which is composed of a *column reference* and a *row reference*. The letters across the top of the worksheet make up the column reference. The numbers down the left side of the worksheet make up the row reference. For example, the address of the top, leftmost cell is A1. This is because the cell is located at the intersection of the A column and row 1.

By default, Excel 2013 opens a new workbook with one blank worksheet. You can add, delete, and rename worksheets within a workbook, as needed.

Explore Worksheets
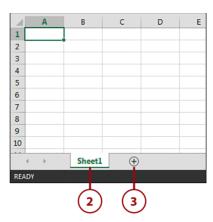
1. Open Excel and open a new Blank workbook.
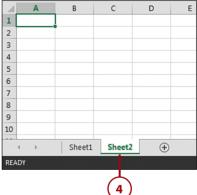
2. The workbook opens with one worksheet called Sheet1. This worksheet contains cells you can use to start entering and editing data.

3. Click the plus icon to add a new worksheet.

4. The new worksheet is added and named Sheet2. Each time you add a worksheet, Excel gives the worksheet a default name of Sheet XX, where XX is the next sequential number.

Workbooks and Worksheets in Depth
You read more about managing workbooks and worksheets in detail in Chapter 2.
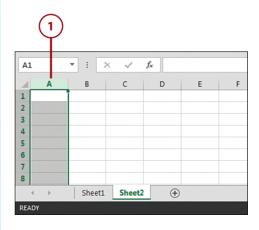
Explore Columns, Rows, and Cells
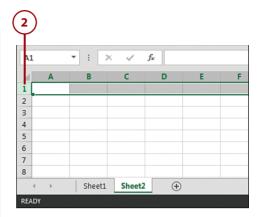
1. Click the column reference and observe how the entire column is selected.
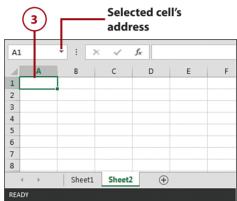
2. Click the row reference and observe how the entire row is selected.

3. Click the cell intersecting at column A and row 1. You can select a single cell on the worksheet area.

The Big Grid

Excel 2013 has 1,048,576 rows and 16,384 columns. This means that there are more than 17 billion cells in a single Excel worksheet. So what happens if you need more than 1 million rows and more than 16,000 columns? Well, the short answer is that you'll have to start a new worksheet.

Selected cell's address

Understanding Contextual Tabs

Contextual tabs are special types of tabs that appear only when a particular object is selected, such as a chart or a shape. These contextual tabs contain commands specific to whatever object you are currently working on.

For example, after you add a shape to a spreadsheet, a new Format tab appears. This is not a standard tab, but a *contextual* tab—meaning it activates only when you work with a shape.

Follow the steps in this task to add a shape and view the contextual tab.
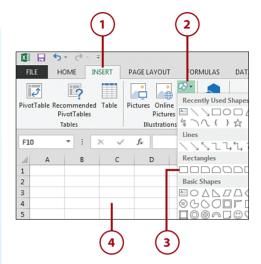
Work with Contextual Tabs
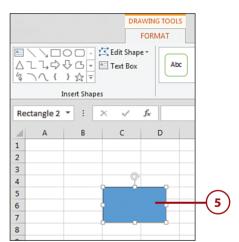
1. Open Excel and start a new blank workbook (as described earlier in the "Explore Worksheets" task). When the new workbook is open, click the Insert tab.

2. Click the Shapes command button.

3. Click the rectangle shape in the menu of shapes.

4. Click anywhere on your spreadsheet to embed your selected shape.

All Shapes and Sizes
Read more about shapes and other graphics in Chapter 7, "Working with Graphics."

5. Click your new shape to select it, and you see the contextual tab, labeled Format, that appears any time the shape is selected.

Customizing the Quick Access Toolbar

The small buttons at the top-left side of the Excel screen are known collectively as the Quick Access Toolbar. By default, the Quick Access Toolbar contains three commands: Save, Undo, and Redo. If you click the drop-down selection arrow next to the Redo button on the Quick Access Toolbar, you see that these 3 commands are only 3 of 11 available as prebuilt or built-in options. Placing a check next to any of the options that you see here automatically adds each option to the Quick Access Toolbar. So, clicking New adds the New command, and clicking Open adds the Open command. You can

actually click each one of these and have all 11 options available to you in the Quick Access Toolbar. In the tasks that follow, you find out how to add commands to the Quick Access Toolbar to customize it to fit your workflow.
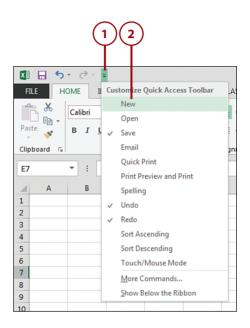
Add a Command to the Quick Access Toolbar
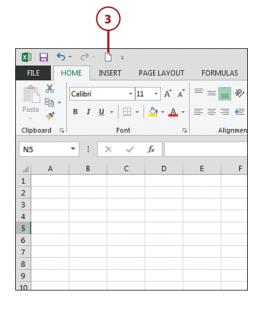
Follow these steps to add to the Quick Access Toolbar one of the commands available from the drop-down menu.

1. Click the arrow to open the Quick Access toolbar drop-down to see all the existing commands you can add.

2. Click the New command to add it to the Quick Access toolbar.

3. You now have a New icon on your Quick Access toolbar.

Moving the Quick Access Toolbar

You can move the Quick Access Toolbar under the Ribbon by right-clicking its drop-down arrow and selecting Show Below the Ribbon. This moves the entire Quick Access Toolbar under the Ribbon—closer to your worksheet.
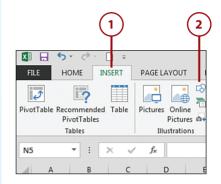
Add Other Commands to the Quick Access Toolbar
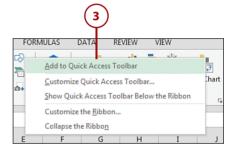
You're not limited to only the commands available on the Quick Access Toolbar drop-down menu—you can add all kinds of commands. For example, if you often work with shapes, you can add the Shapes command to the Quick Access Toolbar.

1. Click the Insert tab.

2. Right-click the Shapes command button.

3. Click the Add to Quick Access Toolbar menu item.

Adding Entire Groups

You can add entire groups to your Quick Access Toolbar. For instance, if you often use the Illustrations group on the Insert tab, you can right-click that entire group and select the Add to Quick Access Toolbar option.
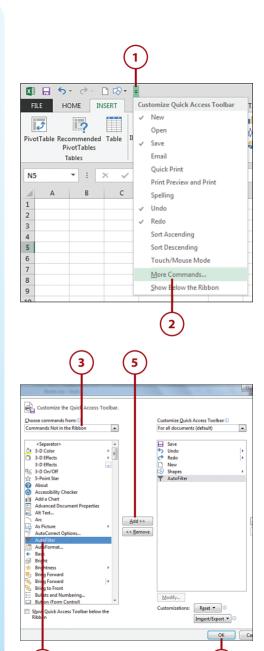
Add Hidden Commands to the Quick Access Toolbar

A closer look at the drop-down menu next to the Quick Access Toolbar reveals an option called More Commands. When you click this item, the Excel options dialog box opens up with the Quick Access Toolbar panel activated. This panel enables you to select from the entire menu of Excel commands, adding them to your Quick Access Toolbar. This comes in handy if you want to add some of the more obscure Excel commands that aren't found on the Ribbon.

1. Click the Quick Access Toolbar drop-down arrow.

2. Select the More Commands option.

3. Choose Commands Not in the Ribbon.

4. Find and select the AutoFilter command.

5. Click the Add button.

6. Click OK to add the command to the Quick Access Toolbar.

Removing Commands
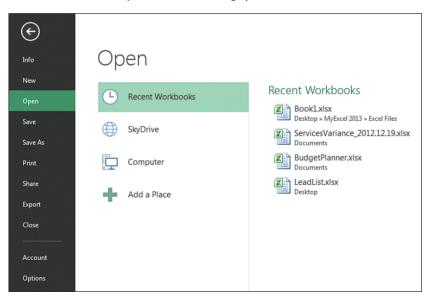
To remove a command from the Quick Access Toolbar, simply right-click it and select the Remove from Quick Access Toolbar option.
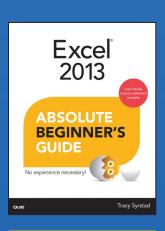
Click the File tab and enter the Backstage view

Use the commands in the Backstage view to print, save, and manage your Excel files

Working with Workbooks, Sheets, Rows, Columns, and Cells

Excel 2013 made easier than ever: the best practical guide for every Excel beginner and hundreds of millions of upgraders!

- Practical, approachable, and packed with useful tips!

- Covers all the basics: workbook creation, editing, formatting, formulas, functions, charts, sharing, security, and much more

- Helps you get started fast with Excel 2013's best new features, from simpler data import and analysis to easier sharing and workbook presentation

- By Tracy Syrstad, one of the world's most experienced and successful Excel tutorial authors!

ABOUT THE AUTHOR

Tracy Syrstad is the project manager for the MrExcel consulting team and also handles many of the tech support calls. She's been helping people with Microsoft Office issues since 1997 when she discovered free online forums where anyone could ask and answer questions. She discovered she enjoyed teaching others new skills, and when Tracy began working with the MrExcel team, she was able to integrate the fun of teaching with one-on-one online desktop sharing sessions.

- Understand the difference between a workbook and a sheet.
- Use a template.
- Insert a new sheet.
- Move Rows, Columns, and Ranges.

2

WORKING WITH WORKBOOKS, SHEETS, ROWS, COLUMNS, AND CELLS

An Excel file is called a workbook and each workbook can contain one or more sheets. This chapter shows you how to create workbooks from scratch or from templates. It shows you how to add, delete, and move sheets within a workbook or between workbooks.

Even a well-planned-out sheet layout may be missing something, such as a date column. Or you might change your mind in the middle of the design, deciding that you want a table elsewhere on a sheet. Instead of starting over, this chapter shows you how to insert rows or columns and move your tables to a new location on your sheet.

Managing Workbooks

Workbook is another name for an Excel file. It's the workbook you open, save, and close. This section shows you how to create a new workbook, open an existing workbook, and save and close a workbook.

Creating a New Workbook

When you open Excel, it already makes a new workbook available (if the Start screen is disabled) or you can create one by clicking Blank Workbook from the Start screen (if the Start screen in enabled). To create a new workbook if you already have another workbook open, you can click File, New and then click on Blank Workbook, or you can press Ctrl+N.

Opening an Existing Workbook

If the Start screen is turned on, when you open Excel you will see various template icons on the right side of the window and a list of recent workbooks and the option to Open Other Workbooks on the left side of the window, as shown in Figure 2.1.

If you want to open a workbook that was recently open, select it from the list on the left. Notice that the Recent Workbooks list shows the workbook name and its location. See the section "Using the Recent Workbooks list" for information on keeping a workbook in the list.

TIP   By default, Excel shows the most recent 25 workbooks opened on the Recent list. You can change this by clicking File, Options, Advanced; then under Display, change the value for Show This Number of Recent Workbooks.
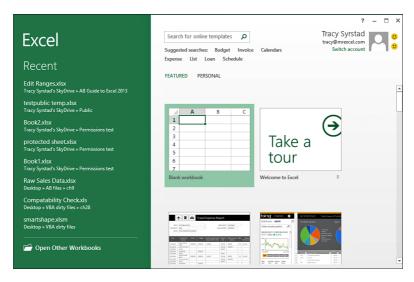
FIGURE 2.1

The left side of the Start screen lists recently opened workbooks. The right side shows various templates you can start working with.

If the workbook you are looking for is not listed, select Open Other Workbooks. A new window opens in which you can choose the place where the workbook is located, as shown in Figure 2.2. The default places are as follows:

- **Recent Workbooks**—The most recently opened workbooks and any workbooks you have pinned to the list

- **Your SkyDrive**—A location on the Internet on Microsoft's SkyDrive server

- **Computer**—Your local computer and any networks attached

- **Add a Place**—A new location you can add, such as a network drive, so it is readily available

The area to the right of the Places list will update to reflect your places selection. For the SkyDrive and Computer options, recent folders will be listed, or you can browse to another folder.

When you select a folder or the Browse option, the Open dialog box, shown in Figure 2.3, appears, allowing you to search for a specific Excel file and offering multiple ways to open it.

In the upper-right corner of the dialog box is a Search field in which you can enter terms to search the current folder and its subfolders. Type in a keyword and press Enter.
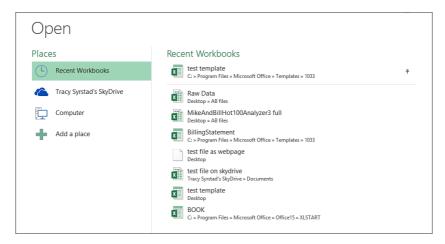
FIGURE 2.2

Excel makes it easy to share your workbook online with access to your SkyDrive and Microsoft's online server, or you can save it to your local drive or network.
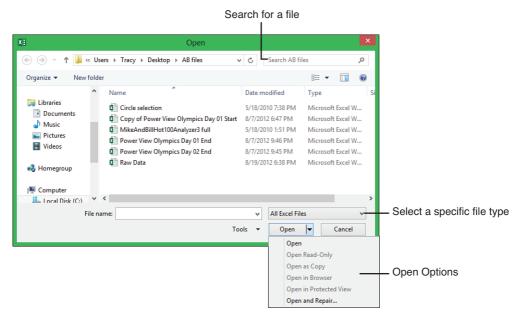
FIGURE 2.3

The Open dialog box offers multiple options to search for a specific file.

Excel can open a variety of files, both Excel files and non-Excel. To change the file type, click All Excel Files and a drop-down of other file types appears.

Once you've found the file you want to open, you can choose how Excel should open it. By default, Excel opens the file normally, but you can also choose to open it as read-only, as a copy, and more:

- **Open**—Opens the workbook normally.

- **Open Read-Only**—Opens a read-only copy of the workbook. You must save the workbook with a different filename if you want to save any changes.

- **Open as Copy**—Opens as a new copy of the workbook. Excel appends "Copy" to the beginning of the filename.

- **Open in Browser**—Opens a workbook saved as a web page in your browser.

- **Open in Protected View**—Opens in Protected view, allowing you to only view the workbook. To make changes, you have to click the Enable Editing button that appears in the alert bar at the top of the sheet.

- **Open and Repair**—Performs checks while opening the workbook and tries to repair any issues it finds. Or, Excel tries to extract the formulas and values.

Using the Recent Workbooks List

When you go to File, Open a list of workbooks recently opened is shown. The Recent Workbooks list makes it easy to open files you're currently working on. A pushpin appears when you place your cursor over a workbook listed under Recent Workbooks. The pushpin points to the left when not in use, but place your cursor over the pushpin and click and the pushpin will face downward and the workbook will be moved to the top section of the list, as shown in Figure 2.2. A workbook with a pin facing down is stuck—that is, the workbook will not be removed from the list of recent workbooks until you unstick it by clicking the pushpin again.

TIP A workbook you've deleted or moved may appear in the list for a while. If you've moved the file, you might get confused by seeing the old location in the list. To remove the workbook from the list, right-click on it and select Remove from List.

Saving a Workbook

When you click the File menu, you see two options for saving your workbook: Save and Save As. Selecting Save saves the workbook with the current name, overwriting the existing file. If the workbook is new and doesn't have a filename yet, clicking Save opens the Save As dialog box. The Save As dialog box allows you to save a workbook with a new name in a new location. You would select

Save to create a copy of the workbook for changes, but want to keep the original workbook intact.

When you select Save As, the area to the right of the File menu changes to show various places where the workbook can be saved to. The default places are as follows:

• **Your SkyDrive**—A location on the Internet on Microsoft's SkyDrive server

• **Computer**—Your local computer and any networks attached

• **Add a Place**—A new location you can add, such as a network drive, so it is readily available

The area to the right of the **Places** list updates to reflect recent folders associated with your places selection. You can select from the list or you can browse to another folder, opening the Save As dialog box, as shown in Figure 2.4.
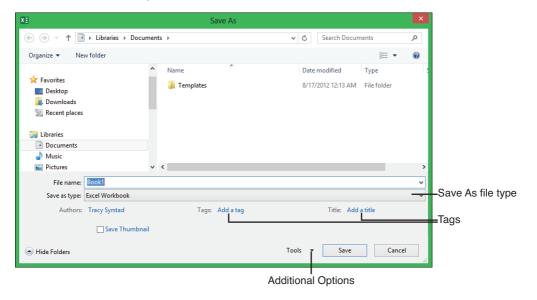
FIGURE 2.4

The Save As dialog box offers multiple options when saving your workbook.

By default, Excel saves your workbook with an XLSX extension. But if you need to share with Excel 97–2003 users or have macros in your workbook, you must choose another Save As file type. Commonly used extensions are as follows:

• **Excel 97–2003 Workbook**—XLS: Excel 97–2003 workbook with or without macros

- **Excel Workbook**—XLSX: Excel 2007–2013 workbook without macros

- **Excel Macro-Enabled Workbook**—XLSM: Excel 2007–2013 workbook with macros

- **PDF**—PDF: A Portable Document Format that can be used to distribute your report but prevent users from making changes

> **TIP** If you find yourself often changing the Save As file type to something else, such as Excel 97–2003 Workbook, you can change the default setting by clicking File, Options, Save and changing the Save Files in This Format value. You will still have the ability to select another file type when saving.

Closing a Workbook

When you're done working with a workbook, you need to close it. To close a workbook, click File, Close or click the X in the upper-right corner. Excel prompts you to save changes if changes have recently been made to the workbook, including recalculating formulas, and the workbook hasn't been saved. When the final workbook is closed, Excel closes.

Using Templates to Quickly Create New Workbooks

Templates are a great way for keeping data in a uniform design. You could simply design your workbook and reuse it as needed, but if you accidentally save data before you have renamed the file, your blank workbook is no longer clean. When using a template, there's no risk of saving data in the template because you are working with a copy of the original workbook, not the workbook itself.

Using Microsoft's Online Templates

Microsoft offers a variety of templates, such as budgets, invoices, and calendars, which can help you get a start on a project. You can search for templates to fit your needs either from the Start screen or by clicking File, New and selecting Featured. A selection of templates appears in the window or you can enter your own keywords in the Search field in the upper left. A search returns matches, which you can filter from a list on the right.

TIP   When you place your cursor over a template preview, a pushpin appears in its lower-right corner, giving you the option to pin it so that it always appears in the list. If you found your template by searching for it, this is a way to make it easily accessible the next time you need it.

When you click on a template, a preview window opens up, as shown in Figure 2.5. The preview shows a larger example of the template, a brief explanation of its purpose, and how the template has been rated by other users. If you click Create, the template is downloaded (for first-time use) and opened.
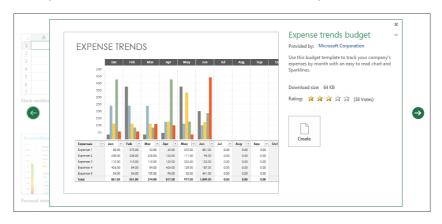
FIGURE 2.5

Use the left and right arrows in the template preview window to scroll through other available templates before opening one.

Opening a Locally Saved Template to Enter Data

There are multiple ways to open a template to enter data and you can use the method that is most convenient to you. To open a template to create a new workbook, use one of the following methods:

- Double-click the file from its saved location.

- From its saved location, right-click the file and select New.

- Click File, New and select the template from the list of online templates.

- Click File, New, Personal and select from the list of templates saved locally. See the "Saving a Template" section for details on how to configure the location. The Personal option will not appear unless the location is set up.

Opening a Locally Saved Template to Make Changes to the Design

If you need to make changes to the design of a template, the previous methods for opening a copy of the template won't work. If you do open a template using one of those methods and make changes, you must use the Save As option and save it as a new template. Instead, to open the original template to make changes, use one of the following methods:

- Right-click the file and select Open.
- Click File, Open, select the file and click Open.

Saving a Template

NOTE   You do not have to save templates to the configured templates location if you plan to use the double-click or right-click method to open a copy of the file.

Before you can use your own template or one provided to you and not downloaded from the Office Store, you have to configure where your templates will be stored. This will be a folder where you will place all your templates you want to access through the File, New, Personal dialog box. To set up the location, click File, Options, Save and enter the path in the Default Personal Templates Location field. Once you have a location configured, Personal will appear to the right of Featured. Featured shows Microsoft's online templates; Personal shows templates in the locally configured location.

CAUTION   The location must exist before you enter it in the field. To make sure you get the path correct, navigate to it in File Explorer, copy the full path from the address bar, then paste it in the Excel field.

Customizing All Future Workbooks

Excel's default template is fine to start with, but what if there are changes you always make, like changing the margins or the font on the sheet? There isn't a template file you can overwrite with your preferred adjustments. But if you design your template and save it with a specific name in a specific location, Excel will use it instead of its own template when you create a new workbook. To change the default template used by Excel, follow these steps:

1.   Open a blank workbook and make the desired changes.

2.   Click File, Save As, Computer and browse to the XLSTART folder. The location is based on where Office was installed. The default location for 64-bit Office is: C:\Program Files\Microsoft Office\Office15\XLSTART.

> **CAUTION**   You must have the proper permissions to save to this location.

3.   Enter **BOOK** in the File Name text box.

4.   Select Excel Template as the Save As Type.

5.   Click Save. Now, the next time you create a new workbook from scratch, your custom template will be used.

> **NOTE**   If you select the Blank Workbook icon on the Start screen, it uses Excel's default template. To have Excel open directly to a blank workbook based on your template, you'll have to turn the Start screen off. To do this, go to File, Options, General and deselect Show the Start Screen When This Application Starts.
>
> If you prefer to keep the Start screen on, then save a second copy of your template in the template folder configured in the options. The template will then be available when you select Personal from the Start screen.

Working with Sheets and Tabs

A sheet, also known as a spreadsheet or worksheet, is where you enter all your data in Excel. A workbook can have multiple sheets—the number limited only by the power of the computer opening the workbook.

Each sheet has a tab, visible above the status bar, as shown in Figure 2.6. The sheet with data that you're looking at is considered the active sheet. To select another sheet, click on its tab.

FIGURE 2.6

The font of the sheet tab you are working on, Sheet1, is bold compared with the other sheet tabs.

Inserting a New Sheet

The quickest way to add a new sheet to a workbook is to click on the circle with a plus sign that appears to the right of the rightmost sheet tab. This inserts a new sheet to the right of the active sheet.

To insert a new sheet to the left of a specific sheet, select the sheet (refer to the "Activating Another Sheet" section) and click Home, Insert, Insert Sheet.

A third method is to right-click on the tab to the right of where you want the new sheet to appear. From the menu, select Insert, select Worksheet, and click OK. The new sheet appears to the left of the tab you right-clicked on, as shown in Figure 2.7.
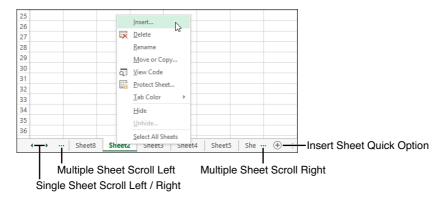
FIGURE 2.7

You can control where a new sheet is inserted by right-clicking on a tab and selecting Insert. In the figure, the new sheet, Sheet8, appears to the left of the selected Sheet2.

Activating Another Sheet

To activate another sheet, click on its tab along the bottom of the Excel window. You can also navigate sheet to sheet using the keyboard. Ctrl+Page Up selects the sheet to the left; Ctrl+Page Down selects the sheet to the right.

If you have more sheets than can fit in the tab area, three small dots will appear on the left and right side of the area, to show there are more sheets available to view, as shown in Figure 2.7. Clicking on the three dots quickly scrolls through multiple sheets. If instead, you need to scroll one sheet at a time, click the left and right arrows located to the left of the sheet tab area.

If you right-click on the left and right arrows instead, a list of all sheets in the workbook appears. You can then select a sheet, click OK, and you'll go straight to that sheet.

Selecting Multiple Sheets

You can select multiple sheets at one time. This doesn't activate them all at once. It groups them together so that an action on one, such as changing the tab color, affects all of them.

To select multiple sheets, select the first one, then while holding down the Ctrl key, select the tabs for the others. As you select each sheet tab, the sheet name will become bold.

To ungroup the sheets, you can either select any sheet other than the active sheet or right-click on any tab and select Ungroup Sheets.

Deleting a Sheet

If you no longer need a sheet, you can delete it, getting it out of your way and reducing the size of your workbook. To delete a sheet, right-click on the sheet's tab and select Delete. You can also delete a sheet by making it your active sheet and clicking Home, Cells, Delete and selecting Delete Sheet from the drop-down. If there is data on the sheet, a prompt appears to verify this is the action you want to take.

CAUTION Deleting a sheet cannot be undone.

Moving or Copying Sheets Within the Same Workbook

You might want to reorganize sheets within the current workbook to group similar sheets together. Or you might need to copy a sheet to run tests on the data but don't want to change the original. To move a sheet within a workbook, use the Move or Copy Sheet option found under Home, Cells, Format. The dialog box can also be accessed by right-clicking on a sheet tab and selecting Move or Copy.

In the Move or Copy dialog box, shown in Figure 2.8, make sure the To Book field is the current workbook. If you're making a copy of the sheet, select the Create a Copy box. Then, in the Before Sheet dialog box, select the sheet you want the active sheet to be placed in front of (to the left of). You can also reorganize the sheets in a workbook by clicking, holding the mouse button down, and dragging the sheet's tab to a new location. If you want to copy the sheet, hold down the Ctrl key as you move the sheet tab.
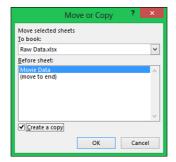
FIGURE 2.8

The Move or Copy dialog box allows you to copy or move sheets within a workbook or between two different workbooks.

Moving or Copying Sheets Between Workbooks

Many database programs will export data to an Excel workbook or compatible file but might not let you choose an existing workbook to export to. If you have a report workbook designed in which you need the new data, you can copy or move the exported data from its workbook to your report workbook. To move a sheet to a new workbook, use the Move or Copy Sheet option found under Home, Cells, Format. The dialog box can also be accessed by right-clicking on a sheet tab and selecting Move or Copy.

To move or copy a sheet to another workbook, the second workbook must already be open. Make sure the sheet you want to move or copy is the active sheet. Then, using the Move or Copy dialog box, shown in Figure 2.8, select from the To Book field the workbook to move or copy the sheet to. If you're making a copy of the sheet, select the Create a Copy box. In the Before Sheet dialog box, select the sheet you want the active sheet to be placed in front of (to the left of).

CAUTION   When you move or copy a sheet from one workbook to another, any formulas on that sheet linked to another sheet in the original workbook will remain linked to the sheet in the original workbook unless you also include the linked sheet(s) in the move or copy procedure.

Renaming a Sheet

Excel's sheet names, Sheet1, Sheet2, etc., aren't very descriptive. Also, when you copy a sheet in a workbook that already has a sheet by that name, Excel will copy the sheet name, appending a number to the end of it. For example, if you copy

Sheet1 within the same workbook, Excel renames it to Sheet1 (2). To give a sheet a more meaningful name, go to Home, Cells, Format and select Rename Sheet from the drop-down. Excel selects the current sheet name in the sheet's tab and you can then type in the new name. You can also rename a sheet by right-clicking on the sheet's tab and selecting Rename.

Coloring a Sheet Tab

If you have several sheets you want to visually group together, you can color their tabs. For example, if you have some sheets that are for data input and others for reports, you could color your data input sheet tabs a light yellow and the report tabs a light blue.

To change the color of a tab, go to Home, Cells, Format, and select Tab Color from the drop-down. A selection of colors appears from which you can select the desired color. The tab color of the active sheet appears as a gradient but fills the entire tab for the other sheets. You can also change the tab color by right-clicking the tab and selecting Tab Color.

Working with Rows and Columns

It's hard to create the perfect data table the first time around. There's always something you've forgotten, such as the date column. Or maybe you did remember the date column but put it in the wrong place. Thankfully, you don't have to start over. You can insert, delete, and move entire columns and rows with a few clicks of the mouse.

Selecting an Entire Row or Column

When you select an entire row or column on a sheet, you are selecting beyond what you can see on the sheet. You really do select the entire row or column, and whatever you do to the part you can see affects the entire selection.

To select an entire row, place your cursor on the header until it turns into a single black arrow then click on the header for the row. For example, to select row 5, as shown in Figure 2.9, click on the header 5. To select an entire column, click on the header for the column.

3	2008	2000	123	42	21	0	TRUE	1
4	2008	2000	233	42	26	9	TRUE	1
5	2008	2000	258	30	11	0	TRUE	1
6	2008	2000	264	23	4	0	TRUE	1
7	2008	2000	310	21	12	0	TRUE	1

FIGURE 2.9

Click on the headers to quickly select an entire row or column.

If you need to select multiple, contiguous rows (or columns), click on the first header, then as you hold down the mouse button, drag the mouse down (or up) until you have selected all the rows you need. Let go of the mouse button and perform your desired action on the selection.

If you need to select multiple, noncontiguous rows (or columns), click on the first header, then as you hold down the Ctrl key, carefully click on the headers of the other rows. When done selecting, release the Ctrl key and perform your desired action on the selection.

Inserting an Entire Row and Column

You might want to insert a row if your data table is missing titles over the data. Or, you might need to insert a column if the data table is missing a column of information and you don't want to enter the information at the end (right side) of the table. When you insert a row, Excel shifts the existing data down. When you insert a column, Excel shifts the existing data to the right.

TIP If you select multiple rows or columns, Excel will insert that many rows or columns. For example, if you need to insert five new rows in your data, instead of highlighting one row, right-clicking, selecting Insert, and then repeating this four more times, select five rows, right-click, and select Insert and— Voila!—you have five blank rows. Note that the selection doesn't have to be contiguous. For example, if you've selected several noncontiguous rows, Excel will insert one row beneath each selected row.

To insert a new row or column within an existing data set, select a cell in the row or column where you want the inserted row or column to go. For example, if you need to insert a new row 10, select a cell in row 10. Go to Home, Cells. From the Insert drop-down, select Insert Sheet Rows or Insert Sheet Columns. For a shortcut, you can right-click on the row or column header and select Insert, as shown in Figure 2.10.

NOTE If you don't have the entire row or column selected and you right-click over the selection, Excel displays the Insert dialog box, prompting you to specify whether you want to shift the selected cells right or down, or if you want to insert an entire row or column. See the "Inserting Cells and Ranges" section for more information.

FIGURE 2.10

Right-clicking a header brings up the options to let you quickly insert or delete a row or column.

Deleting an Entire Row and Column

When you import data from another source, you might need to clean it up by deleting rows or columns of data you don't need. When deleting a row, Excel shifts the data below the row up. For example, after deleting row 5, what was in row 6 now appears in row 5. If deleting a column, Excel shifts the data that was to the right of the deleted column to the left.

To delete a row or column, select a cell in the row or column to delete and go to Home, Cells and from the Delete drop-down, select Delete Sheet Rows or Delete Sheet Columns. For a shortcut, you can right-click on the row or column header and select Delete, as shown in Figure 2.10.

> **TIP**   Your selection of rows or columns does not have to be contiguous. Excel will delete what you have selected.

Moving Entire Rows and Columns

You have to be careful when moving rows and columns. Depending on the method you use, you could end up overwriting other cells. To move rows or columns when you aren't worried about overwriting existing data, follow these steps:

1. Select the row or column you want to move.

2. The selection is surrounded by a thick line. Place your cursor over the line until it turns into a four-headed arrow, as shown in Figure 2.11.

FIGURE 2.11

When the cursor changes to a four-headed arrow, you can click and drag the column to a new location.

3. Click on the line and hold the mouse button down.

4. Drag the selection to a new location and let go.

5. If there is any data in the new location, Excel asks if you want to overwrite it. If you say No, the move will be canceled.

6. The original row or column will still be there, but it will be empty.

If you need to keep data in the new location intact, you could insert a row or column in the new location, copy the data to be moved, paste it in the inserted rows or columns and then go back to the original location, select it again, and delete it. Or, you could get it done more quickly by selecting and cutting the row or column, right-clicking on the new location, and selecting Insert Cut Cells. Excel moves the data in the new location, inserting the cut data and also deleting the data in the old location.

To perform these actions using the ribbon, after selecting and cutting the row or column, select a cell in the new location and go to Home, Cells, Insert and select Insert Cut Cells from the drop-down.

> **TIP**   With one small addition, you can use both of the previous methods to copy a row to a new location. If using the drag method, hold down the Ctrl key as you drag the row to the other location. If using the menu method, copy the row instead of cutting it and select Insert Copied Cells from the menu.

Working with Cells

Just like you can insert, move, and delete entire rows and columns, you can also insert, move, and delete cells. This section will show you how to quickly jump to a cell you can't currently see on your sheet, select noncontiguous ranges, and insert, move, and delete cells.

Selecting a Cell Using the Name Box

Chapter 1, "Understanding the Microsoft Excel Interface," explained how each cell has an address made up of the column and row it resides in. The basics of navigation were touched upon—you can either click on a cell or use the keyboard to select a cell.

You can also jump to a cell by typing the cell address in the Name box, located above the sheet in the left corner. This is a great way of traveling quickly to a specific location on a large sheet.

To use this method, you must enter the cell address properly—that is, the column then the row. For example, if you want to jump to cell AA33, you must type in AA33; you can't type in 33AA. After typing the address, press Enter and the currently selected cell will move to cell AA33.

> **TIP**   If you have R1C1 notation turned on—you'll have numbers instead of letters in the column header—A1 notation won't work. You will have to enter the cell address using the R1C1 system of addressing cells. See the section "Using R1C1 Notation to Reference Cells" in chapter 5, "Using Formulas" for more information on R1C1 notation.

Selecting Noncontiguous Cells and Ranges

In Chapter 1, you were introduced to the basics of cell and range selection in the section "Selecting a Range of Cells." In this section, you learn how to select multiple, noncontiguous ranges, as shown in Figure 2.12.

FIGURE 2.12

You can select noncontiguous cells by holding down the Ctrl key as you select each range.

You don't always want to select cells right next to each other. For example, you might want to select only the cells with negative values and make them bold. You could select each cell one at a time and apply the bold format, or you could hold down the Ctrl key and select each cell. Once all the cells are selected, then you let go of the Ctrl key and select your desired format (or other action).

CAUTION   Once a cell has been selected using the Ctrl key method, you cannot deselect it. If you make a mistake in your cell selection, you have to start over.

Inserting Cells and Ranges

When you insert or delete rows and columns, Excel shifts all the data on the sheet. If you don't want all the data shifted, select a specific range, then either right-click over the selection and choose Insert or go to Home, Cells, Insert and select Insert Cells from the drop-down. A prompt will appear, asking you in which direction you want to shift your existing data, as shown in Figure 2.13. If you want to insert rows, choose Shift Cells Down. If you want to insert columns, choose Shift Cells Right.

FIGURE 2.13

You can control how much of a row or column is inserted by selecting a specific range beforehand.

Deleting Cells and Ranges

When you delete a range, you remove the cells from the sheet, shifting other data over to fill in the empty space. If you aren't careful, you can ruin the careful layout of your sheet, for example, moving data from the credit column to the debit column. If what you really wanted was to delete the data in the cells, leaving the cells intact, Excel calls this clearing the contents. For specifics on clearing contents, see the section "Clearing the Contents of a Cell," in Chapter 3, "Getting Data onto a Sheet."

To delete a range from a sheet, select the range and go to Home, Cells, Delete and select Delete Cells from the drop-down. From the Delete dialog box that opens, choose whether you want to Shift Cells Left or Shift Cells Up. You can also open the dialog box by right-clicking on the selected range and choosing Delete from the menu.

Moving Cells and Ranges

You have to be careful when moving a range—you could end up overwriting other cells. If you need to move rows or columns and aren't worried about overwriting existing data, use the following method. If you don't want to overwrite existing data in the new location, you need to use the Insert Cells method explained in the section "Inserting Cells and Ranges." To move rows or columns when you aren't worried about overwriting existing data, follow these steps:

1. Select the range you want to move.

2. The selection is surrounded by a thick line. Place your cursor over the line until it turns into a four-headed arrow, as shown in Figure 2.14.

3. Click on the line and hold the mouse button down.

4. Drag the selection to a new location and let go.

5. If there is any data in the new location, Excel asks if you want to overwrite it. If you select No, the move will be canceled. Select Yes to complete the move and clear the contents of the original range.

FIGURE 2.14
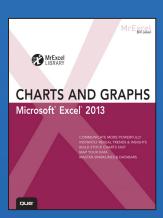
Be careful when dragging your selection to a new area as it will overwrite anything there.

THE ABSOLUTE MINIMUM

You don't always have to start from scratch when creating a workbook—you can use templates offered online at Microsoft or create your own. Nor are you limited to a single sheet. You can add sheets as needed, for example one for your data and one for your report on that data.

You know how to insert new rows and delete columns. With what you've learned in this chapter, you're ready to learn how to enter data on a sheet and put those inserting, deleting, and moving skills to the test.

CHARTS AND GRAPHS

Microsoft® Excel® 2013

COMMUNICATE MORE POWERFULLY
INSTANTLY REVEAL TRENDS & INSIGHTS
BUILD STOCK CHARTS FAST
MAP YOUR DATA
MASTER SPARKLINES & DATABARS

ISBN: 9780789748621

CHAPTER 1:

Introducing Charts in Excel 2013

The first practical, hands-on guide to getting the most out of Excel 2013's jaw-dropping charts, graphs, and other visual tools!

- Everything Excel users need to know to communicate visually - from trend analysis to stock charting, geographical mapping to displaying product pictures

- Shows how to make the most of Excel 2013's new Recommended Charts, Quick Analysis, and chart formatting features

- From basic through leading-edge techniques - including automatic generation of custom charts with VBA

- Part of Bill Jelen's popular MrExcel's Library series!

ABOUT THE AUTHOR

Bill Jelen, Excel MVP and the host of MrExcel.com, has been using spreadsheets since 1985, and he launched the MrExcel.com website in 1998. Bill was a regular guest on Call for Help with Leo Laporte and has produced more than 1,500 episodes of his daily video podcast, Learn Excel from MrExcel. He is the author of 39 books about Microsoft Excel and writes the monthly Excel column for Strategic Finance magazine. His Excel tips appear regularly in CFO Excel Pro Newsletter and CFO Magazine . Before founding MrExcel.com, Bill Jelen spent 12 years in the trenches--working as a financial analyst for finance, marketing, accounting, and operations departments of a $500 million public company. He lives near Akron, Ohio, with his wife, Mary Ellen.

QUE®

Introducing Charts in Excel 2013

1

What's New in Excel 2013 Charts

Several charting improvements are introduced in Excel 2013. My favorite new item is the 153 professionally designed chart styles that you can use to format your charts. These styles enable you to apply modern formatting to the chart in a few clicks. They are useful compared to the old gallery of 48 chart styles introduced in Excel 2007.

Other new items in Excel 2013 charting include the following:
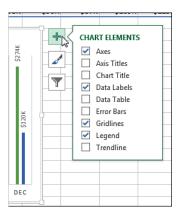
- **Recommended charts**—Your starting point for creating a chart should be to select the data and choose Insert, Recommended Chart. Excel applies heuristics and rules to offer charts that will work for your data. There will always be a few different options to choose from. Excel will never offer the bad chart types like pyramid, cone, or 3D charts. If nothing in the Recommended Charts pane works for you, you can use the All Charts tab in the dialog to access another chart style.

- **Charting unsummarized data sets**—Typically, charts should be created from summaries of your detailed data. However, if you select the entire data set and ask for a recommended chart, Excel looks for columns with repeating values and offers to create a pivot chart that summarizes revenue by product or region, for instance.

- **No legend in a single-series chart**—The new default chart always has a chart title. It has a legend only if you have more than one series. This prevents the redundant chart title and data legend that have appeared in Excel charts over the years.

- **Labels inherit number formatting from the source data**—If you format your source data as currency with 0 decimal places, the data labels and value axis labels have a similar format. This saves you the extra step of figuring out how to change the number format. It is the way it should have worked all along, and hopefully the pivot table team at Microsoft will pick up the cue and make pivot tables work this way in the future.

- **Quickly add elements using the Plus icon**—When a chart is selected, three new helper icons appear to the right of the chart. The Plus icon (+) holds elements that you can quickly add or remove from the chart. This tiny icon effectively replaces the entire Layout ribbon tab from Excel 2007–2010. (See Figure 1.1.)

- **Apply formatting using the Paintbrush icon**—The Paintbrush icon holds thumbnails

Figure 1.1
Three new icons hold old and new chart formatting features.

of new chart styles and new color schemes. You can quickly convert a boring chart to a good-looking chart in a few clicks using this icon.

- **Filter unwanted total rows**—The Funnel icon lets you remove rows or columns from the chart without having to hide them in the worksheet. Suppose that you have 12 monthly data points with four quarterly subtotals interspersed. Typically, you should select the noncontiguous ranges to create the chart. Now you can select the whole data set, create a bad-looking chart, but quickly correct the chart by filtering out the quarterly columns. (See Figure 1.2.)

- **Data label callouts from cell values**—A new data label callout shape debuts in Excel 2013. Even better, the values for data labels can now come from a range of cells in the worksheet. This enables you to create dynamic chart labels using formulas in the worksheet. (See Figure 1.3.)

- **Formatting commands move to a task pane**—The Format task pane replaces the Formatting dialog box. You will like that the Format task pane takes up less space on your screen than the old Format dialog box did. However, the commands are now hidden across a three-tier menu system, which takes some getting used to.

Figure 1.2
Rather than select non-contiguous cells, chart all the data and then filter out the subtotal cells.
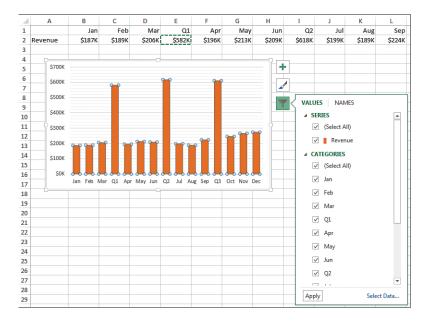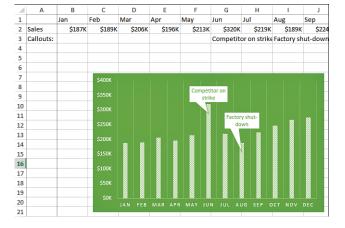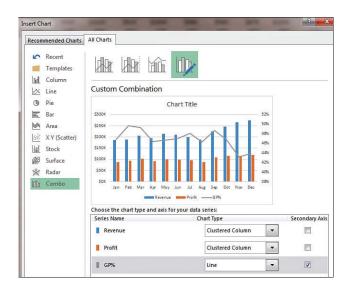
Figure 1.3
Callout text can come from cells on the worksheet.

- **New interface for creating combo charts**—The Excel team finally found a logical way to specify combo charts. Excel 2003 offered a few inflexible combo charts. Excel 2007–2010 did away with the canned examples, hoping people would figure out the confusing steps needed to convert a chart into a combo chart. Few people did. Now the Insert Chart dialog offers a custom combo interface. You can choose exactly which series belongs on the secondary axis, and you can choose the chart type for each series. (See Figure 1.4.)

- **Bad charts types are hidden**—Cone, pyramid, and cylinder Charts have been banished to an underground cavern in Kentucky. They were bad chart types. They never should have been in Excel. The chart team finally removed them from the Insert Chart interface. If you are desperate to create misleading charts, you might be able to find similar settings in the Format task pane.

Figure 1.4
A new interface for creating combo charts.

Choosing Among the Three Ways to Create a Chart
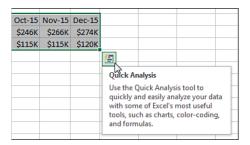
There are three entry points for creating a chart. After using Excel 2013 for a while, I believe the Recommended Charts icon is the fastest way to get a good chart. This section reviews all three methods:

■ Quick Analysis icon

■ Recommended Charts

■ Insert tab of the ribbon

Creating a Chart from the Quick Analysis Icon

When you select a range of data in Excel 2013, the Quick Analysis icon appears below the data. (See Figure 1.5.)
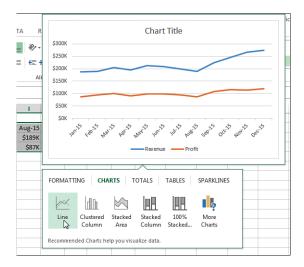
Figure 1.5
Select a range of data and the Quick Analysis icon appears.

Click the Quick Analysis icon and a menu appears with choices for Formatting, Charts, Totals, Tables, and Sparklines. Click the Charts menu, and the Quick Analysis tool offers a

few recommended charts. Hover over any thumbnail to see a live preview of that chart. (See Figure 1.6.)

Figure 1.6
After choosing the Charts category, you can hover, hover, hover to see what each chart would look like.

If none of the thumbnails offer what you want, you can click the More Charts icon at the end, which is equivalent to selecting the Recommended Charts icon.

It requires two clicks to get to the Charts section of the Quick Analysis tool, followed by several hovers to review and reject the various suggested charts, and then a third click to get to the recommended charts. It seems easier to skip this entire process and start with Recommended Charts.

Inserting a Recommended Chart

You will likely prefer the Recommended Charts method for creating charts. The icon appears as a large icon on the Charts group on the Insert tab of the ribbon. (See Figure 1.7.) If you end up being a fan of this icon, you can get here by pressing Alt followed by N then R. (Note that you have to press and release Alt before pressing N.) Or you can right-click the Recommended Charts icon and add it to the Quick Access toolbar.

Select the range of data and choose Recommended Charts. Excel displays the Insert Chart dialog and starts on the Recommended Charts tab of the dialog. Here you can see all the thumbnails without having to hover over each one. (See Figure 1.8.) This is better than the Quick Analysis lens.

To blindly accept the first suggestion, use the keystrokes Alt, N, R, Enter. This takes slightly longer than the old Alt+F1 shortcut key for embedding a default chart, but the results are better.

Figure 1.7
The Recommended
Charts icon is new in
Excel 2013.
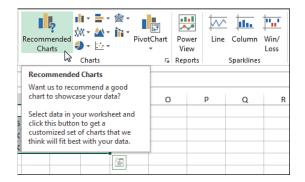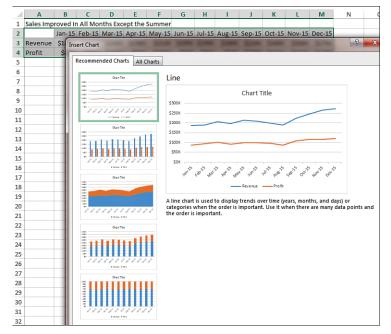
Figure 1.8
You can see thumbnails of
each recommended chart.

If none of the recommended charts suit you, click the All Charts tab in this dialog to get to all 73 built-in chart types. (See Figure 1.9.) This tab is easier to use than the various icons in the Charts group of the ribbon. The left navigation panel offers all the chart categories, even the Stock and Surface types, which are hidden on the ribbon. The seven thumbnails across the top offer seven types of charts, and the two large thumbnails enable you to choose whether your series should be in rows or columns.

Figure 1.9
You can see thumbnails of
each recommended chart.

Decoding the Charting Thumbnails

Most of the charting categories offer thumbnail icons across the top of the dialog. These icons usually represent four basic types of charts:

- **Clustered**—This is the top-left chart in Figure 1.10. Each series gets a column, bar, or point that starts from the category axis of the chart. This type of chart makes it easy to compare how each series did over time or across the axis.

- **Stacked**—This is the top-right chart in Figure 1.10. It shows how the three series add up to reach a total value. It is great for communicating the total sales of the three regions, but horrible for seeing how series 2 or series 3 changed over time.

- **100% Stacked**—The actual numbers in the cells are converted to percentages. Each column/bar/line adds up to 100%. This method is horrible for showing total sales and horrible for showing how each series changed over time. This is the bottom-left chart in Figure 1.10.

- **3D Column**—Whereas the preceding three items are available in 3D versions, the column and area charts offer a fourth 3D choice: the 3D column, where the series are shown in a front-to-back arrangement. This works when the series in the front are smaller than the series in the back. Frequently, though, the items in the back series are hidden by items in the front series.

The row of icons in the middle of Figure 1.10 shows how Excel attempts to represent these four chart types for column charts using thumbnails.

As you move to other chart categories, Excel attempts to convey the Clustered, Stacked, and 100% Stacked concepts in a thumbnail representative of that chart category. Figure 1.11 shows the thumbnails used for Column, Line, Bar, and Area charts.

Figure 1.10
These thumbnails repre-
sent Clustered, Stacked,
100% Stacked, and 3-D
Column.
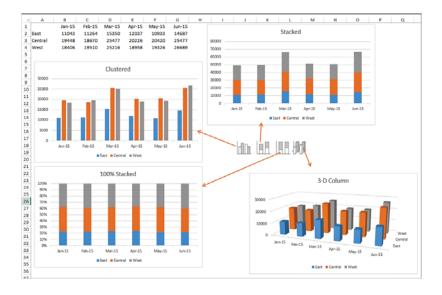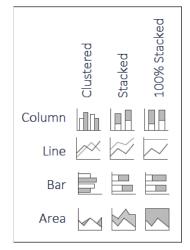
Figure 1.11
The Clustered, Stacked,
and 100% Stacked are
available in Column, Line,
Bar, and Area charts.
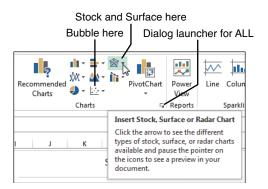
Creating a Chart Using the Other Icons on the Insert Tab

Old habits are hard to break. If you don't embrace the new Recommended Charts icon, you
can head directly to one of the eight other charting drop-down menus on the Insert tab of
the ribbon. (See Figure 1.12.)

There really should be more icons than the eight that are shown. Bubble charts have been
moved to the Scatter Chart icon. Stock and Surface charts are in the Radar chart icon.
Cone, Pyramid, and Cylinder charts have been removed from the Column Chart icon and
moved to the Formatting task pane. There is no icon for Templates or Recent.

Figure 1.12
The Excel team tried to squeeze 11 charting categories into eight icons.

To reach the categories of Templates or Recent, you need to open the All Charts tab of the Insert Chart dialog. Even the tiny dialog launcher shown in Figure 1.12 takes you to the Recommended Charts tab of the Insert Chart dialog. It is easier to click the big Recommended Charts icon and then click All Charts.

Creating a Chart Using Alt+F1

There is a fourth way to create a chart. Select the data and press Alt+F1. Excel creates a default chart in the middle of the current window. Normally, the default chart is a clustered column chart.

If you frequently create a different type of chart, you can set that chart as the default chart that is tied to Alt+F1. In the All Charts dialog, click the small icons across the top and choose Set as Default Chart. (See Figure 1.13.)

Figure 1.13
The command to set the default chart moved in Excel 2013.

TIP
A custom chart template can be defined as the default chart. This enables you to define custom colors, effects, and settings as the default. See "Creating a Chart Template," near the end of Chapter 2, "Customizing Charts," for more information.

Changing the Chart Title

If you create a chart with a single series, Excel 2013 uses the series label as the chart title. In all other cases, the newly created chart has a generic title of "Chart Title."

In the past, you would typically get no title in a chart.

It was easy to ignore a title that did not exist.
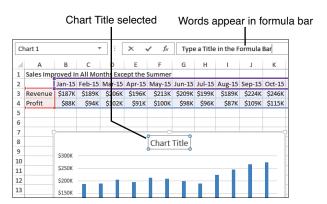
Now that all your charts are going to have a chart title that reads "Chart Title," you are going to have to get very good at changing the title.

Editing a Title in the Formula Bar

If you click the chart title and start typing, the words that you type appear in the Formula Bar. This is disconcerting if you don't know where to look for the words to appear. You might think that nothing is happening or that you accidentally kicked the keyboard cable from the back of the computer.

In Figure 1.14, the Chart Title is selected. Words have been typed, but they are appearing in the Formula Bar instead of in the title. When you finally press Enter to complete the title, the title appears in the chart.
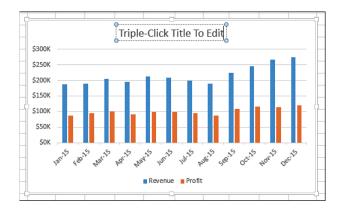
Figure 1.14
Single-click the title, start typing, and the words appear in the Formula Bar until you press Enter.

Editing a Title in the Chart

Triple-click the title in the chart and start typing words (see Figure 1.15). The new words that you type replace "Chart Title."
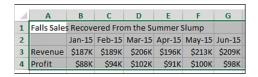
Figure 1.15
Triple-click the title and you can type the words right in the title box.

Why Can't Excel Pick Up the Title from the Worksheet?

You might wonder why Excel cannot pick up the title from the worksheet. Before you create a chart, what if you include the title in A1 or A2? Surely, Excel will figure out that you want the long text in that cell to be used as the title. (See Figure 1.16.)
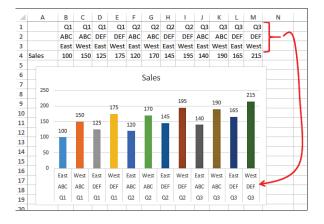
Figure 1.16
You would think titles typed in A1 or A2 would be applied to the title by Excel. They are not.

Unfortunately, Excel does not understand that you want that cell to be the title. I've tried it many ways and it never works. Scott Ruble from the Excel team finally confirmed that they will never pick up the title from a data set with two or more series from the worksheet.

Scott did point out that multiple rows of data labels above the chart do get picked up as a hierarchical set of labels, as shown in Figure 1.17. This is a neat trick that has been in Excel for many versions, but it doesn't help with the title.

Figure 1.17
Multiple rows of cell titles get picked up as category axis labels.
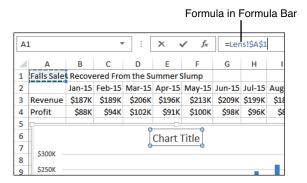
Assigning a Title from a Worksheet Cell

There is a great trick for applying a title from a worksheet cell. Because Excel won't pick up the chart title, you should try this method:

1. Leave space for a chart title somewhere on the worksheet.

2. Create the chart. If there is more than one series, it starts with the default title of Chart Title.

3. Study the chart and find something interesting in the chart. I like to tell my chart audience something about the data. In the current example, type the title **Fall Sales Recovered from the Summer Slump** in cell A1.

4. Click the Chart Title once to select the title.

5. Click in the Formula Bar. Type an equal sign. Click the cell that contains the title. Note that Excel builds the formula as if you are pointing to another worksheet. Because the embedded chart is a foreign object on the worksheet, a formula that points back to the parent worksheet is essentially pointing to another worksheet. (See Figure 1.18.)
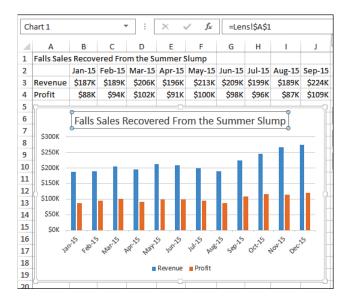
Figure 1.18
Select the title. Type an equal sign and then click a cell that contains the title.

Formula in Formula Bar

6. Press Enter. The words from the cell appear in the chart title. (See Figure 1.19.) If you later change the text in cell A1, the chart title updates.

Using this trick enables you to create a title that would change over time. For example, you could have a formula in A1 with =“Sales as of “&TEXT(TODAY(),”MMM D, YYYY”).

Figure 1.19
It takes four clicks to set up a dynamic title.

Handling Special Situations

You will frequently run into situations where the source data is not set up perfectly for charting. The techniques in this section assist you in creating the perfect chart from your existing data.

Charting Noncontiguous Data

Sometimes you have some extra rows or columns in your data that do not belong in the chart. The new strategy for Excel 2013 is to first create the chart with the extra cells and then use the Funnel icon to filter those rows or columns out of the chart.

Consider Figure 1.20. The data shows sales for four quarters. Someone added subtotal columns to total Half 1 and Half 2 as well as a grand total column. It is simple enough to not select columns G and H when you create the chart. But the subtotal for Half 1 in column D causes a spike in the middle of your chart.

The easy solution is to select the chart. Three icons appear to the right of the chart. Click the Funnel icon. A list of all your Series and Categories appears. Now uncheck H1 from the list. (See Figure 1.21.)

Figure 1.20
A subtotal for Half 1 is
right in the middle of
your contiguous data.
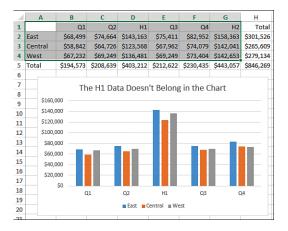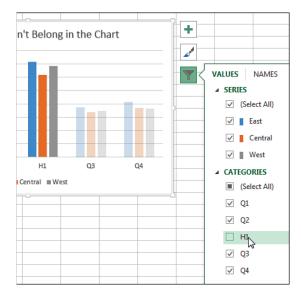
Figure 1.21
Remove H1 from the
chart using the Funnel
icon.

The result is a chart that omits the extra subtotal column. (See Figure 1.22.)

The old strategy is to use the Ctrl key to select noncontiguous data. Suppose you want to create a pie chart showing Profit by Region from the data in Figure 1.23. Follow these steps to select data for creating a chart:

1. Click in cell A1 and drag to cell A10 to select the range of category labels.

2. While holding down the Ctrl key, click in cell F1 and drag down to cell F10 to add F1:F10 to the selection.

3. If you have additional series to plot, repeat step 2 for each additional series.

Figure 1.22
After removing H1, only the quarters remain.

	A	B	C	D	E	F	G	H
1		Q1	Q2	H1	Q3	Q4	H2	Total
2	East	$68,499	$74,664	$143,163	$75,411	$82,952	$158,363	$301,526
3	Central	$58,842	$64,726	$123,568	$67,962	$74,079	$142,041	$265,609
4	West	$67,232	$69,249	$136,481	$69,249	$73,404	$142,653	$279,134
5	Total	$194,573	$208,639	$403,212	$212,622	$230,435	$443,057	$846,269

Use the Funnel Icon to Remove a Row or Column
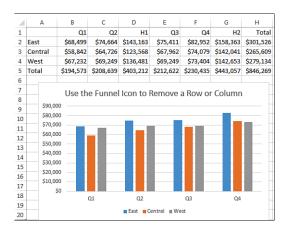
Figure 1.23
Selecting noncontiguous data requires a bit of dexterity because you attempt to drag while holding down the Ctrl key.
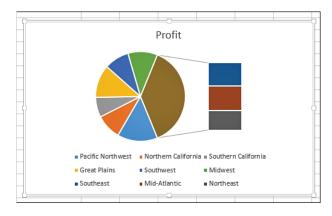
	Region	Units	Avg Price	Sales	GP%	Profit
1	Region	Units	Avg Price	Sales	GP%	Profit
2	Pacific Northwest	8564	87.44	748,836	45.7%	342,218
3	Northern California	5291	81.35	430,423	49.8%	214,351
4	Southern California	4977	75.54	375,963	45.0%	169,183
5	Great Plains	7010	81.26	569,633	49.1%	279,690
6	Southwest	5073	91.76	465,498	45.3%	210,871
7	Midwest	6756	76.06	513,861	48.8%	250,764
8	Southeast	7040	90.01	633,670	48.0%	304,162
9	Mid-Atlantic	8177	78.94	645,492	50.0%	322,746
10	Northeast	5644	92.58	522,522	49.7%	259,693

CAUTION

Excel remembers the order in which you selected the data. Although choosing cell A1 and then individually Ctrl+clicking cells F10, F9, B7, B3, F2, and so on could lead to a selection that looks the same as Figure 1.23, it would not create an acceptable chart. You must select the category labels first and then Ctrl+click and drag to select the first series.

The resulting chart is based on column A and column F. (See Figure 1.24.)

Figure 1.24
By selecting noncontiguous ranges, you can avoid using the Funnel icon.

Profit

- Pacific Northwest
- Northern California
- Southern California
- Great Plains
- Southwest
- Midwest
- Southeast
- Mid-Atlantic
- Northeast

Charting Nonsummarized Data

The new Recommended Charts icon can detect when your data is not summarized and will offer to create a summary chart from the data.
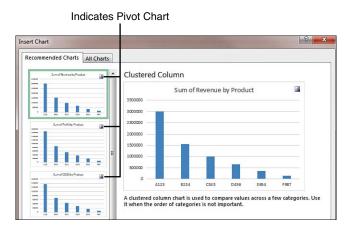
Consider Figure 1.25. This data set contains 500+ rows of detailed data. It looks like a sales database. Creating a chart with 500+ detailed data points would not be a useful chart.

Figure 1.25
This data needs to be summarized before you create the chart.

It is surprising that Excel 2013 handles the summarization for you. Select one cell in the data. Go to Insert, Recommended Chart. The Insert Chart dialog is offering Sum of Revenue by Product as the first chart. Excel is analyzing your data, looking for text columns that have a few values that repeat. In this data set, the six products and four regions become the focus of the recommended charts. (See Figure 1.26.)

Figure 1.26
Excel recommends a summary by product.

Indicates Pivot Chart

When you insert the chart, a new worksheet is inserted to the left of your worksheet. A summary table appears in A3, and a chart is built from that summary table. (See Figure 1.27.)

Figure 1.27
Excel summarizes the
data on a new worksheet
and produces a chart.

> **NOTE** The chart in Figure 1.27 is a special type of chart called a PivotChart. You can change the fields in the chart using the PivotChart Fields pane. Read more about pivot charts in Chapter 8, "Creating Pivot Charts and Power View Dashboards."

CAUTION

The algorithm for choosing which fields to summarize is a bit annoying. In this case, Excel chose Region and Product. Excel avoids Customer with 26 unique values, and it avoids Year with 3 unique numeric values. I've tried various tricks to get Year in the chart, to no avail. Rather than going to great lengths to trick Excel into choosing a column, Chapter 8 shows you how to remove Product in the pivot chart and replace it with Year.

Charting Differing Orders of Magnitude Using a Custom Combo Chart
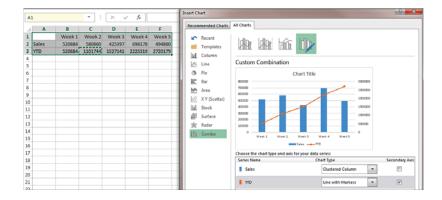
Combo charts let you combine series that have different orders of magnitude. In Figure 1.28, the columns are showing weekly sales. The line is showing accumulated sales. By the end of Week 5, the line is going to be 500% of the size of an average week. As the quarter progresses, the weekly columns will get smaller and smaller compared to the line chart.

The solution is a combo chart with the line plotted on a secondary axis.

Excel 2013 introduces an easy-to-use dialog for creating combo charts. Follow these steps:

1. Select your data.
2. There is a Combo Chart icon on the Insert tab. The icon shows three columns and a line chart above the column chart. Open the Combo Chart drop-down and choose Create a Custom Combo Chart.
3. Excel automatically chooses a Clustered Column chart for Sales and a Line for YTD. Open the Chart Type drop-down for YTD and choose Line with Markers.
4. Check the box for Secondary Axis for the YTD chart.
5. Click OK to create the chart.

Figure 1.28
The new interface for creating combo charts removes the complexity from creating these charts.
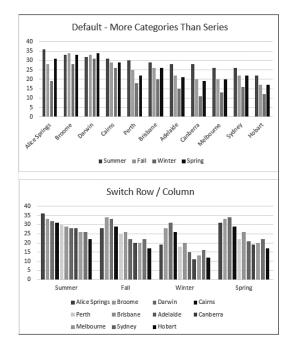
Reversing the Series and Categories in a Chart

Every data set can be plotted with the series coming from the rows or from the columns. When there are more columns than rows, Excel assumes each row is a series. When there are more rows than columns, Excel assumes each column is a series. When the data set is square, each row becomes a series. Although these choices often work, you might need to reverse them.

Figure 1.29 compares two ways to plot data with four seasons and 11 cities. By default, the seasons become the series and the cities become the categories, as in the top chart. You might want to show the data with the seasons as the cities, as in the bottom chart.

Figure 1.29
Both charts are valid ways to plot the data.
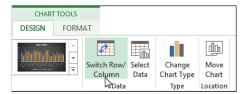
If you are creating the chart using the Insert Chart dialog, Excel 2013 shows you color thumbnails of both choices. The first thumbnail follows the rules for choosing the rows or columns as the series. The second chart has the data reversed. (See Figure 1.30.)

Figure 1.30
The Insert Chart dialog offers two choices, with rows as series or columns as series.

If you used Alt+F1 or any of the small icons in the Charts group to create the chart, you are going to get the default series. A large Switch Rows/Columns icon appears on the Chart Tools Design tab of the ribbon. Click this to switch the series to categories. (See Figure 1.31.)

Figure 1.31
Click this icon to switch the series and categories.
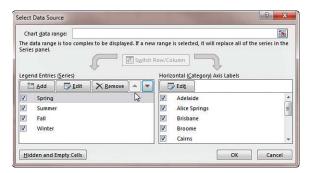
Changing the Data Sequence by Using Select Data

The Select Data icon on the Design tab enables you to change the rows and columns of your data set. The Select Data icon also enables you to resequence the order of the series. When you click this icon, the Select Data Source dialog appears.

As shown in Figure 1.32, buttons in the Legend Entries side of the Select Data Source dialog enable you to add new series, edit a series, remove a series, or change the sequence of a series. A single Edit button on the right side of the dialog enables you to edit the range used for category labels. (See Figure 1.32.)

Figure 1.32
You can use the Select Data Source dialog box for more control over data series.

Using the Charting Tools

When you select a chart, a wide variety of charting tools are available to help you format the chart. You've seen a few of these tools already, such as the Funnel icon. All the charting tools are discussed in depth in Chapter 2. This section introduces the tools at your disposal.

Introducing the Three Helper Icons

When you select a chart, three icons appear to the right of a chart:

- **Plus icon** (aka Chart Elements)—This icon lets you quickly toggle various chart elements. Each element has a check box to turn the element on or off. If you hover to the right of an item, an arrowhead appears. Click the arrowhead to open a fly-out menu with more choices. (See Figure 1.33.)

- **Paintbrush icon** (aka Chart Styles)—This icon is possibly the best new icon in all of Excel 2013. Microsoft had 153 new chart styles professionally designed for Excel 2013. Choices in the Paintbrush icon let you apply a style in a few clicks. Your charts will look modern and new. A second menu under the icon lets you change colors. (See Figure 1.34.)

- **Funnel icon** (aka Chart Filters)—This icon lets you hide or show individual rows or columns in the data. You've used this icon previously in Figure 1.21.

Figure 1.33
The Plus icon lets you add elements to your chart.
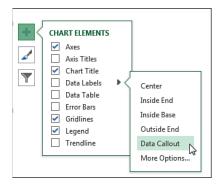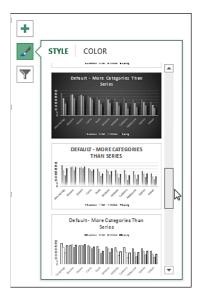
Figure 1.34
The Paintbrush icon quickly applies a style to a chart or changes the colors.

Introducing the Format Task Pane

The old Format dialog box has been replaced by the Format task pane in Excel 2013. If you double-click any element in the chart, you can edit dozens of settings for that element using the choices in the Format task pane.

When the task pane is visible, you can switch to formatting a new element by clicking that element. The choices and menus in the task pane change in response to the current selection.
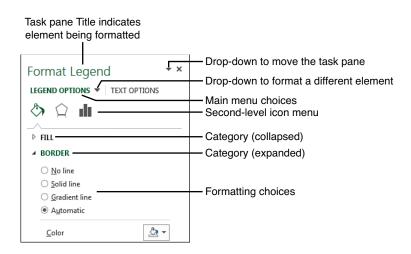
> **NOTE** In fact, if you click away from the chart, the task pane remains. All the choices are grayed out until you again click a chart, shape, picture, or other object.

The task pane is smaller than the old Format dialog box. This led the Excel team to use a new three-level menu system in the task pane. It becomes quite the guessing game to figure out where you can find a particular setting in the task pane. Refer to Figure 1.35 to locate all three levels of menus:

■ The words Format Legend tell you what you are formatting. If you click another element in the chart, these words would change. The drop-down arrow next to the title lets you move the task pane, but it does not let you switch to another element.

■ Underneath the large Format Legend are two menus choices, one for Legend Options and one for Text Options. The currently selected choice is green. The drop-down arrow next to the first menu choice lets you change the element that you are

formatting. This arrow really seems like it should be appearing next to the title of the task pane.

■ Underneath the two main menu choices, a row of icons appears. Each of these icons is a choice in the second-level menu for the task pane. In Figure 1.35, the icons are a Paint Bucket, Pentagon, and a Chart. The row of icons changes if you choose a different element and when you change to a different level 1 menu item. See Table 1.1 for an attempt to explain the icons you see here.

■ Beneath the icons, you see one or more expandable categories in all caps. Click the arrow next to the category to display all the choices for that category. In Figure 1.35, you see Format Legend, Legend Options, Paint Bucket, Border expanded. There is another category for Fill that could be expanded.

Figure 1.35
Formatting choices are hidden among a four-level menu in the task pane.

So far, I've managed to get eight different icons to appear in the level-2 menu in the Format task pane. I would not be completely surprised if there are more. Table 1.1 shows each icon and where it is generally used.

Table 1.1   Format Task Pane Icons

Icon	Looks Like	Official Name	Uses
	Paint Bucket	Fill & Line	Fill and border color
	Pentagon	Effects	Shadow, glow, soft edges, 3D format
	Resize	Size & Properties	Location
	Chart	Variously: Table, Axis, Legend, Series, Trendline, Data Label, Error Bar	Good stuff. Always start here.

Icon	Looks Like	Official Name	Uses
A	Text Fill	Text Fill & Outline	Text color
	Text Effects	Text Effects	Shadow, reflection, glow, soft edges, 3D format, 3D rotation
	Text Box	Textbox	Alignment
	Picture	Picture	Corrections, color, crop

> **NOTE**
> Even though my friend and tech editor Bob Umlas will protest when this book does not use the correct name for something, I am going to use the names in the Looks Like column of Table 1.1 when referring to the icons in the task pane. The official names simply are not descriptive enough.

> **TIP**
> When you start looking for a setting, see if you have the Chart icon. If you do, start there. The best settings are always under the Chart icon.

Table 1.2 identifies which icons will appear as you are formatting various chart elements.

Table 1.2 Icons Available per Chart Element

Element Being Formatted	Paint Bucket	Pentagon	Resize	Chart	Text Fill	Text Effects	Text Box	Picture
Chart Area	•	•	•		•	•	•	
Plot Area	•	•			•	•	•	
Series	•	•		•				
Data Point	•	•		•				
Data Labels	•	•	•	•				
Chart Title	•	•	•		•	•	•	
Axis (any)	•	•	•	•	•	•	•	
Axis Title	•	•	•		•	•	•	
Legend	•	•		•	•	•	•	
Gridlines (any)	•	•						
Data Table	•	•		•	•	•	•	
Trendline	•	•		•				

1

Element Being Formatted	Paint Bucket	Pentagon	Resize	Chart	Text Fill	Text Effects	Text Box	Picture
Series Lines	•	•						
High-Low Lines	•	•						
Error Bar	•	•		•				
Drop Lines	•	•						
Down Bars	•	•						
Up Bars	•	•						
Shape	•	•	•		•	•	•	•
Picture/ Clipart	•	•	•					•
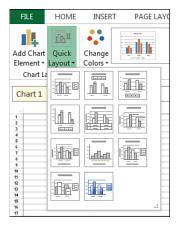
Using Commands on the Design Tab

In most cases, the tools available on the Design tab in the ribbon are redundant with commands previously discussed.

The Add Chart Element drop-down is similar to the Plus icon to the right of the chart. The version in the ribbon adds choices for High-Low Lines, Error Bars, and Up/Down Bars.

The Quick Layout drop-down (see Figure 1.36) was a new idea in Excel 2007, and it never caught on. It offers 4 to 12 ways to format the chart. The charts were not professionally designed. These layouts will rarely be exactly what you need. There are 780 quadrillion possible ways to configure a chart, and this layout gallery offers anywhere from 4 to 12 possible layouts. The odds that your desired layout is represented exactly are slim. For the most part, the choices here are not relevant.

Figure 1.36
These old Quick Layout choices never caught on.

The Change Colors drop-down is identical to the Color menu choice in the Paintbrush icon to the right of the chart.

The Chart Styles gallery is a slightly larger view of the Styles menu in the Paintbrush icon.

The two icons in the Data group let you change the data in the chart.

Change Chart Type is a good place if you decide that you want to change a column chart to an area chart. The same command is in the right-click menu when you click the chart.

Move Chart lets you move a chart to a dedicated chart sheet, in case you want to feel like it is 1985 all over again.

Micromanaging Formatting Using the Format Tab

Everything on the Format tab of the ribbon is redundant with the Formatting task pane. It might be easier to find Shape Fill to change the fill color or Shape Outline to change a border color using the Format tab in the ribbon.
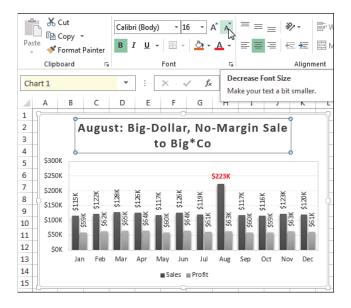
All of the Shape Effects add items like glow, bevel, and so on. Professor Edward Tufte characterizes all this as "chart junk"—effects that add no useful information to the chart.

Using Commands on the Home Tab

If you need to change the font size, color, bold, italic, or typeface, it is easy to select an element in the chart and then use the Font group in the Home tab to format the text.

In Figure 1.37, the title spills to a second line. Click the title once so there is a solid box around the title. You can use the Decrease Font Size icon to make the title size smaller.

Figure 1.37
Use the icons on the Home tab to change the text in the selected chart element.

1

Changing the Theme on the Page Layout tab
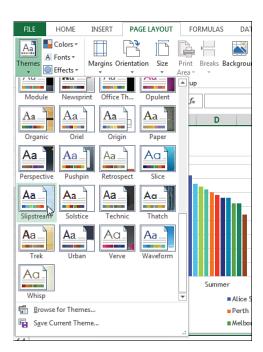
The built-in color choices for a chart include variations of six colors. In Excel 2013, the default colors are Blue, Orange, Gray, Gold, Dark Blue, and Green.

You are not stuck with those colors.

Those six colors are part of a theme. Although everyone starts with the same theme, there are 48 themes to choose from on the Page Layout tab of the ribbon.

On the Page Layout tab, click the Themes drop-down to choose from the 48 built-in themes. If you choose a new theme from the Themes drop-down, Excel applies a new color and set of effects to all the charts in the current workbook. If you want to change only the colors or effects, use the Colors or Effects drop-down in the Themes group. Figure 1.38 shows the themes drop-down.

Figure 1.38
Change colors using the Themes drop-down.

Moving Charts

After creating a chart, you might need to rearrange the data or move the chart to a new location. The topics in this section assist you with these tasks.

Moving a Chart Within the Current Worksheet

In the following case study, you see that Excel had an annoying habit of locating new charts near the bottom of your data set. With a large data set, you might need to move the chart to the proper location thousands of rows away.

There are several ways to move a chart within the current worksheet. Although it is easy to move the chart a few rows in the active window, you should switch to using cut and paste if you have to move the chart hundreds of rows.

Dragging to a New Location in the Visible Window Is Easy

You can drag a chart anywhere in the visible window. Click the border of the chart, avoiding the eight resize handles, and drag. If you try to drag the chart outside of the visible window, the mouse pointer changes to a red "no" symbol. If you release the mouse while it is the no symbol, the chart boomerangs back to its original position.

Dragging Outside of the Visible Window Is Frustrating

If you need to move a chart outside the visible window, drag so that your cell pointer is within one-half row of the edge of the window. Excel slowly starts to scroll in the appropriate direction. As with other Windows programs, you can speed up the scroll by rapidly moving your mouse left and right. However, it is difficult to keep the mouse within the one-half row tolerance while moving left and right.

Adjust the Zoom So New Location Is Within the Visible Window

You are likely to find the scrolling action to be so slow that you will not want to use the previous technique to move the chart to a new location. Another option is to use the Zoom slider to show the worksheet at a 10% zoom. This has the effect of putting about 375 rows in the visible window. If you need to move a chart anywhere from 50 to 375 rows, setting the zoom to 10% and then dragging the chart within the visible window is a fast way to go.

Cut and Paste to Move Thousands of Rows

If your chart is created at the bottom of a 50,000-row data set and you need to move it to the top of the data set, cut and paste might be the fastest way to go. Follow these steps to move a chart within the current worksheet using cut and paste:

1. Select a chart.
2. Press Ctrl+X to cut the chart from the worksheet.
3. Press Ctrl+Home to move to cell A1.
4. Select the cell where you want the top-left corner of the chart to appear.
5. Press Ctrl+V to paste the chart in the new location.

1

CASE STUDY: LOCATING A CHART AT THE TOP OF YOUR DATA SET

The best way to avoid having to move the chart thousands of rows is to have the top of your data set visible before you create the chart.

Suppose that you have 3,000 data points to show on a line chart. If you start in cell A1 and use Ctrl+Shift+down arrow and Ctrl+Shift+right arrow to select the data, you will have row 3000 visible in the window. If you create a chart, the chart appears near row 3000 instead of at the top of the data set.

By changing your workflow slightly, you can make sure the top of the data set is visible before creating the chart.

If you've already selected the data set using Ctrl+Shift+Arrow, use the Ctrl+. (that is, Ctrl+period) shortcut key to move the focus to a new corner of the data. After pressing the shortcut key between one and three times, you will see the top of the data set. Use Alt+F1 or Insert, Recommended Chart to insert the chart near the top of the workbook.

Alternatively, start using Ctrl+* to select the data set. If you start in A2 and press Ctrl+*, Excel selects down to row 3000, but the window stays focused on the top of the worksheet.
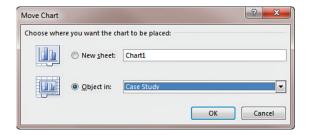
Moving a Chart to a Different Worksheet

In Excel 2013, charts always start out as objects embedded in a worksheet. However, you might want to display a chart on its own full-page chart sheet. When this occurs, there are two options for moving a chart:

- Choose the Move Chart icon at the right edge of the Design tab.
- Right-click any whitespace near the border of the chart and choose Move Chart.

Either way, the Move Chart dialog appears, offering the options New Sheet and Object In. The Object In drop-down lists all the worksheets in the current workbook. The New Sheet option enables you to specify a name for a new sheet (see Figure 1.39).

Figure 1.39
Move a chart to a different worksheet using the Move Chart dialog.

Next Steps

Chapter 2 describes how to use the Layout tab of the ribbon to toggle on or off individual elements of the chart. It also describes how you can micromanage individual elements using the Format ribbon.

CHAPTER 2:

Creating a Basic Pivot Table

The practical, realistic, easy-to-understand guide to Excel's most powerful and underutilized features: PivotTables and PivotCharts

- Fully updated for Excel 2013, including the powerful new Recommended Pivot Table feature

- Teaches PivotTables, PivotCharts, and PowerPivot through real-world scenarios that make intuitive sense to users at all levels

- Includes proven PivotTable recipes for instant on-the-job results

- Helps Excel users avoid common pitfalls and mistakes

- Part of the popular MrExcel's Library series, edited by Excel legend Bill Jelen

ABOUT THE AUTHOR

Bill Jelen, Excel MVP and the host of MrExcel.com, has been using spreadsheets since 1985, and he launched the MrExcel.com website in 1998. Bill was a regular guest on Call for Help with Leo Laporte and has produced more than 1,500 episodes of his daily video podcast, Learn Excel from MrExcel. He is the author of 39 books about Microsoft Excel and writes the monthly Excel column for Strategic Finance magazine. His Excel tips appear regularly in CFO Excel Pro Newsletter and CFO Magazine . Before founding MrExcel.com, Bill Jelen spent 12 years in the trenches--working as a financial analyst for finance, marketing, accounting, and operations departments of a $500 million public company.

Mike Alexander is a Microsoft Certified Application Developer (MCAD) and author of several books on advanced business analysis with Microsoft Access and Excel. He has more than 15 years of experience consulting and developing Office solutions. Mike has been named a Microsoft MVP for his ongoing contributions to the Excel community. In his spare time, he runs a free tutorial site, www.datapigtechnologies.com, where he shares basic Access and Excel tips to the Office community.

MrExcel
LIBRARY
Bill Jelen
Michael Alexander

MrExcel
LIBRARY

PIVOT TABLE
DATA CRUNCHING

Microsoft® Excel® 2013

SLICE & DICE
MASHUPS WITH POWERPIVOT
POWER VIEW DASHBOARDS
INTERACTIVE SLICERS
AUTOMATE WITH VBA
POWER TECHNIQUES

que®

Creating a Basic Pivot Table

2

Preparing Your Data for Pivot Table Reporting

When you have a family portrait taken, the photographer takes time to make sure that the lighting is right, the poses are natural, and everyone smiles his or her best smile. This preparation ensures that the resulting photo is effective in its purpose.

When you create a pivot table report, you're the photographer, taking a snapshot of your data. Taking time to make sure your data looks its best ensures that your pivot table report is effective in accomplishing the task at hand.

One of the benefits of working in a spreadsheet is that you have the flexibility of laying out your data to suit your needs. Indeed, the layout you choose depends heavily on the task at hand. However, many of the data layouts used for presentations are not appropriate when used as the source data for a pivot table report.

As you read the next section, which goes into preparing your data, keep in mind that pivot tables have only one hard rule pertaining to your data source: Your data source must have column headings, which are labels in the first row of your data describing the information in each column. If this is not the case, your pivot table report cannot be created.

However, just because your pivot table report is created successfully does not mean that it's effective. A host of things can go wrong as a result of bad data preparation—from inaccurate reporting to problems with grouping and sorting.

Let's look at a few of the steps you can take to ensure you end up with a viable pivot table report.

Ensure Your Data Is in a Tabular Layout

A perfect layout for the source data in a pivot table is a tabular layout. In tabular layout, there are no blank rows or columns. Every column has a heading. Every field has a value in every row. Columns do not contain repeating groups of data.

Figure 2.1 shows an example of data structured properly for a pivot table. There are headings for each column. Even though the values in D2:D6 are all the same model, the model number appears in each cell. Month data is organized down the page instead of across the columns.

Figure 2.1
This data is structured properly for use as a pivot table source.

	A	B	C	D	E	F
1	REGION	MARKET	STORE	MODEL	MONTH	REVENUE
2	North	Great Lakes	65061011	4055T	April	$2,354
3	North	Great Lakes	65061011	4055T	February	$3,040
4	North	Great Lakes	65061011	4055T	January	$3,454
5	North	Great Lakes	65061011	4055T	March	$2,675
6	North	Great Lakes	65061011	4055T	May	$2,071
7	North	New England	2105015	2500P	April	$11,851
8	North	New England	2105015	2500P	February	$15,304
9	North	New England	2105015	2500P	January	$17,391
10	North	New England	2105015	2500P	March	$13,468
11	North	New England	2105015	2500P	May	$10,429
12	North	New England	22022012	3002C	April	$256
13	North	New England	22022012	3002C	February	$330
14	North	New England	22022012	3002C	January	$375
15	North	New England	22022012	3002C	March	$200

Tabular layouts are *database centric*, meaning you would most commonly find these types of layouts in databases. These layouts are designed to store and maintain large amounts of data in a well-structured, scalable format.

> **TIP**
>
> You might work for a manager who demands that the column labels be split into two rows. For example, he might want the heading Gross Margin to be split with Gross in row 1 and Margin in row 2. Because pivot tables require a unique heading one row high, your manager's preference can be problematic. To overcome this problem, start typing your heading; for example, type **Gross**. Before leaving the cell, press Alt+Enter and then type **Margin**. The result is a single cell that contains two lines of data.

Avoid Storing Data in Section Headings

Examine the data in Figure 2.2. This spreadsheet shows a report of sales by month and model for the North region of a company. Because the data in rows 2 through 24 pertains to the North region, the author of the worksheet put a single cell with North in C1. This approach is effective for display of the data, but not effective when used as a pivot table data source.

Figure 2.2
Region and model data
are not formatted prop-
erly in this data set.

	A	B	C	D
1			**North**	
2	Model 2500 P	January	33,073	
3		February	35,880	
4		March	90,258	
5		April	13,250	
6		May	100,197	
8	Model 3002 P	January	29,104	
9		February	31,574	
10		March	79,427	
11		April	11,660	
12		May	88,173	
14	Model 4055 T	January	35,880	
15		February	25,612	
16		March	27,785	
17		April	69,896	
18		May	10,261	
20	Model 4500 T	January	33,073	
21		February	25,612	
22		March	27,785	
23		April	69,896	
24		May	10,261	

Also in Figure 2.2, the author was very creative with the model information. The data in
rows 2 through 6 applies to Model 2500P, so the author entered this value once in A2 and
then applied a fancy vertical format combined with Merge Cells to create an interesting
look for the report. Again, although this is a cool format, it is not useful for pivot table
reporting.

Also, the worksheet in Figure 2.2 is missing column headings. You can guess that column A
is Model, column B is Month, and column C is Sales, but for Excel to create a pivot table,
this information must be included in the first row of the data.

Avoid Repeating Groups as Columns

The format shown in Figure 2.3 is common. A time dimension is presented across several
columns. Although it is possible to create a pivot table from this data, this format is not
ideal.

The problem is that the headings spread across the top of the table pull double duty as col-
umn labels and actual data values. In a pivot table, this format would force you to manage
and maintain six fields, each representing a different month.

Figure 2.3
This matrix format is
common but not effec-
tive for pivot tables. The
Month field is spread
across several columns of
the report.

	A	B	C	D	E	F	G	H
1								
2	**North**	**MODEL**	**JANUARY**	**FEBRUARY**	**MARCH**	**APRIL**	**MAY**	**JUNE**
3		4054T	$2,789	$2,454	$2,160	$1,901	$1,673	$1,472
4		4500C	$32,605	$28,692	$25,249	$22,219	$19,553	$17,207
5		3002P	$52,437	$46,145	$40,607	$35,734	$31,446	$27,673
6		2500P	$17,391	$15,304	$13,468	$11,851	$10,429	$9,178
7		4055T	$2,468	$2,172	$1,911	$1,682	$1,480	$1,302
8		3002C	$375	$330	$290	$256	$225	$198

Eliminate Gaps and Blank Cells in Your Data Source

Delete all empty columns within your data source. An empty column in the middle of your data source causes your pivot table to fail on creation because the blank column, in most cases, does not have a column name.

Delete all empty rows within your data source. Empty rows may cause you to inadvertently leave out a large portion of your data range, making your pivot table report incomplete.

Fill in as many blank cells in your data source as possible. Although filling in cells is not required to create a workable pivot table, blank cells in and of themselves are generally errors waiting to happen. So a good practice is to represent missing values with some logical missing value code wherever possible.

> **NOTE** Although this might seem like a step backward for those of you who are trying to create a nicely formatted report, it pays off in the end. When you are able to create a pivot table, there will be plenty of opportunities to apply some pleasant formatting. In Chapter 3, "Customizing a Pivot Table," you discover how to apply styles formatting to your pivot tables.

Apply Appropriate Type Formatting to Your Fields

Formatting your fields appropriately helps you avoid a whole host of possible issues, from inaccurate reporting to problems with grouping and sorting.

Make certain that any fields to be used in calculations are explicitly formatted as a number, currency, or any other format appropriate for use in mathematical functions. Fields containing dates should also be formatted as any one of the available date formats.

Summary of Good Data Source Design

The attributes of an effective tabular design are as follows:

- The first row of your data source is made up of field labels or headings that describe the information in each column.
- Each column in your data source represents a unique category of data.
- Each row in your data source represents individual items in each column.
- None of the column names in your data source double as data items that will be used as filters or query criteria (that is, names of months, dates, years, names of locations, or names of employees).

CASE STUDY: CLEANING UP DATA FOR PIVOT TABLE ANALYSIS

The worksheet shown in Figure 2.4 is a great-looking report. However, it cannot be effectively used as a data source for a pivot table. Can you identify the problems with this data set?

Figure 2.4
Someone spent a lot of time formatting this report to look good, but what problems prevent it from being used as a data source for a pivot table?

	A	B	C	D	E	F	G
1	Region	Market		Jan	Feb	Mar	Apr
2	**Bar Equipment**						
3	Midwest	Chicago		132	106	110	90
4		Kansas City		413	504	2,571	505
5		Omaha		332	543	372	424
6	North	Dakotas		130	136	106	90
7		Great Lakes		488	445	4,140	517
8							
9	**Commercial Appliances**						
10	Midwest	Chicago		780	76	851	76
11		Kansas City		3,352	76	8,442	2,831
12		Omaha		228	17,628	76	304
13	North	Dakotas		0	0	2,608	0
14		Great Lakes		990	76	11,435	76
15							
16	**Concession Equipment**						
17	Midwest	Chicago		808	0	3,912	0
18		Kansas City		824	1,761	11,181	1,616
19		Omaha		0	8,147	2,968	3,118
20	North	Dakotas		0	0	5,463	2,370
21		Great Lakes		751	808	13,814	1,632

- The model information does not have its own column. Product Category information appears in the Region column. To correct this problem, insert a new column for Product Category and include the category name on every row.

- There are blank columns and rows in the data. Column C should be deleted. The blank rows between models (such as rows 8 and 15) also should be deleted.

- Blank cells present the data in an outline format. The person reading this worksheet would probably assume that cells A4:A5 fall into the Midwest region. These blank cells need to be filled in with the values from above.

> **TIP**
> Here's a trick for filling in the blank cells. Select the entire range of data. Then select the Home tab on the ribbon and choose the Find & Select icon from the Editing group. This brings up a menu from which you select Go To Special. In the Go To Special dialog, select Blanks. With all the blank cells selected, start a formula by typing the equal sign (=), press the up arrow on your keyboard, and then press Ctrl+Enter to fill this formula in all blank cells. Remember to copy and paste special values to convert the formulas to values.

- The worksheet presents the data for each month in several columns (one column per month). Columns D through G need to be reformatted as two columns. Place the month name in one column and the units for that month in the next column. This step either requires a fair amount of copying and pasting or a few lines of VBA macro code.

> **TIP** For a great book on learning VBA macro programming, read Que Publishing's *VBA and Macros for Microsoft Excel 2013* by Bill Jelen and Tracy Syrstad.

After you make the four changes described here, the data is ready for use as a pivot table data source. As you can see in Figure 2.5, every column has a heading. There are no blank cells, rows, or columns in the data. The monthly data is now presented down column E instead of across several columns.

Figure 2.5
Although this data will take up six times as many rows, it is perfectly formatted for pivot table analysis.

	A	B	C	D	E
1	Product Category	Region	Market	Month	Units
2	Bar Equipment	Midwest	Chicago	Jan	132
3	Bar Equipment	Midwest	Kansas City	Jan	413
4	Bar Equipment	Midwest	Omaha	Jan	332
5	Bar Equipment	North	Dakotas	Jan	130
6	Bar Equipment	North	Great Lakes	Jan	488
7	Commercial Appliances	Midwest	Chicago	Jan	780
8	Commercial Appliances	Midwest	Kansas City	Jan	3,352
9	Commercial Appliances	Midwest	Omaha	Jan	228
10	Commercial Appliances	North	Dakotas	Jan	0
11	Commercial Appliances	North	Great Lakes	Jan	990
12	Concession Equipment	Midwest	Chicago	Jan	808
13	Concession Equipment	Midwest	Kansas City	Jan	824
14	Concession Equipment	Midwest	Omaha	Jan	0
15	Concession Equipment	North	Dakotas	Jan	0
16	Concession Equipment	North	Great Lakes	Jan	751
17	Bar Equipment	Midwest	Chicago	Feb	106
18	Bar Equipment	Midwest	Kansas City	Feb	504
19	Bar Equipment	Midwest	Omaha	Feb	543

Creating a Basic Pivot Table

Now that you have a good understanding of the importance of a well-structured data source, let's walk through creating a basic pivot table.

> **TIP** The sample data set used throughout this book is available for download at www.MrExcel.com/pivotbookdata2013.html.

To start, click any single cell in your data source. This ensures that the pivot table captures the range of your data source by default. Next, select the Insert tab and find the Tables group. In the Tables group, select PivotTable and then choose PivotTable from the drop-down list. Figure 2.6 demonstrates how to start a pivot table.

Choosing these options activates the Create PivotTable dialog, shown in Figure 2.7.

Figure 2.6
Start a pivot table by
selecting PivotTable from
the Insert tab.

Figure 2.7
The Create PivotTable
dialog.
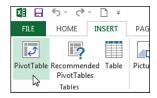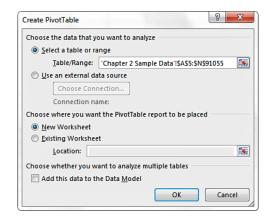
> **TIP**
> You can also press the hotkeys Alt+N+V to start a pivot table.

As you can see in Figure 2.7, the Create PivotTable dialog asks you only two fundamental questions: Where's the data that you want to analyze, and where do you want to put the pivot table?
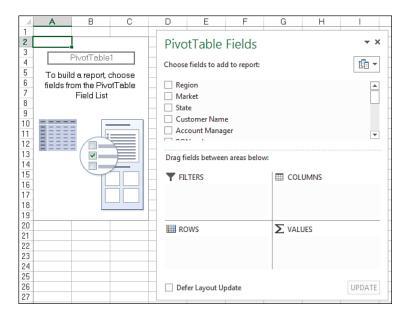
■ **Choose the Data That You Want to Analyze**—In this section, you tell Excel where your data set is. You can specify a data set that is located within your workbook, or you can tell Excel to look for an external data set. As you can see in Figure 2.7, Excel is smart enough to read your data set and fill in the range for you. However, you always should take note of this to ensure you are capturing all your data.

■ **Choose Where You Want the PivotTable Report to Be Placed**—In this section, you tell Excel where you want your pivot table to be placed. This is set to New Worksheet by default, meaning your pivot table will be placed in a new worksheet within the current workbook. You will rarely change this setting because there are relatively few times you'll need your pivot table to be placed in a specific location.

> **NOTE** Note the presence of another option in the Create PivotTable dialog shown in Figure 2.7: the Add This to the Data Model option. You would select this option if you were trying to consolidate multiple data sources into one single pivot table. We cover this option in detail in Chapter 7, "Analyzing Disparate Data Sources with Pivot Tables," and in Chapter 10, "Mashing Up Data with PowerPivot." For now, we cover the steps to create a pivot table from a single source, which means you can ignore this particular option.

After you have answered the two questions in the Create PivotTable dialog, simply click the OK button. At this point, Excel adds a new worksheet that contains an empty pivot table report. Next to that is the PivotTable Field List dialog, illustrated in Figure 2.8. This dialog helps you build your pivot table.

Figure 2.8
You use the PivotTable Field List dialog to build your pivot table.

Finding the PivotTable Field List

The PivotTable Field List dialog is your main work area in Excel 2013. This is the place where you add fields and make changes to your pivot table report. By default, this dialog pops up when you place your cursor anywhere inside your pivot table. However, if you explicitly close this dialog, you override the default and essentially tell the dialog not to activate when you are in the pivot table.

If clicking on the pivot table does not activate the PivotTable Field List dialog, you can manually activate it by right-clicking anywhere inside the pivot table and selecting Show Field List. You can also click anywhere inside the pivot table and then choose the large Field List icon on the Analyze tab under PivotTable Tools in the ribbon.

Adding Fields to the Report

The idea here is to add the fields you need into the pivot table by using the four "drop zones" found in the PivotTable Field List: Filters, Columns, Rows, and Values. These drop zones, which correspond to the four areas of the pivot table, are used to populate your pivot table with data.

> **TIP** Review Chapter 1, "Pivot Table Fundamentals," for a refresher on the four areas of a pivot table.

- **Filters**—Adding a field to the Filters drop zone enables you to filter on its unique data items. In previous versions of Excel, this area was known as the Report Filters area.
- **Columns**—Adding a field into the Columns drop zone displays the unique values from that field across the top of the pivot table.
- **Rows**—Adding a field into the Rows drop zone displays the unique values from that field down the left side of the pivot table.
- **Values**—Adding a field into the Values drop zone includes that field in the values area of your pivot table, allowing you to perform a specified calculation using the values in the field.

Now let's pause a moment and go over some fundamentals of laying out your pivot table report. This is generally the point where most new users get stuck. How do you know which field goes where?
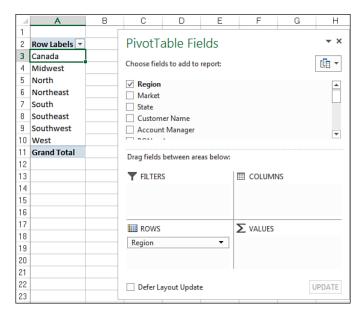
Before you start dropping fields into the various drop zones, ask yourself two questions: "What am I measuring?" and "How do I want to see it?" The answer to the first question tells you which fields in your data source you need to work with, and the answer to the second question tells you where to place the fields.

For your first pivot table report, you want to measure the dollar sales by region. This automatically tells you that you need to work with the Sale Amount field and the Region field. How do you want to see it? You want regions to go down the left side of the report and the sales amount to be calculated next to each region.

To achieve this effect, you need to add the Region field to the Rows drop zone and add the Sale Amount field to the Values drop zone.
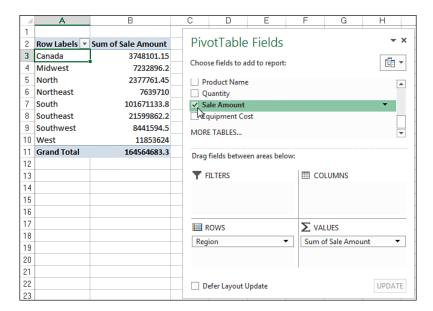
Find the Region field in the field list and place a check in the check box next to it. As you can see in Figure 2.9, not only is the field automatically added to the Rows drop zone, but your pivot table is updated to show the unique region names.

Figure 2.9
Place a check next to
the Region field to auto-
matically add that field to
your pivot table.

Now that you have regions in your pivot table, it's time to add in the dollar sales. To do
that, simply find the Sale Amount field and place a check next to it. As Figure 2.10 illus-
trates, the Sale Amount field is automatically added to the Values drop zone, and your pivot
table report now shows the total dollar sales for each region.

Figure 2.10
Place a check next to the
Sale Amount field to add
data to your pivot table
report.

At this point, you have created your first pivot table report!

How Does Excel Know Where Your Fields Go?

As you've just experienced, the new PivotTable Field List interface enables you to add fields to your pivot table by simply placing a check next to each field name. Excel automatically adds the checked fields to the pivot table. The question is how Excel knows in which drop zone to put the fields you check. The answer is that Excel doesn't really know which drop zone to use; it makes a decision based on data type. Here's how it works: When you place a check next to a field, Excel evaluates the data type for that field. If the data type is numeric, Excel places the field into the Values drop zone; otherwise, Excel places the field into the Rows drop zone. This placement obviously underlines the importance of correctly assigning the data types for your fields.
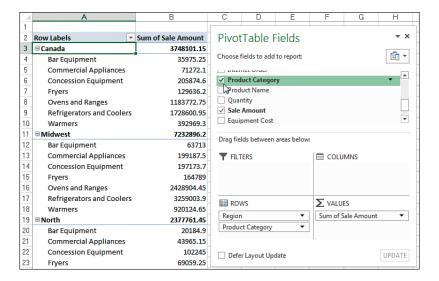
CAUTION

Watch out for blanks in your numeric fields. If you have even one blank cell in a numeric field, Excel reads that cell as a Text field and therefore places it in the Rows drop zone!

Adding Layers to Your Pivot Table

Now you can add another layer of analysis to your report. This time you want to measure the amount of dollar sales each region earned by product category. Because your pivot table already contains the Region and Sales Amount fields, all you have to do is place a check next to the Product Category field. As you can see in Figure 2.11, your pivot table automatically added a layer for Product Category and refreshed the calculations to include subtotals for each region. Because the data is stored efficiently in the pivot cache, this change took less than a second.
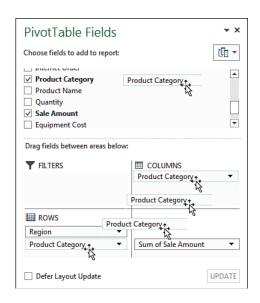
Figure 2.11
Before pivot tables, adding layers to analyses would have required hours of work and complex formulas.

Rearranging Your Pivot Table

Suppose this view doesn't work for your manager. He wants to see Product Categories across the top of the pivot table report. To rearrange them, simply drag the Product Category field from the Rows drop zone to the Columns drop zone, as illustrated in Figure 2.12.

Figure 2.12
Rearranging a pivot table is as simple as dragging fields from one drop zone to another.

You don't have to move your fields into a drop zone to be able to drag them around. You can actually drag fields directly from the field list into the desired drop zone. You can also move a field into a drop zone by using that field's context menu: Click the black triangle next to the field name and then select the desired drop zone.

Instantly, the report is restructured, as shown in Figure 2.13.

Figure 2.13
Your product categories are now column oriented.
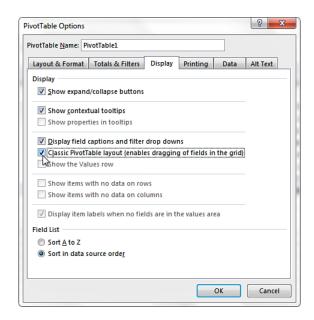
	A	B	C	D	E
1					
2	Sum of Sale Amount	Column Labels			
3	Row Labels	Bar Equipment	Commercial Appliances	Concession Equipment	Fryers
4	Canada	35975.25	71272.1	205874.6	129636.2
5	Midwest	63713	199187.5	197173.7	164789
6	North	20184.9	43965.15	102245	69059.25
7	Northeast	68407.1	285103.35	258557.4	190152.4
8	South	1191742.35	5923096.6	6971902.45	2420745
9	Southeast	283902.15	1525894.3	1569948.6	452579.35
10	Southwest	73528.5	289526.5	355845.45	268554.85
11	West	68684.65	296291.55	422201.2	276443.05
12	Grand Total	1806137.9	8634337.05	10083748.4	3971959.1

Longing for Drag-and-Drop Functionality?

If you're upgrading from Excel 2003, you will notice that you can no longer drag and drop fields directly onto the pivot table layout. This functionality is allowed only within the PivotTable Field List dialog (dragging into drop zones). The good news, however, is that Microsoft has provided the option of working with a classic pivot table layout, which enables the drag-and-drop functionality.

To activate the classic pivot table layout, right-click anywhere inside the pivot table and select Table Options. In the Table Options dialog, select the Display tab and place a check next to Classic PivotTable Layout, as demonstrated in Figure 2.14. Click the OK button to apply the change.

Figure 2.14
Place a check next to Classic PivotTable Layout.
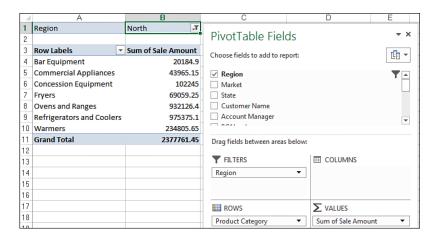
At this point, you can drag and drop fields directly onto your pivot table layout.

Unfortunately, this setting is not global. That is, you have to go through the same steps to apply the classic layout to each new pivot table you create. However, this setting persists when a pivot table is copied.

Creating a Report Filter

Often, you might be asked to produce reports for one particular region, market, or product. Instead of building separate pivot table reports for every possible analysis scenario, you can use the Filter field to create a report filter. For example, you can create a region-filtered report by simply dragging the Region field to the Filters drop zone and the Product Category field to the Rows drop zone. This way, you can analyze one particular region at a time. Figure 2.15 shows the totals for just the North region.

Figure 2.15
With this setup, you not only can see revenues by product clearly, but you also can click the Region drop-down to focus on one region.

Understanding the Recommended PivotTables Feature

With Excel 2013, Microsoft introduces a new feature called Recommended PivotTables. You can find this feature next to the PivotTable icon on the Insert tab (see Figure 2.16).

Figure 2.16
The new Recommended PivotTables icon helps you start a pivot table faster.

This feature is Microsoft's way of getting you up and running with pivot tables by simply creating one for you. The idea is simple. Place your cursor in a tabular range of data and then click the Recommended PivotTables icon. Excel shows you a menu of pivot tables it thinks it can create for you based on the data in your range (see Figure 2.17). When you find one that looks good to you, click it, and click OK to have Excel create it.

Another way to get to a recommended pivot table is to right-click anywhere in your data range and choose the Quick Analysis option. The context menu shown in Figure 2.18 activates, and you can select a recommended pivot table under the Tables section.

The reviews on this new feature are mixed. On one hand, it does provide a quick-and-easy way to start a pivot table, especially for those of us who are not that experienced. On the other hand, Excel's recommendations are rudimentary at best. You will often find that you need to rearrange, add, or manipulate fields in the created pivot table to suit your needs. Although Excel might get it right the first time, it's unlikely you will be able to leave the pivot table as is.

Figure 2.17
Choose from the menu of recommended pivot tables to have Excel automatically create a pivot table for you.
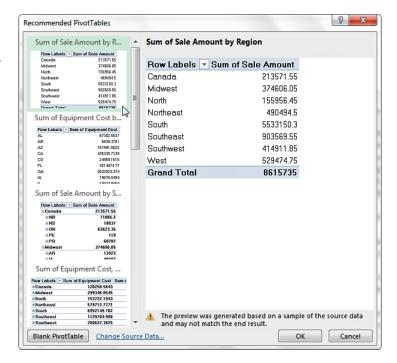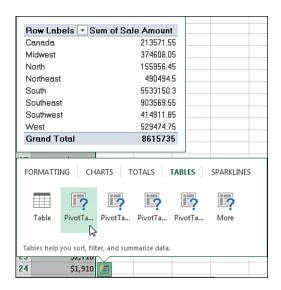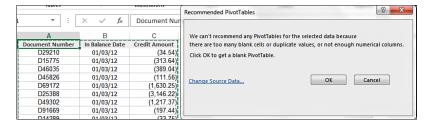
Figure 2.18
You can also choose a recommended pivot table from the Quick Analysis context menu.

There is also a chance that Excel will simply not like the data range you pointed to. For example, the data in Figure 2.19 is a valid range, but Excel doesn't like the repeating dates. So it simply pops up a message to indicate that there are too many duplicates to recommend a pivot table; however, it will gladly create a blank one.

Figure 2.19
When Excel can't recommend a pivot table based on your data, it throws an error and gives you the option of starting with a blank pivot.

All in all, the Recommended PivotTables feature can be a nice shortcut to getting a rudimentary pivot table started, but it's not a replacement for knowing how to create and manipulate pivot tables on your own.

Using Slicers

With Excel 2010, Microsoft introduced a feature called *slicers*. Slicers enable you to filter your pivot table similar to the way Filter fields filter a pivot table. The difference is that slicers offer a user-friendly interface that enables you to easily see the current filter state.

Creating a Standard Slicer

To understand the concept behind slicers, place your cursor anywhere inside your pivot table, and then select the Insert tab on the ribbon. Click the Slicer icon (see Figure 2.20).

Figure 2.20
Inserting a slicer.

The Insert Slicers dialog shown in Figure 2.21 opens. The idea is to select the dimensions you want to filter. In this example, the Region and Market slicers are selected.

After the slicers are created, you can simply click the filter values to filter your pivot table. As you can see in Figure 2.22, clicking Midwest in the Region slicer not only filters your pivot table, but the Market slicer responds by highlighting the markets that belong to the Midwest region.

Figure 2.21
Select the dimensions for which you want to create slicers.
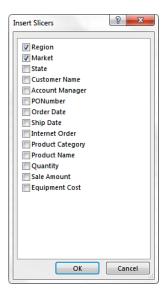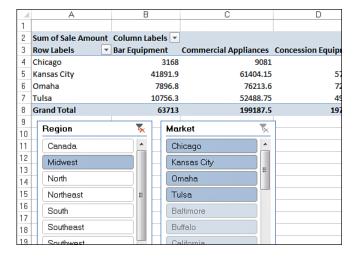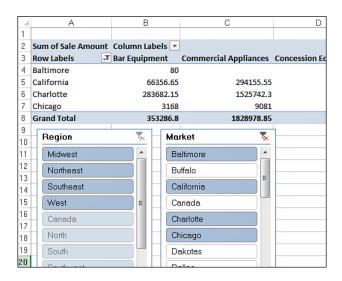
Figure 2.22
Select the dimensions you want to filter using slicers.

You can also select multiple values by holding down the Ctrl key on your keyboard while selecting the needed filters. In Figure 2.23, the Ctrl key was held down while Baltimore, California, Charlotte, and Chicago were selected. This not only highlights the selected markets in the Market slicer, but it also highlights their associated regions in the Region slicer.
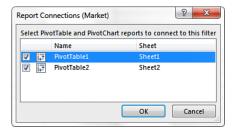
Another advantage you gain with slicers is that you can tie each slicer to more than one pivot table. In other words, any filter you apply to your slicer can be applied to multiple pivot tables.

2

Figure 2.23
The fact that you can visually see the current filter state gives slicers a unique advantage over the Filter field.

To connect your slicer to more than one pivot table, simply right-click the slicer and select Report Connections. The Report Connections dialog shown in Figure 2.24 opens. Place a check next to any pivot table that you want to filter using the current slicer.

Figure 2.24
Choose the pivot tables you want to filter by this slicer.

At this point, any filter applied via the slicer is applied to all the connected pivot tables. Again, slicers have a unique advantage over Filter fields in that they can control the filter state of multiple pivot tables. Filter fields can only control the pivot table in which they live.

> **NOTE** It's important to note that slicers are not part of the pivot table object. They are separate objects that you can use in a variety of ways. For a more detailed look at slicers, their functionality, and how to format them, pick up Que Publishing's *Special Edition Using Microsoft Excel 2013*, by Bill Jelen.

TIP
Notice that in Figure 2.24, the list of pivot tables is a bit ambiguous (PivotTable1, PivotTable2). Excel automatically gives your pivot tables these generic names; which it uses to identify them. You can imagine how difficult it would be to know which pivot table is which when working with more than a handful of pivots. In that light, you might want to consider giving your pivot tables user-friendly names so you can recognize them in dialog boxes such as the one you see in Figure 2.24.

You can easily change the name of a pivot table by placing your cursor anywhere inside the pivot table, selecting the Analyze tab, and entering a friendly name in the PivotTable Name input box found on the far left.

Creating a Timeline Slicer

The Timeline slicer is new in Excel 2013. The Timeline slicer works in the same way as a standard slicer in that it lets you filter a pivot table using a visual selection mechanism instead of the old Filter fields. The difference is the Timeline slicer is designed to work exclusively with date fields, providing an excellent visual method to filter and group the dates in your pivot table.

NOTE
For you to create a Timeline slicer, your pivot table must contain a field where *all* the data is formatted as a date. It's not enough to have a column of data that contains a few dates. All the values in your date field must be formatted as a valid date. If even only one value in your date column is blank or not a valid date, Excel does not create the Timeline slicer.

To create a Timeline slicer, place your cursor anywhere inside your pivot table, select the Insert tab on the ribbon, and then click the Timeline icon (see Figure 2.25).

Figure 2.25
Inserting a Timeline slicer.

The Insert Timelines dialog seen in Figure 2.26 opens, showing you all the available date fields in the chosen pivot table. Here, you select the date fields for which you want to create slicers.

After your Timeline slicer is created, you can filter the data in your pivot table by using this dynamic data-selection mechanism. As you can see in Figure 2.27, clicking the April slicer filters the data in the pivot table to show only April data.

2

Figure 2.26
Select the date fields for which you want slicers created.
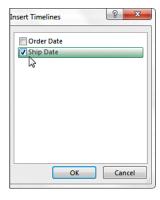
Figure 2.27
Click a date selection to filter your pivot table.
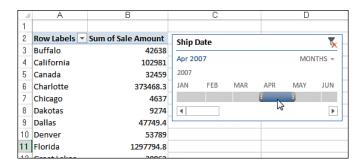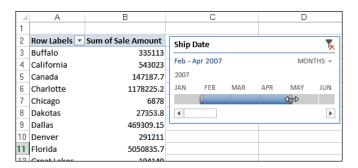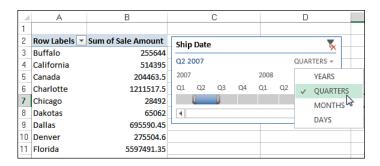
Figure 2.28 demonstrates how you can expand the slicer range with the mouse to include a wider range of dates in your filtered numbers.

Figure 2.28
You can expand the range on the Timeline slicer to include more data in the filtered numbers.

Want to quickly filter your pivot table by quarters? Well, you can easily do it with a Timeline slicer. Click the time period drop-down and select Quarters. As you can see in Figure 2.29, you also have the option of switching to Years or Days if needed.

Figure 2.29
Quickly switch between
Years, Quarters, Months,
and Days.

> **NOTE** Be aware that Timeline slicers are not backward compatible, meaning they are only usable in Excel 2013. If you open a workbook with Timeline slicers in Excel 2010 or earlier versions, the Timeline slicers are disabled.

CASE STUDY: ANALYZING ACTIVITY BY MARKET

Your organization has 18 markets that sell seven types of products. You have been asked to build a report that breaks out each market and highlights the dollar sales by each product. You are starting with an intimidating transaction table that contains more than 91,000 rows of data. To start your report, do the following:

1. Place your cursor inside your data set, select the Insert tab, and click PivotTable.
2. When the Create PivotTable dialog activates, click the OK button. At this point, you should see an empty pivot table with the field list.
3. Find the Market field in the PivotTable Field List and check the box next to it.
4. Find the Sale Amount field in the PivotTable Field List and check the box next to it.
5. To get the product breakouts, find the Product Category field in the PivotTable Field List and drag it into the Columns drop zone.

In five easy steps, you have calculated and designed a report that satisfies the requirements. After a little formatting, your pivot table report should look similar to the one shown in Figure 2.30.

Figure 2.30
This summary can be created in less than a minute.

	A	B	C	D	
1					
2					
3	**Sum of Sale Amount**	**Column Labels** ▾			
4	**Row Labels** ▾	**Bar Equipment**	**Commercial Appliances**	**Concession Equipment**	**Fryer**
5	Baltimore	80			
6	Buffalo	37397.9	237297.85	187711	127
7	California	66356.65	294155.55	416379.2	2610
8	CANADA	35975.25	71272.1	205874.6	129
9	Charlotte	283682.15	1525742.3	1569948.6	4509
10	Chicago	3168	9081	18218	9
11	Dakotas	3386.5	7287.05	30619.2	180
12	Dallas	59580.3	255146	306037	151
13	Denver	34558.85	70068.8	75712.8	1200
14	Florida	1132162.05	5667950.6	6665865.45	2269
15	Great Lakes	16798.4	36678.1	71625.8	51
16	Kansas City	41891.9	61404.15	57300.7	50
17	Knoxville	220	152		

Lest you lose sight of the analytical power you just displayed, keep in mind that your data source has more than 91,000 rows and 14 columns, which is a hefty set of data by Excel standards. Despite the amount of data, you produced a relatively robust analysis in a matter of minutes.

Keeping Up with Changes in Your Data Source

Let's go back to the family portrait analogy. As years go by, your family will change in appearance and might even grow to include some new members. The family portrait that was taken years ago remains static and no longer represents the family today. So another portrait needs to be taken.

As time goes by, your data might change and grow with newly added rows and columns. However, the pivot cache that feeds your pivot table report is disconnected from your data source, so it cannot represent any of the changes you make to your date source until you take another snapshot.

The action of updating your pivot cache by taking another snapshot of your data source is called *refreshing* your data. There are two reasons you might have to refresh your pivot table report:

- Changes have been made to your existing data source.
- Your data source's range has been expanded with the addition of rows or columns.

You handle these two scenarios in different ways.

Changes Have Been Made to Your Existing Data Source

If a few cells in your pivot table's source data have changed due to edits or updates, you can refresh your pivot table report with a few clicks. Simply right-click inside your pivot table report and select Refresh. This selection takes another snapshot of your data set, overwriting your previous pivot cache with the latest data.

> **N O T E**
>
> You can also refresh the data in your pivot table by selecting Analyze from the PivotTable Tools tab in the ribbon and then choosing Refresh.
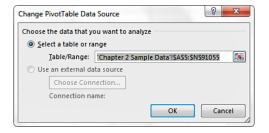
> **T I P**
>
> Clicking anywhere inside your pivot table activates the PivotTable Tools tab just above the main ribbon.

Your Data Source's Range Has Been Expanded with the Addition of Rows or Columns

When changes have been made to your data source that affect its range (for example, you've added rows or columns), you have to update the range being captured by the pivot cache.

To do this, click anywhere inside your pivot table and then select Analyze from the PivotTable Tools tab in the ribbon. From here, select Change Data Source. This selection triggers the dialog shown in Figure 2.31.

Figure 2.31
The Change PivotTable Data Source dialog enables you to redefine the source data for your pivot table.

All you have to do here is update the range to include new rows and columns. After you have specified the appropriate range, click the OK button.

Sharing the Pivot Cache

Many times, you need to analyze the same data set in multiple ways. In most cases, this process requires you to create separate pivot tables from the same data source. Keep in mind that every time you create a pivot table, you are storing a snapshot of your entire data set in a pivot cache. Every pivot cache that is created increases your memory usage and file size. For this reason, you should consider sharing your pivot cache. In other words, in those situations when you need to create multiple pivot tables from the same data source, you can use the same pivot cache to feed multiple pivot tables. By using the same pivot cache for multiple pivot tables, you gain a certain level of efficiency when it comes to memory usage and files size.

In legacy versions of Excel, when you created a pivot table using a data set that was already being used in another pivot table, Excel actually gave you the option to use the same pivot cache. However, Excel 2013 does not give you such an option.

Instead, each time you create a new pivot table in Excel 2013, Excel automatically shares the pivot cache. Most of the time, this is beneficial; you can link as many pivot tables as you want to the same pivot cache with a negligible increase in memory and file size. On the flip side, when you group one pivot table by month and year, all of the pivot tables are grouped in a similar fashion. If you want one pivot table by month and another pivot table by week, you have to force a separate pivot cache.

You can force Excel to create a separate pivot cache by taking the following steps.

1. Press Alt+D+P on your keyboard to launch the old Pivot Table Wizard.
2. Click the Next button to get past the first screen of the wizard.
3. On the second screen, select the range for your pivot table and click the Next button.
4. Excel displays a wordy message saying that you can use less memory if you click Yes. Instead, click No.
5. On the next screen, click the Finish button.

At this point, you have a blank pivot table that pulls from its own pivot cache.

> **TIP**
> If you already have an existing pivot table, you can use this alternative method for creating a separate pivot cache. Copy and paste the existing table to a new workbook, and then copy and paste the pivot table back to a new sheet in the original workbook.

Side Effects of Sharing a Pivot Cache

It's important to note that there are a few side effects to sharing a pivot cache. For example, suppose you have two pivot tables using the same pivot cache. Certain actions affect both pivot tables:

■ **Refreshing your data**—You cannot refresh one pivot table and not the other. Refreshing affects both tables.

■ **Adding a calculated field**—If you create a calculated field in one pivot table, your newly created calculated field shows up in the other pivot table's field list.

■ **Adding a calculated item**—If you create a calculated item in one pivot table, it shows in the other as well.

■ **Grouping or ungrouping fields**—Any grouping or ungrouping you perform affects both pivot tables. For instance, suppose you group a date field in one pivot table to show months. The same date field in the other pivot table is also grouped to show months.

Although none of these side effects are critical flaws in the concept of sharing a pivot cache, it is important to keep them in mind when determining whether using a pivot table as your data source is the best option for your situation.

Saving Time with New Pivot Table Tools
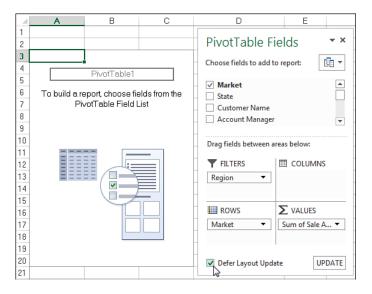
Microsoft has invested a lot of time and effort in the overall pivot table experience. The results of these efforts are tools that make pivot table functionality more accessible and easier to use. The following sections look at a few of the tools that help you save time when managing your pivot tables.

Deferring Layout Updates

The frustrating part of building a pivot table from a large data source is that each time you add a field to a pivot area, you are left waiting while Excel crunches through all that data. This can become a maddeningly time-consuming process if you have to add several fields to your pivot table.

Excel 2013 offers some relief for this problem by providing a way to defer layout changes until you are ready to apply them. You can activate this option by clicking the relatively inconspicuous Defer Layout Update check box in the PivotTable Field List dialog, as shown in Figure 2.32.

Figure 2.32
Click the Defer Layout Update check box to prevent your pivot table from updating while you add fields.

Here's how this feature works: When you place a check in the Defer Layout Update check box, you prevent your pivot table from making real-time updates as you move your fields around without your pivot table. In Figure 2.32, notice that fields in the drop zones are not in the pivot table yet. The reason is that the Defer Layout Update check box is active. When you are ready to apply your changes, click the Update button on the lower-right corner of the PivotTable Field List dialog.

> **NOTE**
>
> Remember to remove the check from the Defer Layout Update check box when you are done building your pivot table. Leaving it checked results in your pivot table remaining in a state of manual updates, preventing you from using the other features of the pivot table (that is, sorting, filtering, and grouping).
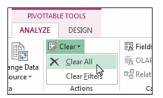
> **TIP**
>
> Incidentally, the Defer Layout Update option is available through VBA. It can help improve the performance of any macro that automates the creation of pivot tables. Using VBA to create pivot tables is covered in Chapter 13, "Using VBA to Create Pivot Tables."

Starting Over with One Click

Often you might want to start from scratch when working with your pivot table layouts. Excel 2013 provides a simple way to essentially start over without deleting your pivot cache. Select Analyze under the PivotTable Tools tab and select the Clear drop-down.

As you can see in Figure 2.33, this command enables you to either clear your entire pivot table layout or remove any existing filters you might have applied in your pivot table.

Figure 2.33
The Clear command enables you to clear your pivot table fields or remove the applied filters in your pivot table.
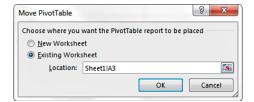
Relocating Your Pivot Table

You might find that, after you have created your pivot table, you need to move it to another location. It might be in the way of other analyses on the worksheet, or you might simply need to move it to another worksheet. Although there are several ways to move your pivot table, Excel 2013 provides a no-frills way to easily change the location of your pivot table.

Select Analyze under the PivotTable Tools tab and select Move PivotTable. This icon activates the Move PivotTable dialog illustrated in Figure 2.34. All you have to do here is specify where you want your pivot table moved.

Figure 2.34
The Move PivotTable dialog enables you to quickly move your pivot table to another location.

Next Steps

In the next chapter, you learn how to enhance your pivot table reports by customizing your fields, changing field names, changing summary calculations, applying formats to data fields, adding and removing subtotals, and using the Show As setting.

2

CHAPTER 8:

Understanding Formulas

The beyond-the-basics, beneath-the-surface guide to Microsoft Excel 2013: comprehensive coverage, real solutions!

- By top Excel expert, trainer, and consultant Bill Jelen, whose MrExcel. com site attracts 10,000,000+ page views every year

- Master Excel 2013's most powerful new features, from Flash Fill to Recommended PivotTable to new collaboration improvements

- Breakthrough techniques, exclusive shortcuts, expert troubleshooting help, and amazingly realistic examples

- For everyone who wants to get the most out of Excel 2013 in any environment

ABOUT THE AUTHOR

Bill Jelen, Excel MVP and the host of MrExcel.com, has been using spreadsheets since 1985, and he launched the MrExcel.com website in 1998. Bill was a regular guest on Call for Help with Leo Laporte and has produced more than 1,500 episodes of his daily video podcast, Learn Excel from MrExcel. He is the author of 39 books about Microsoft Excel and writes the monthly Excel column for Strategic Finance magazine. His Excel tips appear regularly in CFO Excel Pro Newsletter and CFO Magazine . Before founding MrExcel.com, Bill Jelen spent 12 years in the trenches--working as a financial analyst for finance, marketing, accounting, and operations departments of a $500 million public company. He lives near Akron, Ohio, with his wife, Mary Ellen.

UNDERSTANDING FORMULAS

Excel's forté is performing calculations. When you use Excel, you typically use a combination of cells with numbers and cells with formulas. After you design a spreadsheet to calculate something, you can change the numbers used in the assumption cells and then watch Excel instantly calculate new results.

Getting the Most from This Chapter

Do not skip this entire chapter; one particular trick in this chapter can save you daily frustration.

I regularly entertain accountants and auditors with my Power Excel program. Although this program is a fun, laughter-filled tour through the inside tricks of Microsoft Excel, people always learn new things along the way.

I call them "gasp moments."

Imagine this setting: I am in front of 200 managerial accountants who have Excel open for 40 hours each week. You can generally figure that these folks are super-efficient with Excel. Any trick I show might already be in the arsenal of half to three-quarters of the room. A lot of the people nod their heads, while others look surprised. However, a few tricks I show will get a gasp from perhaps 90% of these managerial accountants because they didn't know these tricks before and now realize just how powerful they are. I thrive on the gasp moments.

Most people reading this book believe they know Excel formulas. To a certain extent, this chapter is a primer for the person who is new to Excel. However, even the most astute person using Excel should check out these sections of the chapter:

- Everyone should read the "Double-Click the Fill Handle to Copy a Formula" section. Somehow, most people have learned to drag the fill handle to copy a formula. This leads to horrible frustration on long data sets, as they go flying past the end of the data. This simple but powerful trick is the one that universally amazes attendees of my seminar.

- Honestly answer this question: Do you really understand the difference between cell H1 and cell $H$1? If you think the latter has anything to do with currency, you need to review the "Overriding Relative Behavior: Absolute Cell References" section thoroughly. This isn't a trick, but one of the fundamental building blocks to creating Excel worksheets. Roughly, 5% of the people in a Power Excel seminar do not understand this concept, and about 70% of the people in a community computer club presentation do not understand it. If you don't know when and why to use the dollar signs, you are in good company with 20 million other people using Excel. It is worth taking time to learn this essential technique.

- Finally, there are three ways to enter formulas, and I believe my preferred way is the best. I probably will not convince you to change, but when you understand my way, you can enter formulas far faster than the other two ways. To get a good understanding of the alternatives, read the "Three Methods of Entering Formulas" section later in this chapter.

> **Tip**
>
> Designing a formula that can be written once and then copied to a rectangular range of data is a fantastic way to use Excel more efficiently.

Introduction to Formulas

When Microsoft overhauled Excel for the 2007 version, a number of formula limits were dramatically increased. For example, the number of characters in a formula increased from 1,024 to 8,192. The number of levels of nesting for IF functions increased from 7 to 64. Thanks to those improvements, you can calculate almost anything with a formula in Excel.

This chapter and Chapter 9, "Controlling Formulas," deal with formula basics. The chapters between Chapter 10, "Understanding Functions," and Chapter 15, "Using Trig, Matrix, and Engineering Functions," introduce adding functions to your formulas. Chapter 16, "Connecting Worksheets, Workbooks, and External Data," introduces formulas that calculate data found on other worksheets or in other workbooks. Chapter 17, "Using Super Formulas in Excel," provides interesting examples such as 3D formulas and the all-powerful array formulas.

Because of the record-oriented nature of spreadsheets, you can generally build a formula once and then copy that formula to hundreds or thousands of cells without changing anything in the formula.

Formulas Versus Values

When looking at an Excel grid, you cannot tell the difference between a cell with a formula and one that contains numbers. To see if a cell contains a number or a formula, select the cell. Look in the formula bar. If the formula bar contains a number, as shown in Figure 8.1, you know that it is a static value. If the formula bar contains a formula, you know that the number shown in the grid is the result of a formula calculation (see Figure 8.2). Keep in mind that formulas start with an equal sign.

Figure 8.1
The formula bar reveals whether a value is a static number or a calculation. In this case, cell E2 contains a static number.

Static Value in the Formula Bar

Figure 8.2
In this case, cell E2 contains the result of a formula calculation. A formula starts with an equal sign.

Entering Your First Formula

Your first formula was probably a SUM function, entered with the Quick Analysis tool or the AutoSum button. However, this discussion is talking about a pure mathematical formula that uses a value in a cell that's added, subtracted, divided, or multiplied by a number or another cell.

Billions of variations of formulas can be used. Everyday life throws situations at you that can be solved with a formula. Keep these important points in mind as you start tinkering with your own formulas:

- Every formula starts with an equal sign.

- Entering formulas is just like typing an equation in a calculator with one exception (see the next point).

- If one of the terms in your formula is already stored in a cell in Excel, you can point to that cell's address instead of typing the number into that cell. Using this method enables you to change the value in one cell and then watch all the formulas recalculate. Excel 2013 adds a "slot machine" animation to show the cells in the visible window that are recalculating as the result of changing a cell.

To illustrate these points, see the steps to building a basic formula included in the following example.

Building a Formula

You want to enter a formula to calculate a target sales price, as shown in Figure 8.3. Cell D2 shows the product cost. In column E, you want to calculate the list price as two times the cost plus $3.

Figure 8.3
The formula in cell E2 recalculates if the value in cell D2 changes.

To enter a formula, follow these steps:

1. Select cell E2.

2. Type an equal sign. The equal sign tells Excel that you are starting a formula.

3. Type **2*D2** to indicate that you want to multiply two times the value in cell D2.

4. Type **+3** to add three to the result. There should be no spaces in the formula. If your formula reads =2*D2+3, proceed to step 5. Otherwise, use the backspace key to correct the formula.

5. Press Enter. Excel calculates the formula in cell E2.

By default, Excel usually moves the cell pointer down or to the right after you finish entering a formula. You should move the cell pointer back to cell E2 to inspect the formula, as shown in Figure 8.3. Note that Excel shows a number in the grid, but the formula bar reveals the formula behind the number.

The Relative Nature of Formulas

The formula =2*D2+3 really says, "multiply two by the cell immediately to the left of me and then add three." If you need to put this formula in cells E3 to E999, you do not need to reenter the formula 997 times. Instead, copy the formula and paste it to all the cells. When you do, Excel copies the essence of the formula: "Multiply two by the cell to the left of me and add three." As you copy the formula to cell E3, the formula becomes =2*D3+3. Excel handles all this automatically. Figure 8.4 shows the formula after it is copied.

Figure 8.4
After you paste the formula, Excel automatically updates the cell reference to point to the current row.

Excel's capability to change D2 to D3 in the formula is called *relative referencing*. This is the default behavior of a reference. Sometimes, you do not want Excel to change a reference as the formula is copied, as explained in the next section.

Overriding Relative Behavior: Absolute Cell References

Relative referencing, which is Excel's ability to change a formula as it is copied, is what makes spreadsheets so useful. At times, however, you need part of a formula to always point at one particular cell. This happens a lot when you have a setting at the top of the worksheet such as a growth rate or a tax rate. It would be nice to change this cell once and have all the formulas use the new rate.

The following example sets up a sample worksheet that exhibits this problem and shows how to use an arcane notation style to solve the problem. When you see a reference with two dollar signs, such as $G$1, this indicates an absolute reference to G1. An absolute reference is a cell or range address where the row numbers and the column letters are locked and will not change during copying. Absolute references have a dollar sign before each column letter and each row number. Examples include $G$1 and $T$2:$W$99.

Suppose that you have a sales tax factor in a single cell at the top of a worksheet. After you enter the formula =C2*G1, it accurately calculates the tax in cell D2, as shown in Figure 8.5.

However, when you copy the same formula to cell D3, you get a zero as the result. As you can see in Figure 8.6, Excel correctly changed cell C2 to C3 in the copied formula. However, Excel also changed G1 to G2. Because there is nothing in G2, the formula calculates a zero.

Figure 8.5
This formula works fine in row 2.

Figure 8.6
This formula fails in row 3.

The formula now points to empty cell G2.

Because the sales tax factor is only in G1, you want Excel to always point to G1. To make this happen, you need to build the original formula as =C2*$G$1. The two dollar signs tell Excel that you do not want to have the reference change as the formula is copied. The $ before the G freezes the reference to always point to column G. The $ before the 1 freezes the reference to always point to row 1. Now, when you copy this formula from cell D2 to other cells in column D, Excel changes the formula to =C3*$G$1, as shown in Figure 8.7.

To recap, a reference with two dollars signs is called an *absolute reference*.

D3		:	✕	✓	*fx*	=C3*$G$1	
	A	B	C	D	E	F	G
1	SKU	Mfg Cost	List Price	Sales Tax		Tax Factor	6.25%
2	J41	16.4	35.8	2.24			
3	J20	14.47	31.94	2.00			
4	I51	16.14	35.28	2.21			
5	F69	14.31	31.62	1.98			
6	G61	14.81	32.62	2.04			

Figure 8.7
The dollar signs in the formula make sure that the copied formula always points to cell G1.

If you are never going to copy the formula to the left or right, you can safely use =C2*G$1. This formula freezes only the row number. Given the shape of the current data, it is likely that using a single dollar sign will be valid.

Using Mixed References to Combine Features of Relative and Absolute References

In a number of situations, you might want to build a reference that has only one dollar sign. For example, in Figure 8.8, you want to use the monthly bonus rate in row 3, but you want to allow the column to change. In this case, the formula for cell B13 would be =B6*B$3.

When you copy this formula, it always points to the bonus amount in row 3, but the remaining elements of the formula are relative. For example, the formula in E15 is =E8*E$3, which multiplies Dom's April sales by the April bonus rate.

There are two kinds of mixed references. One mixed reference freezes the row number and allows the column letter to change, as in A$1. The other mixed reference freezes the column letter but allows the row number to change, as in $A1. No one has thought up clever names to distinguish between these references, so they are simply called *mixed references*.

To illustrate the other kind of mixed reference, as shown in Figure 8.9, say you want a single formula to multiply the daily rate from column A by the number of days in row 4. This formula requires both kinds of mixed references.

In this case, you want the cell A6 reference to always point to column A, even when the formula is copied to the right. Therefore, the A6 portion of the formula should be entered as $A6. You also want the C5 portion of the formula to always point to row 5, even when the formula is copied down the rows. Therefore, the C5 portion of the formula should be entered as C$5.

Figure 8.8

By having the dollar sign before the 3 in C$3, you lock the reference to row 3 but allow the formula to point to columns D, E, and so on as you copy the formula.

| B13 | | : | ✕ | ✓ | *fx* | =B6*B$3 | |

▲	A	B	C	D	E	F	G
1		**Widget Sales Bonus Calculation**					
2							
3	**Bonus**	3%	2%	0%	2%	3%	1%
4							
5	**$ Sold**	Jan	Feb	Mar	Apr	May	Jun
6	**Zeke**	6237	9009	5247	9207	7029	6435
7	**Josh**	7722	5544	6138	5445	9603	5544
8	**Dom**	8712	7524	6237	6336	8415	7623
9	**Alec**	5148	7524	9603	5049	5643	5445
10	**Logan**	5841	6534	8118	8910	9009	9405
11							
12	**Bonus**	Jan	Feb	Mar	Apr	May	Jun
13	**Zeke**	187.11	180.18	0.00	184.14	210.87	64.35
14	**Josh**	231.66	110.88	0.00	108.90	288.09	55.44
15	**Dom**	261.36	150.48	0.00	126.72	252.45	76.23
16	**Alec**	154.44	150.48	0.00	100.98	169.29	54.45
17	**Logan**	175.23	130.68	0.00	178.20	270.27	94.05

Figure 8.9

You can create a formula by using a combination of dollar signs to allow cell C6 to be copied to all cells in the table.

| XOR | | : | ✕ | ✓ | *fx* | =A6*C5 | |

▲	A	B	C	D	E	F	G
1		**XYZ Tool Rental**					
2		**Price List**					
3							
4			------- *Number of Hours* ---->				
5	**Per Hour**	**Item**	4	8	24	48	72
6	13.50	Aerator TA-17D Split Drive	=A6*C5				
7	18.00	Aerator TA-25D Split Drive					
8	10.75	Aerator Tow Behind 36" w/weight canisters					
9	10.00	Auger One Man					
10	12.50	Auger Two Man					

Using the F4 Key to Simplify Dollar Sign Entry

In the preceding section, you entered quite a few dollar signs in formulas. The good news is that you do not have to type the dollar signs! Instead, immediately after entering a reference, press the F4 key to toggle the reference from a relative reference to an absolute reference, which automatically has the dollar signs before the row and column. If you press F4 again, the reference toggles to a mixed reference with a dollar sign before the row number. When you press F4 once again, the reference toggles to a mixed reference with a dollar sign before the column letter. Pressing F4 one more time returns the reference to a relative reference. You might find it easier to choose the right reference by looking at the various reference options offered by the F4 key.

The following sequence shows how the F4 key works while you are entering a formula. This particular example was included because it requires two different types of mixed references.

The important concept is that you start pressing F4 after typing a cell reference but before you type a mathematical operator.

1. Type **=A6** (see Figure 8.10).
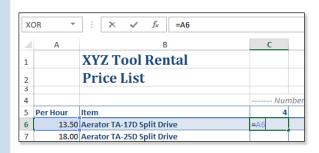
Figure 8.10
This is a relative reference: A6 without any dollar signs.

2. Before typing the asterisk to indicate multiplication, press the F4 key. On the first press of F4, the reference changes to =$A$6, as shown in Figure 8.11.
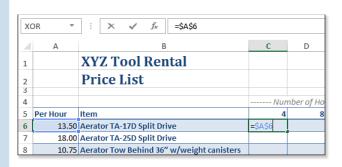
Figure 8.11
After you press F4 once, the reference changes to an absolute reference, with two dollar signs.

3. Press the F4 key again. The reference changes to A$6 to freeze the reference to row 6, as shown in Figure 8.12. This still isn't right because freezing the reference to row 6 will not help.

4. Press F4 one more time. Excel locks just the column, changing the reference to =$A6, as shown in Figure 8.13. This is the version of the reference you want. As you copy the formula across, the formula always points back to column A. As you copy the formula down, the row number in this reference is allowed to change to point to other rows.

Figure 8.12
Press F4 again to switch to a mixed reference with the row number locked.
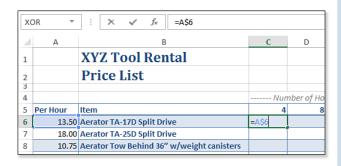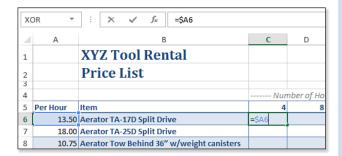
Figure 8.13
Press F4 again to switch to a mixed reference with the column letter locked.
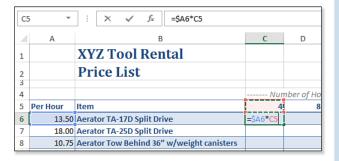
5. To continue the formula, type an asterisk to indicate multiplication and then click cell C5 with the mouse. At the point shown in Figure 8.14, you would press F4 twice to change C5 to a reference that locks only the row (that is, C$5).

6. Press Enter to accept the formula.

Figure 8.14
After clicking C5 with the mouse, press F4 twice to change this reference.

7. When you copy the formula from cell C6 to the range C6:G28, the formula automatically multiplies the rate in column A by the number of days in row 5. Figure 8.15 shows the copied formula in cell E9. The formula correctly multiplies the 10-dollar rate in cell A9 by the 24 hours figure in cell E5.

Figure 8.15
By using the correct combination of row and column mixed references, you can enter this formula once and successfully copy it to the entire rectangular range.

Using F4 After a Formula Is Entered

The F4 trick described in the preceding section works immediately after you enter a reference. If you try to change cell A6 after you type the asterisk, pressing the F4 key has no effect.

However, you can still use F4 by clicking somewhere in the formula bar adjacent to the characters A6. Pressing F4 now adds dollar signs to that reference.

Using F4 on a Rectangular Range

Some functions allow you to specify a rectangular range. For example, in Figure 8.16, you would like to enter a formula to calculate month-to-date sales. Although =SUM(B2:B29) works for cell C29, you cannot copy this formula to the other cells in the column. To copy this formula, you need to change the formula to =SUM(B$2:B29).

Note

After you press F4 again, Excel returns the reference to the relative state A6. As you continue to press F4, Excel toggles between the four modes. It is fine to toggle between them all and then choose the correct one. If you accidentally toggle past the $A6 version, just keep pressing F4 until the correct mode comes up again.

At this point in the figure, you might be tempted to press the F4 key. This does not work. If you select B2:B29 with the mouse or arrow keys, pressing F4 now converts the reference to the fully absolute range $B$2:$B$29. Continuing to press F4 toggles to B$2:B$29, then $B2:$B29, and then B2:B29. Excel does not even attempt to go through the other 12 possible combinations of dollar signs to offer B$2:B29 eventually. If you typed B2:B29, pressing F4 adjusts only the B29 reference.

In this case, you need to click the insertion point just before, just after, or in the middle of the characters B2 in the formula. If you then press F4, toggle through the various dollar sign combinations on the B2 reference. Pressing F4 twice results in the proper combination, as shown in Figure 8.17.

Figure 8.16
Using F4 at this point never produces the desired result of B$2:B29.

	A	B	C	D	E
1	Date	Daily Sales	Sales MTD		
2	2/1/2015	1387			
3	2/2/2015	1956			
4	2/3/2015	1789			
5	2/4/2015	1208			
6	2/5/2015	1096			
7	2/6/2015	1869			
8	2/7/2015	1205			
9	2/8/2015	1858			
10	2/9/2015	1232			
11	2/10/2015	1859			
12	2/11/2015	1014			
22	2/2../2015	1401			
23	2/22/2015	1897			
24	2/23/2015	1781			
25	2/24/2015	1114			
26	2/25/2015	1254			
27	2/26/2015	1401			
28	2/27/2015	1761			
29	2/28/2015	1627	=sum(B2:B29		
30			SUM(**number1**, [number2], …)		
31					
32					
33					

XOR ▼ ⋮ × ✓ *fx* =sum(B2:B29)

Figure 8.17
Using F4 is tricky when your reference is a rectangular range—you must click into the formula.

28	2/27/2015	1761		
29	2/28/2015	1627	=sum(B$2:B29	
30			SUM(**number1**, [number2], …)	
31				

Three Methods of Entering Formulas

In the examples in the previous sections, you entered a formula by typing it. You generally need to start a formula by typing the equal sign (or the plus sign); after that point, you have three options:

- Type the complete formula as described in the previous sections.

- Type operator keys, but use the mouse to touch cell references. In this book, this is referred to as the *mouse method*.

- Type the operator keys and then use the arrow keys to specify the cell references by navigating to the cells. In this book, this method is referred to as the *arrow key method*.

Assume you would like to multiply the merchandise total in cell B2 by the sales tax rate in cell F1, as shown in Figure 8.18.
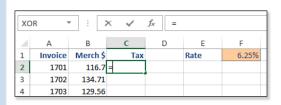
Figure 8.18
You can use three methods to enter the formula
=B2*$F$1.

Enter Formulas Using the Mouse Method

If you started using computers in the last 20 years, it is likely that you use the mouse method for entering formulas. This method is intuitive, but it requires you to move your hand between the keyboard and the mouse several times, as in this example:

1. Type = or +.

2. Click in cell B2.

3. Type *.

4. Click in cell F1.

5. Press F4 to add the dollar signs.

6. Press Enter. This usually moves the cell pointer to cell F2.

This method requires only four keystrokes, but it requires you to move to the mouse twice. Moving to the mouse is the slowest part of entering formulas, but this method is easier than typing the entire formula if you are not a touch typist.

Entering Formulas Using the Arrow Key Method

The arrow key method is popular with people who started using spreadsheets in the days of Lotus 1-2-3 release 2.2. It is worthwhile to learn this method because it is incredibly fast. Almost all formula entry can be accomplished using keys on the right side of the keyboard. Here's how it works:

1. In cell C2, type +.

2. Press the left-arrow key to move the flashing cell border to cell B2. Note that the active cell, which is the one with a green solid border, is still cell C2. The flashing border is like a second cell pointer that you can use to point to the correct cell for the formula. As shown in Figure 8.19, the temporary formula in the formula bar reads +B2.

Tip
If you have a desktop keyboard, you can use the asterisk key on the numeric keypad to avoid pressing the Shift key.
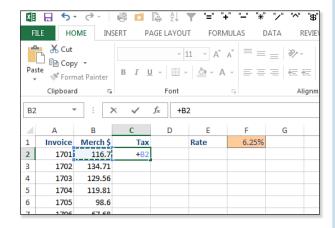
Tip
If you use the mouse method to enter formulas, customize the Quick Access Toolbar (QAT) to icons for Equal Sign, Plus Sign, Minus Sign, Multiplication Sign, Division Sign, Exponentiation Sign, and Dollar Sign. You can then enter most formulas without reaching back to the keyboard. There isn't a QAT icon for the Enter key—use the green check mark to the left of the formula bar for Enter.

Figure 8.19
By using the arrow keys during formula entry, you create a flashing border that can be used to navigate to a cell reference.

Note

As you are moving the flashing cell border with the mouse, ignore the formula bar and watch just the flashing cell border.
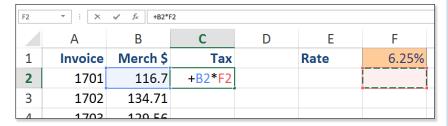
3. To accept cell B2 as the correct reference in the formula, press either an operator key (for example, * or +), a parenthesis, or the Enter key. In this case, type *.

4. Note that the dashed cell pointer disappears, and the focus is now back to the original cell, C2.

5. Press the right-arrow key three times. The flashing cell border moves to D2, E2, and then F2. With each keypress, the temporary formula in the formula bar shows an incorrect formula (+B2*D2, +B2*E2, and +B2*F2). Figure 8.20 shows what the screen looks like after you press the right-arrow key three times.

Tip
Even if you are mouse-centric, you should try this method for half a day. When you get the feel for navigating by using the arrow keys, you can enter formulas much faster by using this method.

Figure 8.20
After step 4, the focus moves to the original cell. Thus, you only have to press the right-arrow key three times instead of four times to arrive at cell F2.

6. Press the up-arrow key to move the flashing cell border to the correct location, cell F1. The temporary formula in the formula bar now shows +B2*F1.

7. Press the F4 key to add dollar signs to the F1 reference.

8. Press Ctrl+Enter to accept the formula and keep the cell pointer in cell C2.

Using this method requires ten keystrokes, with no trips to the mouse. You can enter formulas that have no absolute references, mixed references, parentheses, or exponents by using just the arrow keys and the keys on the numeric keypad.

🔍 Note

Officially, every formula must start with an equal sign. However, to make former Lotus 1-2-3 users comfortable, Excel allows you to start a formula with a plus sign. Power Excel users have discovered that using a plus sign enables them to start a formula by typing on the numeric keypad. Because I routinely start formulas with the plus sign, I am often asked why I start with =+ instead of just =. Even though the formulas appear that way onscreen, I don't actually enter the plus sign. When a formula starts with a plus sign, Excel adds an equal sign and does not remove the plus sign, so you end up with a formula that looks like =+B2*$F$1.

Entering the Same Formula in Many Cells

So far in this chapter, you have entered a formula in one cell and then copied and pasted to get the formula in many cells. To enter the same formula in many cells, you can use three alternatives:

- Preselect the entire range where the formulas need to go. Enter the formula for the first cell and press Ctrl+Enter to enter the formula in the entire selection simultaneously.

- Enter the formula in the first cell and then use the fill handle to copy the formula.

- Beginning with Excel 2007, the method is to define the range as a table. When you use this method, the new formulas are copied down a column automatically.

Copying a Formula by Using Ctrl+Enter

This strategy works when you are entering formulas for one or more screens that are full of data:

1. If you have just a few cells, select them before entering the formula.

2. Click in the first cell and drag down to the last cell, as shown in Figure 8.21. Notice from the name box that the active cell is the first cell.

3. Enter the formula by using any of the three methods described earlier in this chapter. Even if you use the arrow key method, Excel keeps the entire range selected. Figure 8.22 shows a formula after you press F4 to convert the F1 reference to $F$1.

Figure 8.21

Click in the first cell and drag down to the last cell to select a range with the first cell as the active cell.

	A	B	C	D	E	F
1	Invoice	Merch $	Tax		Rate	6.50%
2	1701	116.7				
3	1702	134.71				
4	1703	129.56				
5	1704	119.81				
6	1705	98.6				
7	1706	67.68				
8	1707	175.4				
9	1708	177.92				
10	1709	71.89				
11	1710	75.05				
12	1711	76.62				
13	1712	123.03				
14	1713	151.12				
15						

C2

Figure 8.22

Even with a large range selected, the formula is built only in the active cell.

XOR   =B2*$F$1

	A	B	C	D	E	F
1	Invoice	Merch $	Tax		Rate	6.50%
2	1701	116.7	=B2*$F$1			
3	1702	134.71				
4	1703	129.56				
5	1704	119.81				
6	1705	98.6				
7	1706	67.68				
8	1707	175.4				
9	1708	177.92				
10	1709	71.89				
11	1710	75.05				
12	1711	76.62				
13	1712	123.03				
14	1713	151.12				
15						

4. At this point, you would normally press Enter to complete the formula. Instead, press Ctrl+Enter to enter this formula in the entire selected range. Note that Excel does not enter =B2*$F$1 in each cell. Instead, it converts the formula as if it were copied to each cell. Figure 8.23 shows the formula in cell C10.

Figure 8.23
Pressing Ctrl+Enter tells Excel to enter the formula in the active cell and to copy it to the rest of the selection.

Copying a Formula by Dragging the Fill Handle

If you want to enter a formula in one cell and then copy it to the other cells in a range, you can use the fill handle, which is the square dot in the lower-right corner of the cell pointer. There are two ways to use the fill handle:

- Drag the fill handle.

- Double-click the fill handle.

The dragging method works fine when you have less than one screen full of data:

1. Enter the formula in cell B2.

2. Press Ctrl+Enter to accept the formula and keep the cell pointer in cell B2.

3. Click the fill handle. You know that you are above the fill handle when the mouse pointer changes to a thick plus sign, as shown in Figure 8.24. Drag the mouse down to the last row of data, which in this case is cell C14.
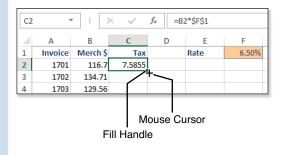
Figure 8.24
You can copy a formula by clicking and dragging the fill handle.

Mouse Cursor

Fill Handle

4. When you release the mouse button, the original cell is copied to all the cells in the selected range.

This method is fine for copying a formula to a few cells. However, if you have thousands or hundreds of thousands of cells, it is annoying to drag to the last row. In Excel 2003 and earlier, you would invariably end up flying past the last row. Note that Excel 2013 automatically slows down and briefly pauses at the last row. However, it is far easier to copy a formula by double-clicking the fill handle.
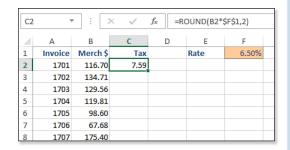
Double-Click the Fill Handle to Copy a Formula

In most data sets, double-clicking the fill handle is the fastest way to copy the formula. Although you will love this method, you need to understand a few shortcomings that can hamper the method when an adjacent column has blank cells among the data.

Figure 8.25 shows a table that has hundreds of rows of data. Suppose you want to copy the formula from cell C2 down to all the rows of data in column B.

Figure 8.25
Dragging the fill handle is frustrating in a table that has hundreds of rows. The double-click method will end that frustration.

In this particular case, cell B2 is nonblank, and column B contains a value in every row down to the end of the data. This is the perfect condition for using the technique of double-clicking the fill handle. Follow these steps:

1. Enter the formula in cell B2.

2. Press Ctrl+Enter to accept the formula and keep the cell pointer in cell B2.

3. Double-click the fill handle.

The active cell is copied down to the last row of your data, as shown in Figure 8.26.

The fill handle double-click method is fast. Excel 2013 follows any diagonal path of nonempty cells to figure out where the bottom of your data set is. You should always press End and then press the down arrow to make sure that the formula was copied far enough.

Figure 8.26
Double-click the fill handle to copy to the last row of your data.

Use the Table Tool to Copy a Formula

The table tool was improved dramatically in Excel 2007. When using this tool, if you tell Excel that your current data set is a table, Excel automatically copies new formulas down to the rest of the cells in the table.

Figure 8.27 shows an Excel worksheet that has headings at the top and many rows of data below the headings.
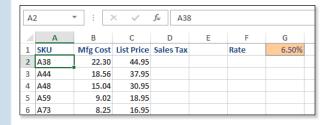
Figure 8.27
This is a typical worksheet in Excel.

To define a range as a table, select a cell within the data set and press Ctrl+T. Excel uses its IntelliSense to guess the edges of the table. If its guess is correct, click OK in the Create Table dialog, as shown in Figure 8.28.
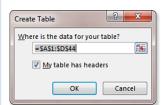
Figure 8.28
The Create Table dialog.

Ctrl+T is one of four entry points for creating a table. You can still use the Excel 2003 shortcut of Ctrl+L. You can choose Format as Table on the Home tab. You can choose the Table icon from the Insert tab.

As shown in Figure 8.29, after Excel recognizes the range as a table, several changes occur:

- The table is formatted with the default formatting. Depending on your preferences, this might include banded rows or columns.

- AutoFilter drop-downs are added to the headings.

- Any formulas you enter use the headings to refer to cells within the table.

Figure 8.29
Defining a range as a table provides formatting and powerful features such as autofilters and natural language formulas.

Now when you enter a formula in the table, Excel automatically copies that formula down to all rows of the table.
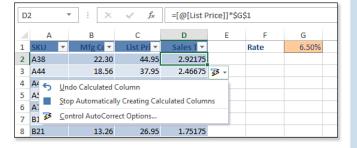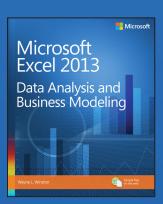
Note

As shown in Figure 8.30, a lightning bolt drop-down appears to the right of cell D3. This drop-down offers you the opportunity to stop Excel from automatically copying the formula down.

Figure 8.30
Thanks to the Table tool in Excel, a new formula entered anywhere in column D is copied automatically to all the cells in column D.

Microsoft
Excel 2013
Data Analysis and
Business Modeling

Wayne L. Winston

Sample files on the web

ISBN: 9780735669130

CHAPTER 8:

IF statements

Solve real business problems with Excel—and sharpen your edge!

- Summarize data with PivotTables and Descriptive Statistics
- Explore new trends in predictive and prescriptive analytics
- Use Excel Trend Curves, multiple regression, and exponential smoothing
- Master advanced Excel functions such as OFFSET and INDIRECT
- Delve into key financial, statistical, and time functions
- Make your charts more effective with the Power View tool
- Tame complex optimization problems with Excel Solver
- Run Monte Carlo simulations on stock prices and bidding models
- Apply important modeling tools such as the Inquire add-in

ABOUT THE AUTHOR

Wayne L. Winston is a professor of Decision Sciences at Indiana University's Kelley School of Business and has earned numerous MBA teaching awards. For 20+ years, he has taught clients at Fortune 500 companies how to use Excel to make smarter business decisions. Wayne and his business partner Jeff Sagarin developed the player-statistics tracking and rating system used by the Dallas Mavericks professional basketball team. He is also a two time Jeopardy! Champion.

Microsoft Press

IF statements

Questions answered in this chapter:

- If I order up to 500 units of a product, I pay $3.00 per unit. If I order from 501 through 1,200 units, I pay $2.70 per unit. If I order from 1,201 through 2,000 units, I pay $2.30 per unit. If I order more than 2,000 units, I pay $2.00 per unit. How can I write a formula that expresses the purchase cost as a function of the number of units purchased?

- I just purchased 100 shares of stock at a cost of $55 per share. To hedge the risk that the stock might decline in value, I purchased 60 six-month European put options. Each option has an exercise price of $45 and costs $5. How can I develop a worksheet that indicates the six-month percentage return on my portfolio for a variety of possible future prices?

- Many stock market analysts believe that moving-average trading rules can outperform the market. A commonly suggested moving-average trading rule is to buy a stock when the stock's price moves above the average of the past 15 months and to sell a stock when the stock's price moves below the average of the past 15 months. How would this trading rule have performed against the Standard & Poor's 500 Stock Index (S&P)?

- In the game of craps, two dice are tossed. If the total of the dice on the first roll is 2, 3, or 12, you lose. If the total of the dice on the first roll is 7 or 11, you win. Otherwise, the game keeps going. How can I write a formula to determine the status of the game after the first roll?

- In most pro forma financial statements, cash is used as the plug to make assets and liabilities balance. I know that using debt as the plug would be more realistic. How can I set up a pro forma statement having debt as the plug?

- When I copy a VLOOKUP formula to determine salaries of individual employees, I get a lot of #NA errors. Then when I average the employee salaries, I cannot get a numerical answer because of the #NA errors. Can I easily replace the #NA errors with a blank space so I can compute average salary?

- My worksheet contains quarterly revenues for Wal-Mart. Can I easily compute the revenue for each year and place it in the row containing the first quarter's sales for that year?

- IF statements can get rather large. How many IF statements can I nest in a cell? What is the maximum number of characters allowed in an Excel formula?

The situations listed here seem to have little, if anything, in common. However, setting up Microsoft Excel 2013 models for each of these situations requires the use of an IF statement. The IF formula is probably the single most useful formula in Excel. With IF formulas, you can conduct conditional tests on values and formulas, mimicking (to a limited degree) the conditional logic provided by computing languages such as C, C++, and Java.

An IF formula begins with a condition such as A1>10. If the condition is true, the formula returns the first value listed in the formula; otherwise, you move on within the formula and repeat the process. The easiest way to show you the power and utility of IF formulas is to use them to help answer each of this chapter's questions.

Answers to this chapter's questions

This section provides the answers to the questions that are listed at the beginning of the chapter.

If I order up to 500 units of a product, I pay $3.00 per unit. If I order from 501 through 1,200 units, I pay $2.70 per unit. If I order from 1,201 through 2,000 units, I pay $2.30 per unit. If I order more than 2,000 units, I pay $2.00 per unit. How can I write a formula that expresses the purchase cost as a function of the number of units purchased?

You can find the solution to this question on the Quantity Discount worksheet in the Ifstatement.xlsx file. The worksheet is shown in Figure 11-1.

◢	A	B	C	D
1		cutoff	price	
2	cut1	500	$ 3.00	price1
3	cut2	1200	$ 2.70	price2
4	cut3	2000	$ 2.30	price3
5			$ 2.00	price4
6				
7				
8	order quantity	cost	per unit cost	
9	450	$1,350.00	$ 3.00	
10	900	$2,430.00	$ 2.70	
11	1450	$3,335.00	$ 2.30	
12	2100	$4,200.00	$ 2.00	

FIGURE 11-1 You can use an IF formula to model quantity discounts.

Suppose cell A9 contains your order quantity. You can compute an order's cost as a function of the order quantity by implementing the following logic:

- If A9 is less than or equal to 500, the cost is 3*A9.

- If A9 is from 501 through 1,200, the cost is 2.70*A9.

- If A9 is from 1,201 through 2,000, the cost is 2.30*A9.

- If A9 is more than 2,000, the cost is 2*A9.

Begin by linking the range names in A2:A4 to cells B2:B4 and linking the range names in cells D2:D5 to cells C2:C5. Then you can implement this logic in cell B9 with the following formula:

```
IF(A9<=_cut1,price1*A9,IF(A9<=_cut2,price2*A9,IF(A9<=_cut3,price3*A9,price4*A9)))
```

To understand how Excel computes a value from this formula, recall that IF statements are evaluated from left to right. If the order quantity is less than or equal to 500 (cut1), the cost is given by price 1*A9. If the order quantity is not less than or equal to 500, the formula checks to see whether the order quantity is less than or equal to 1,200. If this is the case, the order quantity is from 501 through 1,200, and the formula computes a cost of price 2*A9. Next, the formula checks whether the order quantity is less than or equal to 2,000. If this is true, the order quantity is from 1,201 through 2,000, and the formula computes a cost of price 3*A9. Finally, if the order cost has not yet been computed, the formula defaults to the value price 4*A9. In each case, the IF formula returns the correct order cost. Note that three other order quantities are entered in cells A10:A12, and the cost formula is copied to B10:B12. For each order quantity, the formula returns the correct total cost.

An IF formula containing more than one IF statement is called a *nested IF formula*.

I just purchased 100 shares of stock at a cost of $55 per share. To hedge the risk that the stock might decline in value, I purchased 60 six-month European put options. Each option has an exercise price of $45 and costs $5. How can I develop a worksheet that indicates the six-month percentage return on my portfolio for a variety of possible future prices?

Before tackling this problem, I want to review some basic concepts from the world of finance. A European put option allows you to sell a share of a stock at a given time in the future (in this case, six months) for the exercise price (in this case, $45). If the stock's price in six months is $45 or higher, the option has no value. Suppose, however, that the price of the stock in 6 months is below $45. Then you can make money by buying a share and immediately reselling the stock for $45. For example, if in 6 months the stock is selling for $37, you can make a profit of $45 – $37, or $8 per share, by buying a share for $37 and then using the put to resell the share for $45. You can see that put options protect you against downward moves in a stock price. In this case, whenever the stock's price in six months is below $45, the puts start accruing some value. This cushions a portfolio against a decrease in value of the shares it owns. Note also that the percentage return on a portfolio (assume that the stocks in the portfolio do not pay dividends) is computed by dividing the number of the change in the portfolio's value (final portfolio value–initial portfolio value) by the portfolio's initial value.

With this background, look at how the six-month percentage return on this portfolio, consisting of 60 puts and 100 shares of stock, varies as the share price varies between $20 and $65. You can find this solution on the Hedging worksheet in the Ifstatement.xlsx file. The worksheet is shown in Figure 11-2.

The labels in A2:A7 are linked to cells B2:B7. The initial portfolio value is equal to 100($55)+60($5)=*$5,800*, shown in cell B7. By copying the IF(A9<exprice,exprice–A9,0)*Nputs formula from B9 to B10:B18, you can compute the final value of the puts. If the six-month price is less than

the exercise price, you can value each put as *exercise price–six-month price*. Otherwise, each put will have a value of $0 in six months. Copy the Nshares*A9 formula from C9 to C10:C18 and compute the final value of the shares. Copying the ((C9+B9)–startvalue)/startvalue) formula from D9 to D10:D18 computes the percentage return on the hedged portfolio. Copying the (C9–Nshares*pricenow)/(Nshares*pricenow) formula from E9 to E10:E18 computes the percentage return on the portfolio if you are unhedged (that is, you have no puts).

	A	B	C	D	E
1					
2	Nputs	60			
3	Nshares	100			
4	exprice	$ 45.00			
5	pricenow	$ 55.00			
6	putcost	$ 5.00			
7	startvalue	$5,800.00			
8	final stock price	final put value	final share value	percentage return hedged	percentage return unhedged
9	$ 20.00	$1,500.00	$2,000.00	-39.7%	-63.6%
10	$ 25.00	$1,200.00	$2,500.00	-36.2%	-54.5%
11	$ 30.00	$ 900.00	$3,000.00	-32.8%	-45.5%
12	$ 35.00	$ 600.00	$3,500.00	-29.3%	-36.4%
13	$ 40.00	$ 300.00	$4,000.00	-25.9%	-27.3%
14	$ 45.00	$ -	$4,500.00	-22.4%	-18.2%
15	$ 50.00	$ -	$5,000.00	-13.8%	-9.1%
16	$ 55.00	$ -	$5,500.00	-5.2%	0.0%
17	$ 60.00	$ -	$6,000.00	3.4%	9.1%
18	$ 65.00	$ -	$6,500.00	12.1%	18.2%

FIGURE 11-2 This is a hedging example that uses IF statements.

In Figure 11-2, you can see that if the stock price drops below $45, the hedged portfolio has a larger expected return than the unhedged portfolio. Note also that if the stock price does not decrease, the unhedged portfolio has a larger expected return. This is why the purchase of puts is often referred to as *portfolio insurance*.

Many stock market analysts believe that moving-average trading rules can outperform the market. A commonly suggested moving-average trading rule is to buy a stock when the stock's price moves above the average of the previous 15 months and to sell a stock when the stock's price moves below the average of the previous 15 months' price. How would this trading rule have performed against the Standard & Poor's 500 Stock Index (S&P)?

In this example, compare the performance of the moving-average trading rule (in the absence of transaction costs for buying and selling stock) to a buy-and-hold strategy. The strength of a moving-average trading rule is that it helps you follow market trends. A moving-average trading rule lets you ride up with a bull market and sell before a bear market destroys you.

The data set contains the monthly value of the S&P 500 Index for the period of January, 1871, through October, 2002. To track the performance of the moving-average trading strategy, you need to track the following information each month:

- What is the average of the S&P 500 Index over the past 15 months?

- Do I own stock at the beginning of each month?

- Do I buy stock during the month?

- Do I sell stock during the month?

- What is the cash flow for the month (positive if I sell stock, negative if I buy stock, and 0 otherwise)?

The worksheet for this situation requires you to scroll down many rows. It helps to keep columns A and B, as well as the headings in row 8, visible as you scroll down. To do this, in the Matradingrule .xlsx file, move the cursor to cell C9, choose **View** on the ribbon, and then select **Freeze Panes** in the **Windows** group. This presents the choices shown in Figure 11-3.
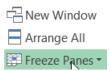
FIGURE 11-3 Choose from the Freeze Panes options.

Here you choose **Freeze Panes**, which keeps columns A and B and rows 6–8 visible as you scroll through the worksheet. **Freeze Top Row** keeps just the top row visible while scrolling through the rest of the worksheet. For example, if the visible top row is row 6, you see row 6 no matter how far down you scroll. If you choose **Freeze First Column**, you always see the leftmost column as you scroll through the worksheet. Selecting **Unfreeze Panes** from the menu returns you to normal worksheet view.

The Matradingrule.xlsx file, shown in Figure 11-4, includes the formulas needed to track the effectiveness of a moving-average strategy. Tackling this problem requires several IF formulas, and some of the IF formulas require an *AND* operator. For example, you buy the stock during a month if and only if you don't own the stock at the beginning of the month, and the current month's price is higher than the 15-month moving average for the stock's price. The first month for which you can compute a 15-month moving average is April 1872, so begin your calculations in row 24.

	S&P				MA profit	$1,319.75
	Comp.				Buy and hold profit	$849.45
Date	P	Own?	MA	Buy?	Sell?	Cash flow
1872.02	4.88					
1872.03	5.04					
1872.04	5.18	Yes	4.739	no	no	0
1872.05	5.18	Yes	4.788	no	no	0
1872.06	5.13	Yes	4.833	no	no	0
1872.07	5.1	Yes	4.868	no	no	0
1872.08	5.04	Yes	4.892	no	no	0
1872.09	4.95	Yes	4.904	no	no	0
1872.1	4.97	Yes	4.913	no	no	0
1872.11	4.95	Yes	4.929	no	no	0
1872.12	5.07	Yes	4.939	no	no	0
1873.01	5.11	Yes	4.955	no	no	0
1873.02	5.15	Yes	4.989	no	no	0
1873.03	5.11	Yes	5.023	no	no	0
1873.04	5.04	Yes	5.048	no	yes	5.04
1873.05	5.05	no	5.06	no	no	0
1873.06	4.98	no	5.071	no	no	0
1873.07	4.97	no	5.067	no	no	0
1873.08	4.97	no	5.053	no	no	0
1873.09	4.59	no	5.039	no	no	0

FIGURE 11-4 The moving-average trading rule beats buy-and-hold.

Assume that you first owned the stock in April 1872, so enter **Yes** in cell C24.

■ By copying the AVERAGE(B9:B23) formula from D24 to D25:D1590, you compute the 15-month moving average for each month.

> **Note** An easy way to copy the formula from D24 to D25:D1590 is to point at the lower-right corner of cell D24 (the pointer appears as a crosshair) and then double-click the left mouse button. Double-clicking copies the formula to all the cells in the column with a value in the column to the left. You can use this trick to copy formulas in multiple columns also.

■ By copying the IF(AND(C24="No",B24>D24),"yes","no") formula from E24 to E25:E1590, you determine for each month whether the S&P share is purchased during the month. Remember that you purchase the share only if you did not own the stock at the beginning of the month, and the current value of the S&P exceeds its 15-month moving average. Notice the *AND* portion of the formula. It contains two conditions (more than two are allowed) separated by a comma. If both conditions are satisfied, the formula returns Yes; otherwise, it returns No. For an IF formula to recognize text, place quotation marks (" ") around the text.

■ By copying the IF(AND(C24="Yes",B24<D24),"yes","no") formula from F24 to F25:F1590, you determine for each month whether the S&P share is sold. The stock is sold if and only if you owned the S&P share at the beginning of the month and the current value of the S&P share is below the 15-month moving average. April 1873 is the first month in which you sell your S&P stock.

■ During any month before October 2002, if you buy a share of the S&P during the month, the cash flow is negative the value of the S&P share you bought. If you sell a share of the S&P during the month, the cash flow equals the value of the S&P. Otherwise, the cash flow is 0. During October 2002, you sell any S&P share you own to get credit for its value. Therefore, by copying the IF(E24="yes",–B24,IF(F24="yes",B24,0)) formula from G24 to G25:G1589, you record the

cash flow for all months before October 2002. Entering the IF(C1590="yes",B1590,0) formula in cell G1590 gives you credit for selling any stock you own at the beginning of the past month.

- In cell G6, compute total profit from the moving-average trading strategy with the SUM(G24:G1590) formula. You can see that the 15-month moving-average strategy earns a profit of $1,319.75.

- The profit from buying and holding shares is simply the October 2002 S&P value minus the April 1872 S&P value. Compute the profit from the buy-and-hold strategy in cell G7 with the B1590–B24 formula.

As you can see, the buy-and-hold profit of $849.45 is far worse than the profit from the moving-average trading rule. Of course, this solution ignores the transaction costs incurred in buying and selling stocks. If transaction costs are large, this might wipe out the excess profits the moving-average trading strategy earns.

In the game of craps, two dice are tossed. If the total of the dice on the first roll is 2, 3, or 12, you lose. If the total of the dice on the first roll is 7 or 11, you win. Otherwise, the game keeps going. How can I write a formula to determine the status of the game after the first roll?

The fact that you lose in craps if you throw a 2, 3, or 12 can be conveniently modeled by placing an OR formula within an IF formula. In cell B5 of the Craps worksheet of the Ifstatement.xlsx file, shown in Figure 11-5, enter the IF(OR(A5=2,A5=3,A5=12),"lose",IF(OR(A5=7,A5=11),"win","keep_going")) formula. This formula is then copied from B5 to B6:B7. The formula displays lose if a 2, 3, or 12 is entered in cell A5. It displays win if a 7 or 11 is entered, and it displays keep going for any other value.

	A	B
2	Craps	
3		
4	Toss	Result
5	3	lose
6	7	win
7	9	keep going

FIGURE 11-5 Use IF statements to model the first roll in craps.

In most pro forma financial statements, cash is used as the plug to make assets and liabilities balance. I know that using debt as the plug would be more realistic. How can I set up a pro forma statement having debt as the plug?

A pro forma statement is basically a prediction of a company's financial future. A pro forma requires construction of a company's future balance sheets and income statements. The balance sheet provides a snapshot of the company's assets and liabilities at any point in time; an income statement tells you how the company's financial status is changing at any point in time. Pro forma statements can help a company determine its future needs for debt; they are also key parts of models that stock analysts use to determine whether a stock is properly valued. In the Proforma.xlsx file, the free cash flows (FCFs) were generated for a company for the next four years. Figure 11-6 shows the balance sheet, and Figure 11-7 shows the income statement.

	B	C	D	E	F	G	H
3	Sales growth	SG	0.02		CA/Sales		0.15
4	Initial sales	IS	1,000.00		CL/Sales		0.07
5	Interest rate on debt	IRD	0.10		NFA/Sales		0.60
6	Dividend payout	DIV	0.05		GFA/Sales		0.90
7	Tax rate	TR	0.53				
8	COGS/Sales	COGS	0.75				
9	Depreciation rate	DEP	0.10				
10	Liquid asset interest rate	LAIR	0.09				
11	Balance sheet						
12			0.00	1.00	2.00	3.00	4.00
13	Cash and marketable securities			0.00	0.00	0.00	52.56
14	Current assets		150.00	153.00	156.06	159.18	162.36
15	Gross fixed assets		900.00	1,001.33	1,115.29	1,243.29	1,386.93
16	Acc. dep.		300.00	400.13	511.66	635.99	774.68
17	Net fixed assets		600.00	601.20	603.62	607.30	612.24
18	Total assets		750.00	754.20	759.68	766.48	827.17
19							
20	Current liabilities		70.00	71.40	72.83	74.28	75.77
21	Debt		180.00	118.96	59.33	1.80	0.00
22	Stock		400.00	400.00	400.00	400.00	400.00
23	Retained earnings		100.00	163.84	227.52	290.39	351.40
24	Equity		500.00	563.84	627.52	690.39	751.40
25	Total liabilities		750.00	754.20	759.68	766.48	827.17

FIGURE 11-6 This is a pro forma assumptions statement and balance sheet.

	B	C	D	E	F	G	H
27	Income statement		0.00	1.00	2.00	3.00	4.00
28	Sales		1,000.00	1,020.00	1,040.40	1,061.21	1,082.43
29	Cost of goods sold		700.00	765.00	780.30	795.91	811.82
30	Depreciation			100.13	111.53	124.33	138.69
31	Operating income			154.87	148.57	140.97	131.92
32	Interest income			0.00	0.00	0.00	4.73
33	Interest expense			11.90	5.93	0.18	0.00
34	Income before taxes			142.97	142.64	140.79	136.65
35	Taxes			75.77	75.60	74.62	72.42
36	Net income			67.20	67.04	66.17	64.22
37							
38	Beg. retained earnings			100.00	163.84	227.52	290.39
39	Dividends			3.36	3.35	3.31	3.21
40	Ending retained earnings			163.84	227.52	290.39	351.40

FIGURE 11-7 This is a pro forma income statement.

Column D contains information about the company's status (during year 0). The basic assumptions are as follows:

- Sales growth (SG) is 2 percent per year.

- Initial sales are $1,000.

- Interest rate on debt is 10 percent.

- Dividend payout is 5 percent of net income.

- The tax rate is 53 percent.

- Cost of goods sold (COGS) is 75 percent of sales.

- Depreciation is 10 percent of gross fixed assets.

- Liquid assets earn 9 percent.

- Current assets are 15 percent of sales.

- Current liabilities are 7 percent of sales.

Assign the names in the cell range C3:C10 to the cells in the range D3:D10. Then, during each Year *t*, basic finance and accounting imply the following relationships, which are then implemented in a series of formulas:

- **Formula 11.1**   Year *t*+1 sales=(Year *t* sales)*(1+SG). You can compute sales during each year by copying the D28*(1+SG) formula from E28 to F28:H28.

- **Formula 11.2**   Year *t* COGS=COGS*(Year *t* sales). Each year's COGS is computed by copying the formula COGS*E28 from E29 to F29:H29.

- **Formula 11.3**   If Year *t* assets>Year *t* liabilities, Year *t* debt must be set equal to Year *t* total assets – Year *t* current liabilities – Year *t* equity. Otherwise, Year *t* debt equals 0. You can compute each year's debt in E21:H21 with the IF(E18>E20+E24,E18−E20−E24,0) formula. If Year *t* total assets are greater than Year *t* total liabilities, this formula sets Year *t* debt to Year *t* total assets – Year *t* current liabilities – Year *t* equity. This equalizes, or balances, assets and liabilities. Otherwise, you set Year *t* debt equal to 0. In this case, Year *t* cash and marketable securities are used to balance assets and liabilities.

- **Formula 11.4**   Year *t* current liabilities = (CL/Sales ratio)*(Year *t* sales). In E20:H20, use the $H$4*E28 formula to compute current liabilities for each year (copying this formula from E20 to F20:H20).

- **Formula 11.5**   Year *t* equity = Year *t* stock +Year *t* retained earnings. In E24:H24, you can compute equity by copying the SUM(E22:E23) formula from E24 to F24:H24.

- **Formula 11.6**   If Year *t* debt is greater than 0, Year *t* cash and marketable securities equals 0. Otherwise, Year *t* cash and marketable securities equals MAX(0,Year *t* total liabilities – Year *t* current assets – Year *t* net fixed assets). In E13:H13, compute cash and marketable securities for each year by copying the IF(E21>0,0,MAX(0,E25−E14−E17)) formula from E13 to F13:H13. If Year *t* debt is greater than 0, you need not use Year *t* cash and marketable securities to balance assets and liabilities. In this case, set Year *t* cash and marketable securities equal to 0. Otherwise, set Year *t* cash and marketable securities equal to Year *t* total assets – Year *t* current liabilities – Year *t* equity. This balances assets and liabilities if Year *t* assets (without cash and marketable securities) are less than Year *t* liabilities. If debt does not balance assets and liabilities, this creates liquid assets as the plug that does balance assets and liabilities.

- **Formula 11.7**   Year *t* interest expense = (Year *t* debt)*IRD. In E33, compute interest expense by using the IRD*E21 formula, copying this formula again to F33:H33.

- **Formula 11.8**   Year *t* interest income = (Year *t* cash and marketable securities)*LAIR. In E32:H32, compute interest income by copying the LAIR*E13 formula from E32 to F32:H32.

- **Formula 11.9**   Year *t* operating income = Year *t* sales – Year *t* COGS – Year *t* depreciation. In E31:H31, operating income is computed by copying the E28−E29−E30 formula from E31 to F31:H31.

- **Formula 11.10**   Year *t* dividends = (Year *t* net income)*DIV. In E39:H39, copy the E36*DIV formula from E39 to F39:H39 to compute dividends for each year.

- **Formula 11.11** Year $t + 1$ beginning retained earnings = Year t ending retained earnings. Compute beginning retained earnings each year in F38:H38, copying the E40 formula from F38 to G38:H38.

- **Formula 11.12** Year t end of year retained earnings = Year t beginning retained earnings + Year t net income – Year t dividends. In E40:H40, compute each year's ending retained earnings by copying the E38+E36–E39 formula from E40 to F40:H40.

- **Formula 11.13** Year t income before taxes = Year t operating income – Year t interest expense + Year t cash income. Compute income before taxes by copying the E31–E33+E32 formula from E34 to F34:H34.

- **Formula 11.14** Year t taxes = (Year t income before taxes)*TR. Compute each year's taxes in E35:H35 by copying the TR*E34 formula from E35 to F35:H35.

- **Formula 11.15** Year t net income after taxes = (Year t income before taxes) – (Year t taxes). In E36:H36, compute each year's net income by copying the E34–E35 formula from E36 to F36:H36.

- **Formula 11.16** Year t gross fixed assets = Year t net fixed assets + Year t accumulated depreciation. In cells E15:H15, compute gross fixed assets for each year by copying the E17+E16 formula.

- **Formula 11.17** Year t depreciation = (Year t net fixed assets)*DEP. For each year, you can use the DEP*E15 formula to compute depreciation, copying the formula from E30 to F30:H30.

- **Formula 11.18** Year t accumulated depreciation = Year $t – 1$ accumulated depreciation + Year t depreciation. For each year, use the D16+E30 formula to compute accumulated depreciation by copying the formula from E16 to F16:H16.

- **Formula 11.19** Year t net fixed assets = Year t gross fixed assets – Year t accumulated depreciation. In row 17, to compute net fixed assets, copy the E15–E16 formula from E17 to F17:H17.

- **Formula 11.20** Year t total assets = Year t liquid assets + Year t net fixed assets + Year t cash and marketable securities. By adding liquid assets, current assets, and net fixed assets, you can compute total assets by copying the SUM(E13,E14,E17) formula from E18 to F18:H18.

- **Formula 11.21** Year t total liabilities = Year t current liabilities + Year t debt + Year t equity. By copying the SUM(E20,E21,E24) formula from E25 to F25:H25, you can compute total liabilities for each period. Each year will balance because of your debt and liquid asset statements.

Formulas 11.3 and 11.6 require the use of IF statements. This worksheet will also contain circular references. (For more information about solving circular references, see Chapter 10, "Circular references.") For example, the following relationships create a circular reference:

- Year t cash affects Year t total assets.

- Year t total assets affect Year t debt.

- Year t debt affects Year t cash.

Because the worksheet contains circular references, you need to click the **File** tab, choose **Options** and then **Formulas**, and select the **Enable Iterative Calculations** check box. As explained in Chapter 10, this enables Excel to resolve circular references. Note that for each Year t, total assets in row 18 equal total liabilities in row 25. This shows the power of IF formulas combined with circular references.

When I copy a VLOOKUP formula to determine salaries of individual employees, I get a lot of #N/A errors. Then when I average the employee salaries, I cannot get a numerical answer because of the #N/A errors. Can I easily replace the #N/A errors with a blank space so I can compute average salary?

The Errortrap.xlsx file (see Figure 11-8) contains salaries and names of five employees in the D3:E7 cell range. In D11:D15 is a list of five people; compute their salary by copying the =VLOOKUP(D11,$D3:$E$7,2,False) formula from E11 to E12:E15. Unfortunately, in cells E13 and E14, Excel displays an #N/A error. NA is the acronym for "not available." Excel returns an #N/A error value when a formula cannot return an appropriate result. Because JR and Josh have no listed salary, VLOOKUP cannot return a salary value for them. Thus, when you compute average salary in E16 with the =AVERAGE(E11:E15) formula, you get an #N/A error. Many people manually replace the #N/A errors with spaces so that their average formula will calculate properly (by ignoring the spaces). The IFERROR function makes replacing errors with a character (such as a space or a 0) easy. The syntax of IFERROR is IFERROR(value,value_if_error).

The first argument is the formula you want calculated, and the second argument is the value to insert in the cell if your formula returns an error value. (Other common error values are #DIV/0, #NAME, #NUM, #REF, #VALUE; more on these later in the section.) Therefore, copying the IFERROR(VLOOKUP(D11,$D$3:$E$7,2,False)," ") formula from F11 to F12:F15 computes salary correctly for each actual employee and enters a blank space for people who are not actual employees. The =AVERAGE(F11:F15) formula now correctly computes the average salary for all listed employees.

	C	D	E	F	G
1					
2			salary		
3		Jane	40		
4		Jack	60		
5		Jill	70		
6		Erica	34		
7		Adam	120		
8					
9					
10		Name	Salary	errortrapped	
11		Erica	34	34	
12		Adam	120	120	
13		JR	#N/A		
14		Josh	#N/A		
15		Jill	70	70	
16		average	#N/A	74.66667	
17					
18	aggregate		74.66667		
19	function				
20	works!				

FIGURE 11-8 Here is an example of how error-trapping formulas trap formulas.

Excel 2010 and Excel 2013 provide a new function, the *AGGREGATE* function, that enables you to perform calculations and ignore rows that contain errors. The syntax of the AGGREGATE function is AGGREGATE(function_number,option,array). *Function number* is a code between 1 and 19 that gives

the function involved in your calculations. For example, 1 = AVERAGE, and 9 = SUM. See the Help topic on this function (by clicking **Help On This Function** in the Function Wizard) for a complete list of available functions. *Option* gives the type of cells that should be ignored in your calculations. (See Table 11-1.) Option 6 ignores error values. *Array* is simply the range of cells used in your calculations. As shown in cell E18, the AGGREGATE(1,6,E11:E15) formula correctly computes the average (74.6667) of the listed salaries.

TABLE 11-1 Types of cells that should be ignored

Option	Behavior
0 or omitted	Ignore nested *SUBTOTAL* and *AGGREGATE* functions..
1	Ignore hidden rows and nested *SUBTOTAL* and *AGGREGATE* functions.
2	Ignore error values and nested *SUBTOTAL* and *AGGREGATE* functions.
3	Ignore hidden rows, error values, and nested *SUBTOTAL* and *AGGREGATE* functions.
4	Ignore nothing.
5	Ignore hidden rows.
6	Ignore error values.
7	Ignore hidden rows and error values.

The Errortypes.xlsx file, shown in Figure 11-9, contains examples of other common error values.

FIGURE 11-9 These are examples of Excel error values.

- In cell D3, the =C3/B3 formula yields a #DIV/0! value because you can't divide by 0.

- In cell D6, the =C6+D6 formula yields a #VALUE! error because *Jack* is not the appropriate type of data for the entered formula. (*Jack* is text.)

- In cell D7, the =SUM(Sales) formula returns a #NAME? error, indicating that the Sales range name referred to in the formula is not defined.

- In cell D8, the =SQRT(–1) formula results in a #NUM! error, which results when you enter an unacceptable argument in a function. Because negative numbers do not have square roots, Excel displays the #NUM! error.

- In cell C9, entering the SUM(A1:A3) formula and then deleting column A results in a #REF! error because the cells that are referred to in the formula (cells A1:A3) are no longer in the worksheet.

You can use the *IFERROR* function to replace any of these error values by a number or text string.

My worksheet contains quarterly revenues for Wal-Mart. Can I easily compute the revenue for each year and place it in the row containing the first quarter's sales for that year?

The Walmartrev.xlsx file contains quarterly revenues (in millions of dollars) for Wal-Mart. (See Figure 11-10.) Rows 6, 10, 14, and so on contain the revenues for the first quarter of each year. In each of these rows, you will compute total revenues for the year in column E. In other rows, column E should be blank. You could enter the SUM(D6:D9) formula in cell E6 and copy this formula to E10, E14, E18, and so on, but there must be a better way. Using an IF statement with two neat Excel functions, (*ROW* and *MOD*), gives you an easy way to enter the formula once and then copy it. The *ROW(cell reference)* function yields the row of reference. The *=ROW(A6)* function yields a 6; if you are in row 6, the *=ROW()* function would also yield a 6. The *MOD(number,divisor)* function yields the remainder when *number* is divided by *divisor*. For example, MOD(9,4) yields 1, whereas MOD(6,3) yields 0. You want your formula to work only in rows that leave a remainder of 2 when divided by 4. Therefore, copying the =IF(MOD(ROW(),4)= 2,SUM(D6:D9)," ") formula from E6 to E7:E57 ensures that you add up revenues for the current year only in rows that leave a remainder of 2 when divided by 4. This means that you compute annual revenues only in the first quarter of each year, which is the goal.

	B	C	D	E
3	Wal-Mart revenues			
4				
5	Year	Quarter	revenue	Annual revenues
6	1991	1	9,281	43886.76096
7	1991	2	10,340	
8	1991	3	10,628	
9	1991	4	13,639	
10	1992	1	11,649	55483.59296
11	1992	2	13,028	
12	1992	3	13,684	
13	1992	4	17,122	
14	1993	1	13,920	67344.29593
15	1993	2	16,237	
16	1993	3	16,827	
17	1993	4	20,361	
18	1994	1	17,686	82493.89093
19	1994	2	19,942	
20	1994	3	20,418	
21	1994	4	24,448	
22	1995	1	20,440	93627
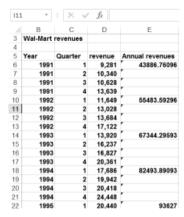
FIGURE 11-10 This is a summary of Wal-Mart, annual revenues.

IF statements can get rather large. How many IF statements can I nest in a cell? What is the maximum number of characters allowed in an Excel formula?

In Excel 2013, you can nest up to 64 IF statements in a cell. In previous versions of Excel, you could nest a maximum of seven IF statements. In Excel 2010 (and Excel 2007), a cell can contain up to 32,000 characters.

Problems

1. Suppose the price of a product will change at dates in the future as follows:

Date	Price
On or before February 15, 2004	$8
From February 16, 2004, through April 10, 2005	$9
From April 11, 2005, through January 15, 2006	$10
After January 15, 2006	$11

   Write a formula that computes the price of the product based on the date the product is sold.

2. The Blue Yonder Airlines flight from Seattle to New York has a capacity of 250 people. The airline sold 270 tickets for the flight at a price of $300 per ticket. Tickets are nonrefundable. The variable cost of flying a passenger (mostly food costs and fuel costs) is $30 per passenger. If more than 250 people show up for the flight, the flight is overbooked, and Blue Yonder must pay overbooking compensation of $350 per person to each overbooked passenger. Develop a worksheet that computes Blue Yonder's profit based on the number of customers who show up for the flight.

3. A major drug company is trying to determine the correct plant capacity for a new drug. A unit of annual capacity can be built for a cost of $10. Each unit of the drug sells for $12 and incurs variable costs of $2. The drug will be sold for 10 years. Develop a worksheet that computes the company's 10-year profit, given the chosen annual capacity level and the annual demand for the drug. Assume demand for the drug is the same each year. You can ignore the time value of money in this problem.

4. The drug company is producing a new drug. The company has made the following assumptions:

   • During Year 1, 100,000 units will be sold.

   • Sales will grow for three years and then decline for seven years.

   • During the growth period, sales will grow at a rate of 15 percent per year. During the decline, sales will drop at a rate of 10 percent per year.

   Develop a worksheet that uses the values for Year 1 sales, the length of the growth cycle, the length of the decline cycle, the growth rate during the growth cycle, and the rate of decrease during the decline cycle to compute unit sales for Years 1–11.
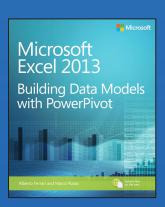
5. You are bidding on a construction project. The low bid gets the project. You estimate the project cost at $10,000. Four companies are bidding against you. It costs $400 to prepare the bid. Write a formula that (given the bids of your four competitors and your bid) computes your profit.

6. We are bidding on a valuable painting. The high bid gets the painting. We estimate the painting's value at $10,000. Four companies are bidding against us. It costs $400 to prepare the bid. Write a formula that (given the bids of our four competitors and our bid) determines whether we get the painting.

7. A drug company believes its new drug will sell 10,000 units during 2004. It expects two competitors to enter the market. The year in which the first competitor enters, the company expects to lose 30 percent of its market share. The year in which the second competitor enters, the company expects to lose 15 percent of its market share. The size of the market is growing at 10 percent per year. Given values for the years in which the two competitors enter, develop a worksheet that computes the annual sales for years 2004–2013.

8. A clothing store has ordered 100,000 swimsuits. It costs $22 to produce a swimsuit. It plans to sell them until August 31 at a price of $40 and then mark the price down to $30. Given values for demand through August 31 and after August 31, develop a worksheet to compute the profit from this order.

9. On each roll of the dice after the first roll in a game of craps, the rules are as follows: If the game has not ended and the current roll matches the first roll, you win the game. If the game has not ended and the current roll is a 7, you lose. Otherwise, the game continues. Develop a worksheet that tells you (given knowledge of the first four rolls) the status of the game after four dice rolls.

10. On the S&P moving-average example, suppose you still buy shares if the current price exceeds the 15-month moving average, but you sell if the current price is less than the 5-month moving average. Is this strategy more profitable than selling when the current price is less than the 15-month moving average?

11. A European call option gives you the right to buy a share of stock at a specified future date for a given exercise price. A *butterfly spread* involves buying one call option with a low exercise price, buying one call option with a high exercise price, and selling two call options with an exercise price midway between the low and high exercise prices. Here is an example of a butterfly spread: The current stock price is $60. You buy a $54 six-month European call option for $9, buy a $66 six-month European call option for $4, and sell two $60 European call options for the price of $6. Compute the profit (in dollars, not percentage) for this transaction as a function of six-month stock prices ranging from $40 to $80. When a trader purchases a butterfly spread, which type of movement in the stock price during the next six months is the trader betting on?

12. Suppose a stock is currently selling for $32. You buy a six-month European call option with an exercise price of $30 for $2.50 and sell a six-month European call option with an exercise price of $35 for $1. Compute the profit of this strategy (in dollars) as a function of a six-month stock price ranging from $25 to $45. Why is this strategy called a *bull spread*? How would you modify this strategy to create a *bear spread*?

13. Reconsider your pro forma example. Suppose the interest rate on your debt depends on your financial well-being. More specifically, suppose that if your earnings before interest and

taxes (EBIT) are negative, your interest rate on debt is 16 percent. If your interest expense is more than 10 percent of EBIT and EBIT is positive, your interest rate on debt is 13 percent. Otherwise, your interest rate is 10 percent. Modify the pro forma to account for this variable interest rate.

14. Do this problem independently of problem 13. Suppose your firm wants a debt-to-equity ratio of 50 percent during each year. How would you modify the pro forma? Hint: You must keep each year's stock nonnegative and use stock and cash or marketable securities to balance assets and liabilities.

15. Martin Luther King Day occurs on the third Monday in January. Write a formula that computes (given the year) the date of Martin Luther King Day. Hint: First determine the day of the week for January 1 of the given year.

16. Thanksgiving occurs on the fourth Thursday in November. Write a formula that computes (given the year) the date of Thanksgiving. Hint: First determine the day of the week for November 1 of the given year.

17. The first quarter of the year is January–March; the second quarter, April–June; the third quarter, July–September; and the fourth quarter, October–December. Write a formula that returns (for any given date) the quarter of the year.

18. Write a formula that returns a person's age, given her date of birth.

19. Labor Day is the first Monday in September. Write a formula that determines the date of Labor Day for a given year.

20. The Nancybonds.xlsx file gives the bond rating for several bonds in a previous and future month. You want to count efficiently how many bonds were downgraded. Unfortunately, each company is listed in more than one row. Assuming you have sorted the data on the company name, how would you determine the number of downgraded bonds?

21. The Addresses.xlsx file gives people's names on one line, their street address on the next line, and their city, state, and ZIP code on the following line. How could you put each person's information on one line?

22. The FormattingDDAnum.xlsx file gives a bunch of text strings, such as DDA, D, DDA1250045, and so on. A cell is properly formatted if the first three characters are DDA and the last seven characters are a number one million or larger. Determine which cells are properly formatted.

23. Suppose the number of Group 1 members is listed in cell B1, the number of Group 2 members is listed in cell B2, and the number of Group 3 members is listed in cell B3. The total number of group members is always 100. Suppose that Group 1 has 50 members, Group 2 has 30 members, and Group 3 has 20 members. Efficiently place a 1 for each Group 1 member in column D, a 2 for each Group 2 member in column D, and a 3 for each Group 3 member in column D. Thus, column D should (in this example) have a 1 in D1:D50, a 2 in D51:D80, and a 3 in D81:D100.

24. The Dividebyprice.xlsx file contains units sold of each product and total revenue. You want to determine the average price for each product. Of course, if units sold is 0, then there is no average price. Error-trap the Dividebyprice.xlsx file to ensure that all products with zero sales yield the No Sales message instead of a #DIV!0 error.

25. The School of Fine Art has 100 lockers numbered 1 to 100. At present, all lockers are open. Begin by closing every locker whose number evenly divides by 3. Then toggle (toggling means opening a closed locker and closing an open locker) each locker whose number is divisible by 4; then toggle each locker whose number is divisible by 5, . . . , and, finally, toggle each locker whose number is divisible by 100. How many lockers are now open?

26. The Matchlist.xlsx file contains a list of people who bought your product in February and a list of people who bought it in March. Determine how many of your February customers purchased your product in March.

27. Set up a calendar worksheet that takes a given month and year as inputs and tells you the day of the week on which each day of the month occurs.

28. Some rows of column C of the Problem28data.xlsx file contain the word *and*. Create formulas in column D that place an X in a row of column D if and only if column C in that row contains the word *and*.

29. Suppose you toss two 20-sided dice. The 400 possible outcomes are (1, 1), (1, 2), . . . , (20, 20). Use IF statements to generate all 400 possible outcomes systematically.

30. Using the data in the Catsanddogs.xlsx file, determine how many times the text string cat or dog occurs.

Microsoft
Excel 2013
Building Data Models
with PowerPivot

Alberto Ferrari and Marco Russo

CHAPTER 4:

Understanding data models

Your guide to quickly turn data into results.

- Understand the differences between PowerPivot for Self Service BI and SQL Server Analysis Services for Corporate BI

- Extend your existing data-analysis skills to create your own BI solutions

- Quickly manipulate large data sets, often in millions of rows

- Perform simple-to-sophisticated calculations and what-if analysis

- Create complex reporting systems with data modeling and Data Analysis Expressions

- Share your results effortlessly across your organization using Microsoft SharePoint

ABOUT THE AUTHOR

Alberto Ferrari has achieved Microsoft SQL Server Analysis Services (SSAS) Maestro status and is a consultant and trainer who specializes in developing Microsoft BI solutions. He is the coauthor of Microsoft PowerPivot for Excel 2010 and Expert Cube Development with Microsoft SQL Server 2008 Analysis Services and SQLBI Methodology.

Microsoft Press

Understanding data models

The first chapters of this book introduced some basic Microsoft PowerPivot features for generating interesting reports from existing data. In several examples, you discovered the need to model your data to make them easier to understand and manage. Nevertheless, because all the examples were introductive ones, data modeling was not an issue. It is now time to perform a deeper analysis of data modeling, so you can learn what it is and how to handle different data models.

Data modeling is not a new concept for database analysts and administrators. These technicians already know that a good data model is the foundation of a good database solution. In the world of Microsoft Excel, the same applies: a good data model is definitely the foundation of a good reporting system.

On the other hand, the problem of data modeling is somewhat new to Excel users. Until the introduction of the data model, Excel users could query only one table at a time, so the very concept of relationship was missing. The Excel data model is a sophisticated columnar database that can store huge amounts of data. It is a real database. Because you are handling a database, you need to understand how different tables are related to each other and which model is the easiest and most effective regarding your needs. In short, you need to learn the basics of data modeling.

You have already learned about the concept of relationship, which is the foundation of data modeling in the previous chapters. Nevertheless, because we believe that relationships and data models need to be well understood by the Excel analysts who are willing to create self-service business intelligence (BI) reporting, we are going to dedicate this full chapter to data modeling. At the beginning, we repeat some concepts that you have already encountered in Chapter 1, "Introduction to PowerPivot," but this time, we are trying to be more precise and follow a complete path that will help you learn how to become a good data modeler.

Understanding the basics of data modeling

You have learned in Chapter 1 that there are two ways to merge information coming from two different tables into a single, coherent view of data:

- Using single-table PivotTables, you need to use *VLOOKUP* to bring data from several tables into a single one, so that the PivotTable can work on only one table and let you analyze data.

- Using the multi-table data model, you do not need to use *VLOOKUP*. You use the creation of relationships between tables so that the PivotTable (and Power View) can analyze data coming from different tables in a single view.

The first method is easy and intuitive, but it falls short as soon as the number of tables you need to analyze grows to more than three or four. Handling and managing all the *VLOOKUP* functions end in a very complex work if *VLOOKUP* is the only available tool in your hands. On the other hand, the second method is much better, but it requires you to understand what a relationship is and how it works.

In reality, *VLOOKUP* and relationships behave much in the same way. The big difference is that *VLOOKUP* is an operation, meaning that you instruct Excel to perform a specific calculation (that is, search for a value in a table where a condition is met), whereas a relationship is a declaration (that is, you instruct Excel on how to search for values when they will be needed). But you do not ask it to perform an operation—you simply declare the existence of a relationship.

This is the big difference between using *VLOOKUP* and using relationships. When creating a relationship, you tell Excel: "Whenever you will need to merge values coming from these two tables in a computation, search for matching values between the two tables using these two columns, which need to be equal in the two tables to make a match." You can think of a relationship as a set of "virtual *VLOOKUPs*," where the lookup operation will happen only when needed for the columns requested.

Now let's learn more about the power of relationships by looking at how to handle a simple scenario with both Excel and PowerPivot. Suppose that you want to count the number of products for each category and subcategory by using a PivotTable. In the *AdventureWorks* database, you have three tables:

- **DimProduct**   Contains information about products and a column called ProductSubcategoryKey, which is the code of the subcategory of the product

- **DimProductSubcategory**   Contains all the product subcategories and, for each category, the code of the category it belongs to, in a column called ProductCategoryKey

- **DimProductCategory**   Contains information about categories

Producing a report without a data model

If you want to create a single Excel table to hold the information of the three database tables, you start by loading the three sets of data in three Excel tables similar to the ones you see in Figure 4-1, where we have removed many columns from the tables, to highlight the useful ones.

You can find this data model in the companion workbook "CH04-01-Understanding Relationships.xlsx."

FIGURE 4-1 Here, you can see the three tables containing the datasets of products, categories, and subcategories.

If you want to bring all the information from the three tables into a single one, you will need to perform these steps:

1. Create a new *VLOOKUP* operation column in DimProductSubcategory to store the ProductCategoryName taken from the DimProductCategory table by searching for the ProductCategoryKey:

```
= VLOOKUP( [@ProductCategoryKey], DimProductCategory, 3 )
```

2. Create two *VLOOKUP* operations in DimProduct to store both the product subcategory name and the product category name, taken from the DimProductSubcategory table using the ProductSubcategoryKey as the value to search for:

```
= VLOOKUP( [@ProductSubcategoryKey], DimProductSubcategory, 5 )
= VLOOKUP( [@ProductSubcategoryKey], DimProductSubcategory, 3 )
```

At the end, you created a single table containing all the columns you need, as shown in Figure 4-2.

ProductKey	ProductAlternateKey	ProductSubcategoryKey	EnglishProductName	ProductCategory	ProductSubcategory
207	TO-2301		Top Tube	#N/A	#N/A
208	TP-0923		Tension Pulley	#N/A	#N/A
209	RC-0291		Rear Derailleur Cage	#N/A	#N/A
210	FR-R92B-58	14	HL Road Frame - Black, 58	Components	Road Frames
211	FR-R92R-58	14	HL Road Frame - Red, 58	Components	Road Frames
212	HL-U509-R	31	Sport-100 Helmet, Red	Accessories	Helmets
213	HL-U509-R	31	Sport-100 Helmet, Red	Accessories	Helmets
214	HL-U509-R	31	Sport-100 Helmet, Red	Accessories	Helmets
215	HL-U509	31	Sport-100 Helmet, Black	Accessories	Helmets
216	HL-U509	31	Sport-100 Helmet, Black	Accessories	Helmets
217	HL-U509	31	Sport-100 Helmet, Black	Accessories	Helmets
218	SO-B909-M	23	Mountain Bike Socks, M	Clothing	Socks

FIGURE 4-2 The DimProduct table now contains all the necessary columns to create a PivotTable.

If you want to create a PivotTable based on this table, you then insert a new PivotTable using the DimProduct table as the source. You can put ProductCategory in the rows, but when it comes to producing the count of products, there is a problem. Producing the sum of any numeric value is as easy as dragging the column name to the VALUES area, but counting rows is somewhat more difficult. To do this, follow these instructions.

First, drag the ProductKey to the VALUES area. After that, you will obtain the sum of ProductKey, which is meaningless because there is no point in summing a key. Nevertheless, Excel uses *SUM* as the default aggregate function for any numeric column. Then, click the Field Settings button on the ANALYZE tab of the ribbon, as shown in Figure 4-3.
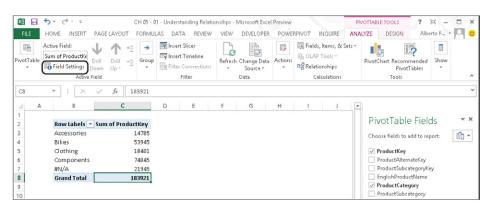
FIGURE 4-3 The Field Settings button lets you modify the way that values are shown in the PivotTable.

Excel will open the Field Settings dialog box (see Figure 4-4), where you can modify the way that values from the ProductKey column are summarized. To count products, you simply need to choose the Count function, as shown in Figure 4-4.
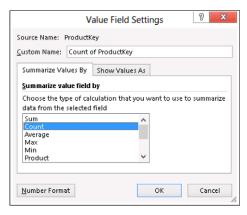
FIGURE 4-4 In the Field Settings dialog box, choose the Count function to count the number of products.

Using the *Count* function to aggregate values, the PivotTable will now show the desired report, as shown in Figure 4-5.

FIGURE 4-5 The PivotTable now shows the count of products instead of the sum of ProductKey.

Clearly, because the information about how to aggregate the ProductKey column is local to the PivotTable, if you want to create a similar report in another worksheet, you will need to repeat the operation. This is not a major issue with small reports, but it might be annoying if you create a very large workbook that covers complex analysis.

Now, let's recall briefly what we have covered so far. The relevant points of this technique are as follows:

- The product category name takes two steps to be brought into the products table. To reach it, you need to bring it first to the product subcategory table and, from there, to the products table. Another technique would have been to use a nested *VLOOKUP,* making the formula harder to read and modify later.

- To bring two columns from the product subcategory table to the product table, you have to create two columns, with two different formulas (although they are very similar), in the products table.

- If you are working with small tables, then speed and memory occupations are not an issue. But as soon as you try loading hundreds of thousands of rows, using multiple *VLOOKUP* functions will start slowing your workbook, and the table size will use up a lot of memory.

- When you need to count the products, instead of performing a *SUM,* you modify the way the PivotTable aggregates data because, by default, it always sums values. This operation needs to be executed for any PivotTable that you create based on the same table.

Building a data model

Producing the same report in PowerPivot is not more difficult; it just follows a different path. Using PowerPivot, you will first create a data model, and then create a PivotTable to produce the report. You can find this data model in the companion workbook "CH04-02-Data Model.xlsx."

To do this, first, you open the PowerPivot window and then you load three tables: DimProduct, DimProductCategory, and DimProductSubcategory. Because the relationships among the tables are already present in the database, PowerPivot creates them during the loading process. If, at this point, you switch to Diagram view, you will see the data model shown in Figure 4-6.
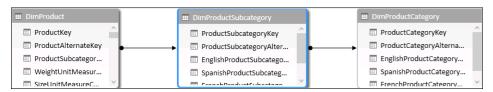
FIGURE 4-6 This data model diagram is created by PowerPivot by detecting existing relationships among tables.

You can immediately create a PivotTable based on this data, but this time, it is better to create the required calculated field in advance because we already know that we are interested in counting the number of products.

Using the measure grid, you can enter this calculated field definition:

```
NumOfProducts := COUNTROWS ( DimProduct )
```

At this point, you can simply create a new PivotTable, put the category name from the DimProductCategory table into the rows, and choose the calculated field NumOfProducts for the values. The result (see Figure 4-7) is very similar to the one obtained with the previous technique.

FIGURE 4-7 The PivotTable working on the data model looks very similar to the previous one.

What are the main differences between the two approaches?

- Using the data model approach, you did not need to use *VLOOKUP* to create columns. Rather, you declared the existence of relationships and then, when the PivotTable needed to slice products using the category name, it had to follow a relationship.

- Having a chain of relationships between tables is not an issue: even if the product category needs two steps to reach the products table (one from products to subcategories and one from subcategories to categories), the data model knows how to handle them in a simple way.

- The definition of the calculated field NumOfProducts is not local to the PivotTable. It is defined in the model and, if you need to use the same calculation in many different PivotTables based on the same data model, you do not need to create the definition in each PivotTable.

- Regarding speed and memory usage, the data model is capable of holding hundreds of millions of rows and computing very complex PivotTables in a matter of seconds.

Having reached this point, it is now time to be more precise. So, what is a data model? A data model is nothing but a group of tables tied together by a set of relationships with some calculated fields and, possibly, some calculated columns.

Thus, it is correct to say that you already know everything that you need to build a data model. Nevertheless, being able to build a data model does not mean that you know how to build a good one. The art of data modeling is complex and, as you are going to learn in the rest of this book, a good data model is the key to producing easy and maintainable DAX code. Thus, the rest of this chapter is dedicated to concepts and topics that will help you build better data models.

More about relationships

You can look at relationships via the diagram view in the PowerPivot window. Graphically, a relationship between tables is represented as an arrow starting from one table and pointing to another one. Being an arrow, a relationship has a direction. At this point, look at Figure 4-8, which shows the data model that you have built right now.
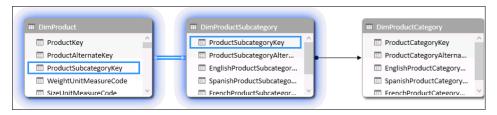
FIGURE 4-8 Relationships are shown as arrows running from one table to another.

Let's focus on the arrow between DimProduct and DimProductSubcategory:

- DimProduct is the source table (that is, the table at which the relationship originates).

- DimProductSubcategory is the target table (that is, the table that contains values related to the source table).

- The ProductSubcategoryKey column in the DimProduct table is the foreign key. It contains the value that needs to be searched in DimProductSubcategory to find the related row.

- The ProductSubcategoryKey column in the DimProductSubcategory table is the primary key. It needs to have a unique value for each row in the table. This is a requirement because, for each row in the DimProduct table, there can be only one related row in the DimProductSubcategory table. If a target table does not have unique values on the column that is used to establish the relationship, this relationship cannot be created.

Because there can be only one subcategory for each product, whereas for a single subcategory there can be many products, we sometimes refer to the tables using different names. The source table is called the "many side" of the relationship, whereas the target table is called the "one side" of the relationship. For this reason, we refer to this kind of relationship as a "one-to-many relationship." Later in this chapter, we will discuss about different kinds of relationships, like many-to-many relationships.

To better understand the direction of this relationship between DimProduct and DimProductSubcategory, you can try to delete and re-create it. To delete the relationship, switch to Diagram view, click the arrow, and hit the Del key. Excel opens a box for you to confirm the deletion. After you click Delete From Model, the relationship will be deleted.

Then, to re-create it, click the Create Relationship button on the Design tab of the PowerPivot ribbon. The Create Relationship dialog box will open and you can fill it in, as shown in Figure 4-9.
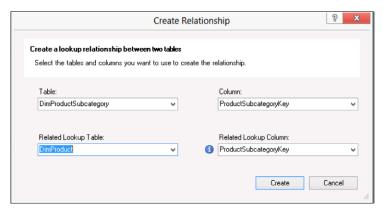
FIGURE 4-9 When you create a relationship in the wrong direction, PowerPivot automatically reverses it.

In the figure, you can see an information circle beside the Related Lookup Column. By placing the mouse on top of it, you will see this message:

The relationship cannot be created in the requested direction. When you click create, the direction of the relationship will be reversed.

PowerPivot analyzed the data in your data model and verified that ProductSubCategoryKey is a primary key in the DimProductSubcategory and is not in the DimProductTable. Thus, the relationship cannot have DimProduct as the target table. It needs to have DimProductSubcategory as the target and DimProduct as the source. For this reason, the relationship will be created in the opposite direction you have requested. This is not only a simple warning: You are going to learn that the direction of a relationship is very important because some DAX function will work only when invoked from the correct side of the relationship.

> **Note** Rare cases do exist where you have a column that is a primary key in both tables. For example, you might have a DimProduct table with several columns and a DimProductAdditionalInfo table, containing additional information for the product. In such a scenario, the ProductKey is a primary key in both columns. This relationship is known as an identity relationship because if the relationship holds, both tables are on the "one" side. Nevertheless, in the PowerPivot data model, all relationships are always one-to-many, and you will need to decide which is the one side and which is the many side. This task is not easy because DAX evaluates the formulas in different ways depending on the side of a relationship where you write the calculation. You will learn more about this topic in Chapter 7, "Understanding evaluation contexts."

Understanding normalization and denormalization

Now that you have taken a first look at what a data model is, you need to learn the difference between the physical and the logical data model. To do that, we are going to show an example of a physical data model, describe its limitations and issues, and introduce the new concept of a logical data model.

The data model containing products, categories, and subcategories is shown again in Figure 4-10.

FIGURE 4-10 This is an example of a normalized data model, each table containing a single entity.

This data model contains three tables, and each table contains a single entity. But there is a tendency to believe that a product category should not be in a table by itself, but is nothing but

an attribute of the product. So, why is it stored inside a table instead of being a column in the DimProduct table? You are facing a very common issue that happens whenever users and technicians deal with a database:

- Users think in a logical way. They divide the database into entities and concepts that are directly related to their view of the topics, which are closely related to the real world. In this example, they see products as entities and categories and subcategories as simple textual descriptions of the products. Users do not need to think of a table as containing all the categories, nor do they understand why there should be a relationship between these two tables.

- Technicians, on the other hand, think in a physical way. They divide the database into technical entities that are closely related to the physical representation of data on disk. In the technical world, it is obvious that if products have categories and subcategories, the most effective way to store this information is by creating a set of relationships among three different tables. The physical representation on disk aims at reducing data duplication to avoid inconsistencies. This technique is known as *data normalization*. Data normalization is the reason to have so many tables in a relational database. You will learn more about normalization in the next part of this text.

Thus, there are always at least two distinct ways to look at a database. We will refer to the two views as *logical* (user) versus *physical* (technician). Because the database is normally created by technicians, you can expect to have to manage normalized data models. Nevertheless, because we want to analyze it, we need a way to turn the normalized data model into a denormalized one, reducing the number of tables and increasing the model readability.

To reduce the normalization of the data model, you create two calculated columns in the DimProduct table, and then you hide the product category and subcategory tables from the model. First, create a calculated column in the DimProduct table with these DAX expressions:

```
ProductCategory    = RELATED ( DimProductCategory[EnglishProductCategoryName] )
ProductSubcategory = RELATED ( DimProductSubcategory[EnglishProductSubcategoryName] )
```

You end up with a DimProduct table that looks like the one shown in Figure 4-11.

EnglishProductName	Color	ListPrice	ProductCategory	ProductSubcategory
Long-Sleeve Logo Jerse...	Multi	$48.07	Clothing	Jerseys
Long-Sleeve Logo Jerse...	Multi	$48.07	Clothing	Jerseys
Long-Sleeve Logo Jerse...	Multi	$49.99	Clothing	Jerseys
LL Fork	NA	$148.22	Components	Forks
ML Fork	NA	$175.49	Components	Forks
HL Fork	NA	$229.49	Components	Forks
LL Mountain Handlebars	NA	$40.49	Components	Handlebars
LL Mountain Handlebars	NA	$44.54	Components	Handlebars
LL Mountain Front Wheel	Black	$60.75	Components	Wheels
ML Mountain Front Wh...	Black	$209.03	Components	Wheels
HL Mountain Front Wh...	Black	$300.22	Components	Wheels
Touring Front Wheel	Black	$218.01	Components	Wheels

FIGURE 4-11 The DimProduct table contains denormalized columns.

The column values from their tables have been brought to the DimProduct one using the *RELATED* function. At this point, there is no need to show the DimProductCategory and DimProductSubcategory tables in the PivotTables because their content is now available in the DimProduct table. For this reason, it is better to hide them from the data model. To do so, select the table in Diagram view, right-click it, and choose Hide From Client Tools, as shown in Figure 4-12.
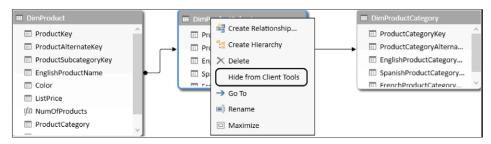
FIGURE 4-12 Technical tables should be hidden using the Hide From Client Tools function.

Performing the same operation for both DimProductCategory and DimProductSubcategory, you will reach the optimal situation, where the normalized tables, although still present in the data model, are no longer visible in the PivotTable. We refer to tables like those, which are needed for the data model but can be safely hidden, as *technical tables* (tables that exist for technical reasons but that are useless from an analytical point of view).

It is worth noting at this point that there also are technical columns (that is, columns that are needed for the data model but that have no analytical significance). For example, if you look at the ProductSubcategoryKey in the DimProductTable, it is useless to see it inside a PivotTable, even if it is necessary because it is the foreign key of the relationship with the DimProductSubcategory table. Following the same procedure, you can right-click the column and hide it from the client tools. The same can be done with many other columns as well. Hiding technical columns from the data model and removing normalization is one of the first and most important operations that improve the usability of a data model.

Recall what has been discussed so far: if you think of the possible ways to store products, categories, and subcategories on a disk file, you have these two options:

■ You can store the textual descriptions of category and subcategory inside each product, as the logical model requires. If you do this, you will end up repeating the same description for each product with a specific category.

■ On the other hand, you can store the textual description of each category only once and then reference it from subcategories using a number, which is usually shorter than the description. Moreover, you can perform the same operation with each subcategory, storing its description only once in the subcategory table and using a reference from products that links to the subcategory. This technique, in databases like Microsoft SQL Server, leads to less space being used on disk, even if it will be harder to read.

The first option is required by the logical data model. It is clear that, due to the nature of the method, the logical model contains a lot of redundancy. Because the description of a category is the same for many different products, the logical model is not optimized due to the high level of data duplication. The technical model, on the other hand, is optimized because redundancy is removed, which leads to taking up less disk size and memory and, in general, to a higher level of database performance for daily operation. In other words, the technical model is normalized (redundancy removed) and the logical one is denormalized (redundancy present).

Normalization is a process that increases the number of tables needed to store the same information, yet it reduces the overall size of the database. The need for normalization has been present in relational databases from the beginning, and database analysts are so used to normalization that they look at denormalized databases with suspicion. Nevertheless, denormalized databases have the great characteristic that they are much easier to query and understand because their structure is much more similar to the logical model.

Moreover, with the advent of modern columnar databases like PowerPivot, the need to normalize tables is greatly reduced, if not removed completely. If a single value is repeated several times in a table, PowerPivot will not store the value as is. Instead, it will provide a kind of automatic normalization by removing the long description and replacing it with numbers that just refer to it. In other words, the process of normalization is carried on internally by the PowerPivot database even if you load a denormalized table.

Is normalization the correct technique?

Note that normalization is neither good nor bad. The same applies for its opposite, denormalization. The need to normalize data comes from the very nature of the database and follows technical rules that are too complex to cover here and yet are very important. An operational database should be normalized, whereas a data model that the user can browse directly should be denormalized.

When considering logical vs. technical data models, you should avoid the temptation to think that one model is "more correct" than the other. Each one could be right or wrong, depending on the kind of operations that you need to perform on it. The point here is that, when you load data from an operational database, you will probably find normalized data. Because you will need to show that information to a user, you will need techniques to remove normalization and produce a denormalized data model, which is the correct one for the analysis.

Finally, even if denormalized data models are normally easier to query, you need to be aware that sometimes you end up over-denormalizing a data model (that is, reaching a point where denormalization starts to be a major issue) because your data model is losing its original shape. We will cover over-denormalization in the "Understanding over-denormalization" section later in this chapter because it happens very often and needs specific handling.

Denormalizing within SQL queries

In order to get acquainted with the concept of denormalization, you now will learn how to denormalize a data model using the SQL editor instead of using the *RELATED* function from inside the data model.

The *RELATED* technique is easy to use and understand, but it has the disadvantage of being annoying to implement for the many normalized columns that normally appear in an operational database. It might be interesting to learn at least one different technique to perform the same operation.

This new technique requires you to use the query builder of PowerPivot to denormalize the original data, loading information inside PowerPivot in an already denormalized way. Thus, we are shifting from denormalization inside PowerPivot (the previous technique) to denormalization outside PowerPivot (the new one).

Having more than one technique to solve the same problem will later raise the issue of deciding when to use each one, but that will come up later in this book. You can find this simple data model in the companion workbook "CH04-03-Understanding Normalization.xlsx."

The PowerPivot query designer

If you want to remove normalization from a data model, you can use the PowerPivot query designer to produce denormalized datasets. Nevertheless, before you start designing a query, it is important to understand a SQL query.

A *SQL query* is a statement written in a special language called SQL, which is understood by most modern relational databases. This book is not going to describe in detail how to write an SQL query. Here, we are interested in discussing the usage of the PowerPivot query designer, which requires you to understand just the basics of SQL and handles much of the hard work for you.

A very simple SQL query looks like this:

```
SELECT
    ProductKey,
    EnglishProductName AS ProductName,
    EnglishProductSubcategoryName AS ProductSubcategory
FROM
    dbo.DimProduct AS Prod
    INNER JOIN dbo.DimProductSubcategory AS Sub
        ON Prod.ProductSubcategoryKey = Sub.ProductSubcategoryKey
WHERE
    Sub.EnglishProductSubcategoryName = 'Helmets'
```

Let's look at this example in greater detail. A query is composed of several parts, the most common and important of which are the following:

- **SELECT**   In the first part of the query, you will declare which columns you are interested in reading. You can refer more than one table in the column list of *SELECT,* which is particularly useful when you need to get data from tables with relationships between them.

- **FROM**   In this second part of the query, you declare which tables you want the database to read to gather the columns declared in the *SELECT* section. If more than one table is referenced, as is the case in this example, then you should specify how to follow relationships through the different tables. This is done using the *JOIN* predicate.

- **WHERE**   The last part of the query lets you create filtering conditions so as not to retrieve all the rows. In the example, the code asks for rows that contain "Helmets" as the subcategory name.

Now that you have seen the overall structure of a query, you are going to create this query and load its data inside PowerPivot so you can appreciate the difference between the query result and the original tables.

Building a query is always a complex step because it requires a good knowledge of the data model to query. SQL technicians normally write queries using a standard text editor because they already know how to retrieve data. You are probably not an SQL technician, however, so you will be interested in using the PowerPivot query designer, which is a tool that lets you build an SQL query through a visual interface. Even if the query designer will handle the hard work for you, understanding the final query structure will let you understand better what the query designer is asking of you.

To open the query designer, go to the PowerPivot window and start the loading procedure to take data from a database. When choosing how to import data, instead of selecting from a list of tables, you will choose the Write A Query That Will Specify The Data To Import option, as shown in Figure 4-13.
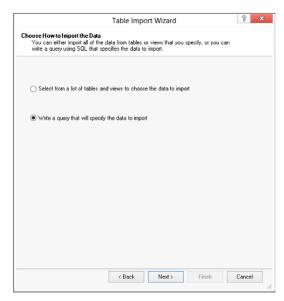
FIGURE 4-13 Choose the right option to start the query editor.

Once you click Next, the Table Import Wizard will open the SQL query editor (see Figure 4-14).
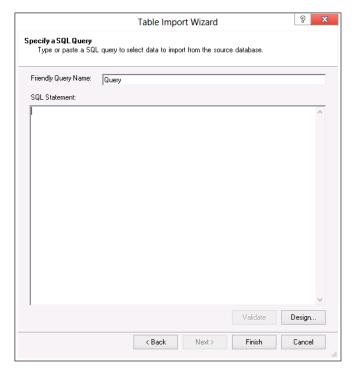
FIGURE 4-14 The SQL query editor lets you write your SQL code in a simple text box.

The query editor is a simple text editor where you can start to write the SQL query and give it a user-friendly name. Because you are probably not an SQL expert, it is much better to click Design to open the query designer.

In the query designer shown in Figure 4-15, the three tables, Product, ProductSubcategory, and ProductCategory, are already selected, and the Relationship pane has been expanded to describe it. The relationship panel is hidden by default; to expand it, you need to use the double-arrow button that appears right below the Selected Fields panel.
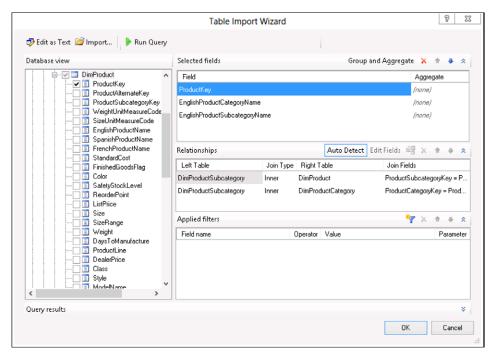
FIGURE 4-15 The query designer has a very intuitive graphical interface to build complex SQL queries.

Let's explore the query designer in more detail now. In the left pane, you can select the tables to add to the query and, expanding the table, you can select the single columns inside them. In the right part of the window, you can see the list of all the columns selected (the SELECT part of the query), the list of the tables involved, along with their relationships (the FROM part of the query) and, in the lower part, the list of filters to apply (the WHERE part of the query).

In the example, we selected a very small number of columns, to make the screenshots easy to read. In the real world, the number of columns to retrieve is normally much higher.

The less intuitive part of the query designer is the Relationship pane. You can see that, having selected three tables, the designer detected the existence of relationships between them: one between DimProduct and DimProductSubcategory and one between DimProductSubcategory and DimProductCategory. Those relationships have been automatically detected by the engine because they are stored inside the database metadata.

To understand how relationships work, you need to understand the following:

- A relationship exists between two tables only. It does not matter how many tables are involved in a query—a single relationship relates only two of them at a time.

- The two tables are ordered: one is called "Left table," and the other is "Right table." This is very important, as the next part of this text will show, because a wrong table ordering might lead to an incorrect query.

- A relationship has a type (the middle column in the relationship pane) that indicates how to retrieve rows from the database. You definitely need to understand how different types of relationships affect the query result to get the data you need.

Thus, reading the result of the auto detection algorithm, you can see that it detected both relationships and has chosen the *INNER* type of relationship for both.

In a few lines, we will analyze how different types of relationships return different results in the query designer. For the moment, though, it is enough to click OK to see the query that gets built. PowerPivot returns to the query editor, which is now filled with a complex query that the designer has written, as shown in Figure 4-16.
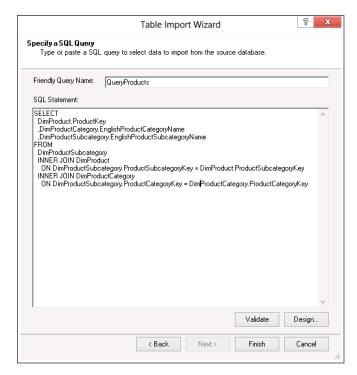
FIGURE 4-16 The SQL query editor, with a query written by the designer.

If you click Finish, you start the loading process, which means that the query is executed against the database and all the rows are imported into the PowerPivot database. Nevertheless, if you look at the number of rows returned, you may not like what you see: only 397 rows have been loaded, whereas the DimProducts table effectively contains 606 different products.

Something is wrong here why has PowerPivot lost 209 rows during the loading process? The reason is that the designer is able to understand that a relationship exists between two tables, but it is not as good in determining the type of relationship, and, having selected the wrong one, the database failed to return all the products. It is definitely time to understand better which kind of relationships you can set and how they affect the final result.

When relating the left and right tables (in this example, we will use left for DimProducts and right for DimProductSubcategory), you have four different options, listed in Table 4-1. The big difference among all the types of *JOIN* is in the handling of *NULL* values in the left or right table.

TABLE 4-1 Different types of relationships

Relationship	Result
INNER JOIN	Only the rows where there is a match between the left and right table are returned. In this example, the query will return only the products that have a subcategory, discarding the ones with the ProductSubcategoryKey column set to *NULL*.
LEFT OUTER JOIN	The left table is preserved. This means that rows coming from the left table that do not have corresponding rows in the right one will be returned anyway, leaving all the columns for the right table as *NULLs*. In the example, all products would be returned and, if they have no subcategory, the subcategory name column would be left as *NULL*.
RIGHT OUTER JOIN	The right table is preserved. In the example, all the subcategories would be returned, even the ones where there are no products inside them. Moreover, because the left table is not preserved, the products that have no subcategory will not be returned.
FULL OUTER JOIN	Both tables are preserved. In this example, you would receive all the products and all the subcategories with *NULL* values in products for categories without products, and in categories for products without categories.

If you carefully read Table 4-1, it will be clear how the error happened. The query designer chose the *INNER JOIN* type for the relationship between Products and Subcategory, thus discarding all the products that do not have a subcategory set. In other words, all the products that contained *NULL* in the ProductSubCategoryKey column will not be returned by the query because no subcategory exists with *NULL* in the key column.

To load all the products, you will need to change the relationships to *OUTER*. You can edit the query again by clicking the Table Properties button in the PowerPivot window, as shown in Figure 4-17.
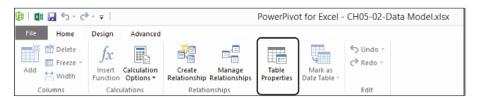
FIGURE 4-17 The Table Properties button lets you modify the properties of a table.

Using the Table Properties button, PowerPivot will open again the query editor and, from here, you can recall the query designer.

In Figure 4-18, you can see the query designer again, this time with the correct relationships set.
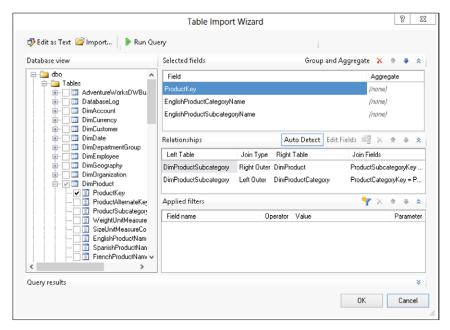
FIGURE 4-18 Here, you can see the SQL query designer, this time with the correct relationships set.

At this point, let's review the relationship definition:

■ The first relationship is between DimProductSubcategory and DimProduct. Note that DimProductSubCategory is the left table, and you want to get all the products, even the ones that contain *NULL* in the ProductSubcategoryKey column. Thus, you need a *RIGHT OUTER JOIN* between the two tables, so that the right table (DimProduct) is preserved.

■ The second relationship is between DimProductSubcategory and DimProductCategory. The DimProductSubcategory table is again the left one, but now you want all the subcategories, even if they do not have any category. Thus, here you need a *LEFT OUTER JOIN* type so that the left table (DimProductSubcategory) is preserved.

If you now save this query and refresh the table, you will get all 606 products inside your PowerPivot database, even those that do not have any category or subcategory set. This is the desired result.

> **Note** *LEFT* and *RIGHT* outer joins are somewhat interchangeable because both types of join save the rows from one table (the left or right one, depending on the type of join). If you swap the two tables, then you will need to change the join type. We find that it is better to have only one kind of join (that is, left or right, and usually, left is our preference) in a query definition because it makes it easier to follow the chain. In this example, we would prefer to swap the tables in the first relationship so that we follow a chain that starts from Products, then *LEFT OUTER JOINs* to ProductSubcategory, and finally *LEFT OUTER JOINs* to ProductCategory. Unfortunately, if you rely on the auto detection algorithm, it will choose the left and right tables individually, making your work a bit harder.

The DimProduct table imported in this way is already denormalized because it contains both the category and subcategory description, without the need to use *RELATED* and hide technical tables to simplify the data structure.

OUTER JOIN and chains of relationships

Note a subtle yet very important point about relationship: if you need to make the type of relationship between DimProduct and DimProductSubcategory a *LEFT OUTER* one, then you definitely need to make the relationship between DimProductSubcategory and DimProductCategory a *LEFT OUTER* as well. In fact, if the second join is an *INNER* one, then all the rows without a subcategory would not be returned because, having no subcategory, they do not have a category either. Whenever you are handling chained relationships, if the first one is an *OUTER JOIN,* then all the next ones need to be of the same type.

Even seasoned SQL programmers often fall into this trap and write queries incorrectly. Nevertheless, because they are used to making such mistakes, they are used to checking the query results and double-checking the chain of joins, if any. PowerPivot users will need to follow the same double-checking strategy whenever they will write a query.

You have seen that using the SQL query designer, you can write queries to the database that transform a simple table into a complex structure, which resolves the relationships at the database level and returns a denormalized table to PowerPivot. Thus, the SQL designer is a very interesting tool that you can use to perform denormalization when needed.

When to denormalize tables

Now that you have learned two distinct techniques to perform denormalization, an interesting question arises: "When do I need to denormalize tables inside queries, and when is it better to perform the same operation in PowerPivot?" Unfortunately, as with many interesting questions, there is no definitive answer. (Maybe this is why we find it interesting.)

In the example of DimProductCategory, the answer is very simple: it is much better to denormalize the column inside a query because this operation leads to a data model that is easier to query and reduces the number of tables inside PowerPivot. The rationale behind this is that the DimProductCategory table has one sole purpose: to provide a description of product categories. Therefore, when you fulfill this requirement by adding a ProductCategory column inside the Products table, the DimProductCategory table can be safely removed from the data model without changing the features of the model. The same, obviously, applies to DimProductSubcategory.

This situation (tables being needed only to provide descriptions) is amazingly common in the world of relational databases. Thus, you will encounter many cases where the application of this simple rule will drastically reduce the number of tables needed and simplify the PowerPivot data model. For simple lookup tables, the denormalization through SQL queries is by far the easiest tool.

On the other hand, when a table contains more information than a simple description, then it is much more difficult to decide whether it is better to import the table and perform the denormalization inside PowerPivot or to perform the denormalization inside a query and import the result of that query. If, for example, the DimProductSubcategory table contained more columns, then you need to add all these columns to the DimProduct table, which might result in a complex product table that is difficult to use. These situations need to be handled delicately to make sure that you make the correct decisions and produce an easy-to-use model.

Finally, complex tables that represent business entities are the pillars of data models and therefore should not be denormalized. For instance, the relationship between sales and products should not be denormalized at all because orders and products are different business entities. Both the entities and their relationships should exist in the PowerPivot database, and denormalizing it would produce an incorrect data model.

The rationale behind this approach is that you might be interested in analyzing products without relating them to sales. Think about this: if you denormalize the product name inside sales, you might want to count the number of products to produce a report. But if you see the product name only inside sales, then you will be able to count only sold products, not all the products. If a product has never been sold, then it will not even appear in the sales table because no sales row exists for it. If you reach the point where you over-denormalize the structure, you might end up creating a data model that returns incorrect results.

> **Note** One of the most frequent causes of over-denormalization is the fact that you sometimes do not have access to the original database. All you have available are prebuilt queries or reports coming from a source system. If these queries already denormalize data, they might do that for reasons that are different from yours when analyzing data. For example, you might have access to a query that returns products and sales in the same dataset, but you do not have a query for products only. Handling these scenarios is always a problem because although normalization is easy to remove, over-denormalization is much more complex.

As a rule of thumb, you will denormalize tables when the number of useful columns inside the table is very small (for lookup tables, it is often only one) and you will not denormalize tables when the number of useful column exceeds three. For two or three columns, it is up to the data modeler (that is, you) to choose the best denormalization strategy.

Keep in mind that the task of defining a data model is very complex, and denormalization is just one of the many tough decisions that need to be made. As a PowerPivot user, you are going to create data models more and more, so get ready to make lots of difficult decisions; and do not worry too much about making the wrong ones. Even seasoned BI experts often have a hard time deciding whether to keep an attribute normalized or not.

Understanding over-denormalization

Now that you have a solid background of what data modeling is and what normalization and denormalization are, it is time to delve into a complex situation that often happens when there is the need to create self-service BI reports.

Over-denormalization is simply excess denormalization (as you might assume from the name). You have already learned that denormalization works to create a user-friendly database; but over-denormalization is way too much of a good thing. This example shows why.

Figure 4-19 displays a simple data model based on the *AdventureWorks* database.
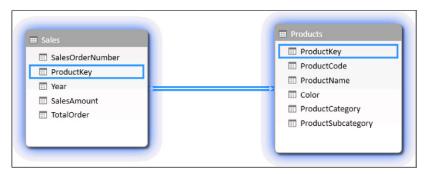
FIGURE 4-19 A very simple, yet over-denormalized data model.

Both tables have been loaded from SQL queries, and their content is as follows:

■ The Products table contains products, categories, and subcategories that already have been denormalized. Moreover, many columns have been removed from the tables, so as to focus on the most important ones.

■ The Sales table contains all the sales from the FactInternetSales table. In this table, a column has been added, TotalOrder, which contains the total of the order. Thus, SalesAmount contains the sales of the specific product (as indicated by ProductKey), whereas TotalOrder is the total sales of the order containing that specific row.

> **Note** This is an ad-hoc example, created just for this demonstration. Nevertheless, queries like this one occur very often when taking data from predefined queries used to build reports (to cite just one scenario). In fact, it is often the case that a whole dataset is stored in a single table, so that products, categories, and subcategories all belong to a single table. However, we used two tables in this example. The limitation of previous versions of Excel that required data to be stored in a single table to pivot over them has probably contributed to the spread of over-denormalization.

What is the problem with this data model? It is not in the Products table, where denormalization has been applied correctly. Rather, the problem is in the Sales table. In fact, the TotalOrder column

is not additive. If, for example, you create a simple PivotTable like the one shown in Figure 4-20, you immediately see the problem.

Row Labels	Sum of SalesAmount	Sum of TotalOrder
Black	8,838,411.96	11,736,713.49
Blue	2,279,096.28	4,496,968.50
Multi	106,470.74	2,932,191.11
NA	435,116.69	17,848,527.87
Red	7,724,330.52	9,629,604.09
Silver	5,113,389.08	5,880,293.60
White	5,106.32	412,379.28
Yellow	4,856,755.63	6,290,502.98
Grand Total	29,358,677.22	59,227,180.92

FIGURE 4-20 Over-denormalization often leads to non-additive columns.

Although the sum of sales amounts is correct, the sum of total orders is not, and you can easily check this by looking at the Grand Total row. The reason is that the total order data is stored at the order level and, thus, it is repeated for all the rows of a single order. If you want to analyze, for example, the percentage of a category or of a color against the total of the order, this data model is of no help.

Note We have spent a lot of time consulting with companies that wanted to create specific reports and failed to complete them. In our experience, we can roughly estimate that 80 percent of the time, the problem was with an incorrect data model and about 50 percent of those problems were caused by over-denormalization. Although it is possible to create very complex DAX code that handles over-denormalization, the formulas tend to become so complex that we rarely trust any scenario where over-denormalization happens. It is always better to stop, rebuild the data model using the correct normalization level, and then write simpler DAX code.

In this specific case, the correct solution is not to denormalize the orders. The correct data model contains one table for the orders, as shown in Figure 4-21.

FIGURE 4-21 Here, you can see the correct, more normalized data model.

Now that the TotalOrder in the correct table, there is a many-to-many relationship between products and orders, so you will need to know how to use many-to-many relationships (which will be discussed in Chapter 13, "Advanced DAX"). However, with the correct formulas in place, you can easily

build a report like the one shown in Figure 4-22, where the non-additive nature of the TotalOrder field is clear: the grand total is not the sum of the order total value repeated for each product on that same order (which would incorrectly overstate the actual order total).

Row Labels	Sum of SalesAmount	SumOfTotalOrder
Black	8,838,411.96	10,805,254.37
Blue	2,279,096.28	4,109,309.77
Multi	106,470.74	2,749,799.81
NA	435,116.69	9,533,098.49
Red	7,724,330.52	9,594,581.35
Silver	5,113,389.08	5,639,690.09
White	5,106.32	412,379.28
Yellow	4,856,755.63	5,621,153.54
Grand Total	29,358,677.22	29,358,677.22

FIGURE 4-22 With the correct data model, the results are as expected, and the formulas are still easy to create.

What is the lesson here? As an Excel user, you are not expected to be a professional data modeler. But because you need to work with data models, you must learn and finally master at least the basics of data modeling. If you do not, sooner or later you will encounter a data model that does not return the correct results.

Being able to quickly detect over-denormalization and find a good solution to the issue will help you most of the time, but there are very rare cases where the model will be much more complicated than this simple example. You will learn how to handle most of the complex scenarios in Chapter 13.

Understanding OLTP and data marts

The task of creating a data model is taxing. Many decisions need to be made, and it is normal to be missing some of the information you need. Creating a model is more an art than a science; it requires a lot of experience and foresight. Nevertheless, a good data model is the key to producing good analysis, and this is why we believe that a good PowerPivot analyst needs to understand data modeling.

As you have already discovered, there are at least two different ways to look at a data model:

- **Technical view**   This is a very compact and normalized way to look at data. It is the best way to model data when it needs to be used by software. Normalized data models are used by OLTP software, which is—more or less—the software that handles everyday tasks for a company. Handling such models is not your job—there are technicians who do this all day long. Nevertheless, you will probably need to read data this way at some time or another.

- **User view**   This is a highly denormalized view of data, much simpler to query and to analyze. As a side effect, however, it contains a lot of redundancy, generates bigger databases, and is very hard for software to update. Denormalized data models are normally present when the company has a data warehouse—that is, a database that contains all the data about the company in a structured way that is useful for querying.

The task of building a data warehouse usually involves several years of labor by highly specialized technicians, whose job it is to build both the data warehouse model and the software systems that fill it with the information found in the various OLTP systems around the company each day. An OLTP system is a transactional database like the one you use in your daily work to do accounting, manage sales, send invoices, and so on.

This is not a book about data warehouse modeling, so we are not going to go into a complete discussion about how this is done. Instead, we would like to give to you a sense of what a data warehouse is and the differences between querying an OLTP system and querying a data warehouse.

Querying the OLTP database

The denormalization process that you have seen up to now works well with simple tables that follow simple or chained relationships. In order to show this, we use the OLTP version of the *AdventureWorks* database, which contains the original OLTP database from which the *AdventureWorks* data warehouse is built. The *AdventureWorks* OLTP database, like any other relational database, often contains much more complex relationships, which will test your judgment.

Note This is the only time where we show data from the OLTP version of *AdventureWorks* instead of loading data from the *AdventureWorks* data warehouse, which is much simpler to use. You can download both versions of *AdventureWorks* from the links provided in Chapter 1.

As an exercise, we are going to use a diagram for the Customer table, which is implemented in *AdventureWorks* in a complex way (see Figure 4-23). The purpose of this example is to show how complex the querying of a database might be, so it is a deliberately difficult sample. You can find this model in the companion workbook "CH04-04-Querying OLTP.xlsx."

FIGURE 4-23 Here, you can see the complex relationship diagram for Customer.

You can see that the business entity Customer is represented by several tables that have complex relationships among them:

- Sales_Customer has a relationship with Sales_Individual, which stores some demographic information about the customer.

- Sales_Individual, in turn, refers to Person_Contact through the ContactID column. Person_Contact contains the personal data of any contact, and from the database point of view, a customer is a contact. Thus, Sales_Customer is in a chained relationship with Person_Contact.

- A customer row might have different addresses, so the table Sales_CustomerAddress contains all the addresses of a Sales_Customer row.

- All addresses have a related Person_AddressType, which categorizes the address. The main address of the customer has a type (a value in AddressTypeID field) called Primary.

- Other information about the address (that is, address line, ZIP code, and so on) are stored in the Person_Address table, which has a relationship with Sales_CustomerAddress, which in turn has a relationship with Sales_Customer, forming a complex relationship that looks very similar to a chained one.

- Person_Address refers to Person_StateProvince, which in turn is in relation to Person_CountryRegion. They form a chained relationship that lets us recover the name of the country through a couple of steps, from Person_Address to Person CountryRegion.

You cannot expose such complexity through PowerPivot because the PivotTable would become a nightmare. Moreover, some tables expose a many-to-many relationship (like CustomerAddress) and will not work without some complex DAX code. So you should start the analysis of this topic to understand what to do to make the model easier to query:

- The Sales_Individual table contains some demographic information. When the information is useful, it should be added to the Sales_Customer table, thus removing the Sales_Individual table from our data model.

- The Person_Contact information can be denormalized in the Sales_Customer table so that it is available directly in the Sales_Customer table.

- The Person_AddressType table is a simple descriptive one, so it can be denormalized, if you need to do that. Moreover, if (as we are supposing) we are not interested in having all the addresses of a customer, we can take only the primary address of a customer and then denormalize all the address information in the Sales_Customer table.

- Because we need to denormalize the address, we denormalize all the chained relationships, starting from the Person_Address and reaching Person_CountryRegion; we move all the information from those tables directly into the Sales_Customer table.

When the starting data model is complex, as this one is, the best (if not the only) option is to use SQL queries. Moreover, the query will be complex because it involves a high number of *JOINs,* all of

which need to be *LEFT OUTER JOIN* because you want to include all customers, whether or not they have an address stored in the database.

Here is the final result of the query:

```sql
SELECT
    Customer.CustomerID AS CustomerID,
    Customer.AccountNumber AS AccountNumber,
    Customer.CustomerType AS CustomerType,
    Contact.FirstName AS FirstName,
    Contact.LastName AS LastName,
    Contact.MiddleName AS MiddleName,
    Contact.EmailAddress AS EmailAddress,
    Contact.Phone As Phone,
    Address.AddressLine1 AS AddressLine1,
    Address.AddressLine2 AS AddressLine2,
    Address.City AS City,
    StateProvince.StateProvinceCode AS StateProvinceCode,
    StateProvince.Name As StateProvince,
    cr.CountryRegionCode AS CountryRegionCode,
    cr.Name AS CountryRegionName
FROM
    Sales.Customer AS Customer
    LEFT OUTER JOIN Sales.Individual AS Individual
        ON Individual.CustomerID = Customer.CustomerID
    LEFT OUTER JOIN Person.Contact AS Contact
        ON Contact.ContactID = Individual.ContactID
    LEFT OUTER JOIN Sales.CustomerAddress AS CustomerAddress
        ON CustomerAddress.CustomerID = Customer.CustomerID
    LEFT OUTER JOIN Person.AddressType AS T
        ON CustomerAddress.AddressTypeID = T.AddressTypeID
            AND T.Name = 'Primary'
    LEFT OUTER JOIN Person.Address AS Address
        ON Address.AddressID = CustomerAddress.AddressID
    LEFT OUTER JOIN Person.StateProvince AS StateProvince
        ON StateProvince.StateProvinceID = Address.StateProvinceID
    LEFT OUTER JOIN Person.CountryRegion AS CR
        ON CR.CountryRegionCode = StateProvince.CountryRegionCode
```

Although it is possible to write such a query using the query designer, we believe that these kinds of queries are beyond the scope of business analysts. Writing such complex SQL code should be done only by database professionals because the chances of making errors are too high.

We made this example complex on purpose. Our goal is not to scare you, or to direct you to one of the many good SQL books that are on the market. Instead, the focus here is the fact that there is a big difference between how data is stored in the source data model and how you want to handle it for your analytical purposes. Having data in the correct form makes it easier to produce interesting queries because they are easier to write using the PivotTable. Moreover, if you ever are faced with such a complex architecture, it might be wisest to ask for help from a SQL technician—you might even want to ask that person to write the query for you.

Data marts, facts, and dimensions

You have seen that querying an OLTP database could turn into a nightmare. It is now worth spending some time describing the difference between an operational database and a data warehouse, which requires us to describe how a data warehouse is composed.

The *AdventureWorks* data warehouse follows the Kimball methodology. There is not room here to provide a full description of the Kimball methodology—at this point, we just want to discuss the basic principles of this methodology to clarify some names and concepts that are present in the *AdventureWorks* database. In Kimball methodology, tables are divided into two big categories: *facts* and *dimensions*. Facts and dimensions are tied together into business units that are called, in Kimball's terminology, *data marts*. Let's quickly review what facts and dimensions are.

The core organizing principle of the data mart structure is that the database is composed of two distinct types of entities:

- **Dimension**   A *dimension* is an analytical object. It can be the list of products or customers, the time space, or any other entity used to analyze numbers.

  - Dimensions have attributes. An attribute of a product may be its color, its manufacturer, or its weight. An attribute of a date may be simply its weekday or its month name.

  - Dimensions have both natural keys and surrogate keys. The natural key is the original product code, customer ID, or real date. The surrogate key is an independent integer used in the data marts to uniquely identify a dimension entity, joining it to related facts.

  - A dimension has relationships with facts. Its purpose is to add qualitative information to the numeric measures of an event contained in facts. Sometimes a dimension might reference other dimensions, or it might correlate to other dimensions, even if its main purpose is that of joining to facts.

- **Fact**   A *fact* is something that happened or that has been measured. A fact may be the sale of a single product to a single customer or the total amount of sales of a specific item during a month. From our point of view, a fact is an event (usually represented by a number) that we want to aggregate in several forms to generate our reports.

  - We normally relate a fact to several dimensions, but we do not relate facts in any way with other facts.

  - Facts definitely have relationships with dimensions via the surrogate keys. This is one of the foundations of Kimball's methodology.

> **Note** Now that you have seen the two main entities in the data warehouse world (that is, facts and dimensions), it should be clear why the tables we have used so far from the *AdventureWorks* data warehouse always have a prefix of either *Fact* or *Dim*. *Dim* stands for dimension and *Fact,* not surprisingly, stands for fact.

Star schemas

Having seen what facts and dimensions are, you are now ready to learn about a very well-known structure called a star schema. A *star schema* is the shape taken by data warehouse diagrams. When you define dimensions and create relationships between facts and dimensions, a star schema is the natural result. The fact table is always at the center of the schema, with the dimensions related to the fact table placed around it. Because the fact table is directly related only to the dimensions, you get a shape resembling a star (see Figure 4-24, which shows a piece of the schema of the *AdventureWorks* data warehouse). You can find this model in the workbook "CH04-05-Star Schemas.xlsx."
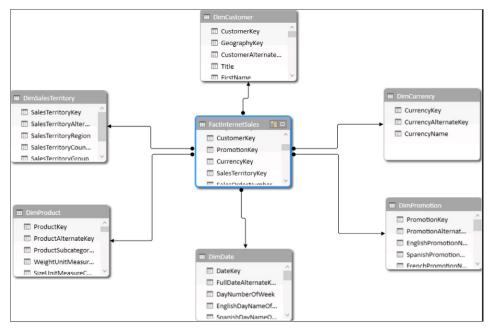
FIGURE 4-24 A star schema has a fact table at the center and dimensions around it.

Looking at this image, you can easily understand that a Customer bought a Product with a specific Currency. The sale is pertinent to a specific Territory. Star schemas are easily understandable by almost anybody at first glance.

So long as a data warehouse contains only star schemas created in accordance with the Kimball methodology, it gains several other very useful characteristics:

- There are no chained relationships. Because relationships exist only between facts and dimensions, the maximum depth of a relationship is 1. This makes it easy to generate relationships in PowerPivot because you know that a dimension is always related to a fact.

- All information is highly denormalized. This means that you never need (apart from some very specific cases) to use the *RELATED* function to load a column from a table into another one.

- If the model is well written, all descriptions in attributes are user-friendly. This means that no columns contain *NULLs* and the user can slice the data by using attributes in a very easy way.

Moreover, we figure that most of the reporting needs in the company are already fulfilled by the data warehouse. When you get your data from there, chances are better that personal reports will be comparable with the corporate ones. If, on the other side, you create your personal reports from scratch, you are in danger of producing computation that differs from that produced by the data warehouse.

Last, but not least, is the fact that by directly querying the OLTP system, we rely on its data model to remain the same over time. Although people building the data warehouse know that their data structure is queried by analysts, OLTP programmers normally do not think that someone is poking around at their data, and so they feel free to change data structure as their needs change, so long as their software works. However, if the underlying OLTP data model changes, you discover it only when someone refreshes your reports for the first time. Chances are that you would need to perform a lot of work to re-create the data model in PowerPivot.

Which database is the best to query?

Deciding whether is it better to load data from the OLTP system or the data warehouse is not easy. Using the data warehouse, you receive data already cleaned up and organized, which reduces the hard work of creating your personal data model. On the other hand, data is already organized, so you cannot model it as you would like because a lot of operations have been already carried out.

If you have a data warehouse, we suggest that you start using it. If this is enough to satisfy your reporting needs, you have completed your work with a minimum of effort. If, on the other hand, you still lack some information, you should import new data into PowerPivot tables and try to relate it with the data warehouse facts and dimensions.

Directly querying the OLTP should be seen as a very last resort because the work of creating a clean data model from an OLTP database is not easy. You could spend most of your time cleaning the model, reducing the time you might prefer to dedicate to analyzing the data.

Using advanced relationships

You have learned that the data warehouse is, by far, the best source for your data analysis. That said, even a data warehouse often hides some complexities. In this section, you are going to learn about some examples of complex relationships and how to handle them.

The first example you need to discover is multiple-column relationships. Up to now, we have said that a relationship holds between two tables and it is based on a single column, which creates the relationship when the two tables are equal. In reality, a relationship can be created between two tables using more than one column.

A good example of this is the relationship between FactInternetSales and FactInternetSalesReason in the *AdventureWorks* data warehouse, which is based on two columns. You can find this example in the companion workbook "CH04-06-Advanced Relationships.xlsx."

Note You might notice at this point that after a long digression about fact tables never being related each other, we are now introducing an example of two fact tables that share a relationship. What is happening here? Nothing special, really. The fact is that in the real world, there are many data warehouses that ignore the strict rules of Kimball methodology, and they do it for good reasons. Sometimes the real world cannot be modeled using only facts and dimensions and, in such cases, strange relationships start to appear, as the one we are showing in this section.

In *AdventureWorks,* each line of an order can have one or many sales reasons associated. For this reason, there is a table (FactInternetSalesReason) that contains a row for each reason linked to a specific sale.

In Figure 4-25, you can see the FactInternetSalesReason table, with the dimension DimSalesReason.
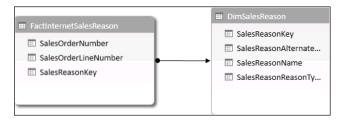
FIGURE 4-25 FactInternetSalesReason is a fact table linked to a single dimension, which is the sales reason.

FactInternetSalesReason, by itself, does not contain any interesting numbers. It contains two other columns, SalesOrderNumber and SalesOrderLineNumber, which are useful to create a relationship with the FactInternetSales table, which contains the same two columns. This time, the only issue is the fact that the relationship is based on two keys, not only one.

In PowerPivot, there is no way to create a relationship between two tables based on more than a single column. This limitation exists due to the columnar nature of the database in PowerPivot. In other databases, like SQL Server or Microsoft Access, this limitation does not exist. For this reason, you are likely to encounter several relationships based on more than one column. These relationships will not be automatically loaded inside PowerPivot.

The key to handle such a scenario is to create a single column out of the two ones present in the original tables. The easiest way to do this is to create a calculated column that uses string concatenation to merge the two columns into a single one. You can create in FactInternetSales a new calculated column with this expression:

```
InternetSalesKey =
FactInternetSalesReason[SalesOrderNumber]
& "_"
& FactInternetSalesReason [SalesOrderLineNumber]
```

The resulting table will look like Figure 4-26.

SalesOrderNumber	SalesOrderLineNumber	SalesReaso...	InternetSalesKey
SO51176	1	1	SO51176_1
SO51178	1	1	SO51178_1
SO51179	1	1	SO51179_1
SO51180	1	1	SO51180_1
SO51181	1	1	SO51181_1
SO51182	1	1	SO51182_1
SO51183	1	1	SO51183_1
SO51184	1	1	SO51184_1
SO51185	1	1	SO51185_1
SO51187	1	1	SO51187_1
SO51189	1	1	SO51189_1

FIGURE 4-26 The new InternetSalesKey is a string concatenation of the two columns used in the relationship.

You need to perform the same operation in the FactInternetSales table, using the columns with the same names. At this point, you will have a column that is a key in the FactInternetSales table and you will be able to create a relationship using that column. In Figure 4-27, you can see the resulting data model.
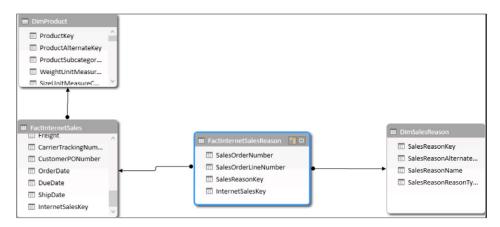
FIGURE 4-27 Here, you can see the data model, with a single column built from the two original ones.

As you have seen, even if the original data model was not perfect for the Excel data model, by means of using simple techniques and calculated columns, you have been able to make it work. You often will encounter scenarios like this one, where some modeling skills are required to design a data model that fits the needs of PowerPivot.

Nevertheless, after all this hard work, there still is a subtle problem. If you create a PivotTable and put the SalesReasonName on rows and some colors on the columns, projecting the SalesAmount on the values, you will get a PivotTable similar to the one shown in Figure 4-28.

Sum of SalesAmount	Column Labels					
Row Labels	Black	Blue	Multi	White	Yellow	Grand Total
Demo Event	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Magazine Advertisement	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Manufacturer	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
On Promotion	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Other	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Price	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Quality	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Review	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Sponsorship	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Television Advertisement	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93
Grand Total	8,838,411.96	2,279,096.28	106,470.74	5,106.32	4,856,755.63	16,085,840.93

FIGURE 4-28 The PivotTable shows incorrect results when slicing with the sales reason.

You can see that the same number is repeated for all the rows. The PivotTable correctly slice data using the color of the product, but it repeats the numbers for all the reasons.

The reason why this happens is that the relationship between DimSalesReason and FactInternet-Sales is not a one-to-many relationship, as the ones you have learned so far have been. Indeed, this is a many-to-many relationship; that is, a single sale reason can be related to many sales but, at the same time, a single line in the FactInternetSales can be related to many rows in DimSalesReasons. PowerPivot does not handle many-to-many relationships without DAX coding. Chapter 13 will cover how to work correctly with many-to-many relationships.

For now, it is enough to recognize that many-to-many relationships often appear in data models and, when they are part of the model, you will need some advanced DAX code to solve them. In reality, the code is not difficult to write—it is only hard to understand. We want to show here, in advance, the formula:

```
SalesM2M :=
CALCULATE (
    SUM ( FactInternetSales[SalesAmount] ),
    FactInternetSalesReason
)
```

Using this new calculated field, the PivotTable now shows correct results, as shown in Figure 4-29.
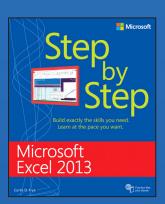
SalesM2M	Column Labels					
Row Labels	Black	Blue	Multi	White	Yellow	Grand Total
Manufacturer		199,293.05			241,430.44	440,723.49
On Promotion	2,494,836.60	545,315.31	8,254.88	665.26	977,700.45	4,026,772.50
Other	130,100.26	16,636.88	6,270.95	206.77	8,354.84	161,569.70
Price	3,899,986.11	1,328,163.92	106,470.74	5,106.32	2,438,305.38	7,778,032.47
Review	98,775.71	3,129.66			460,265.20	562,170.57
Television Advertisement	2,169.38	1,994.43			8,368.45	12,532.26
Grand Total	5,720,481.28	1,576,767.19	106,470.74	5,106.32	3,433,412.64	10,842,238.17

FIGURE 4-29 Using the many-to-many pattern, the PivotTable now shows correct results.

By looking carefully at the numbers, you will see that many-to-many relationships are non-additive, meaning that the total of the columns is not the sum of the single cells. The reason is that a single sale can be related to more than one reason, and thus, it will be accounted more than once for each row, while counting it only once for the grand total. Moreover, keep in mind that this formula uses several advanced features of DAX. We have shown the formula here because we wanted to give you the feeling of what a many-to-many relationship is and that it can be safely used in DAX.

Note Many-to-many relationships often appear due to over-denormalization of your model. In such a case, it is better not to rely on the many-to-many pattern but to remove over-denormalization from the model, creating a correct data model.

CHAPTER 4:

Getting started with Excel 2013

Experience learning made easy—and quickly teach yourself how to manage, analyze, and present data with Excel 2013. With Step by Step, you set the pace—building and practicing the skills you need, just when you them! Includes downloadable practice files and companion eBook.

Discover how to:

- Work with Excel 2013 in touch mode

- Write formulas, create macros, and organize data

- Present data visually and add images to worksheets

- Consolidate multiple sets of data into a single workbook

- Analyze data using the Quick Analysis Lens

- Collaborate with colleagues and present workbooks online
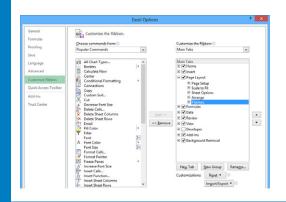
ABOUT THE AUTHOR

Curtis Frye is the author of numerous books on Excel and other Office products, most recently Microsoft Excel 2010 Step by Step. He graduated from Syracuse in 1990 with a degree in political science, and then moved to Washington, DC, where he worked as a defense trade analyst for four years and as the director of sales and marketing for an ISP for one year. He moved to Portland, Oregon, in 1995 to launch his freelance writing career. When Curt's not writing, and often while he is writing, he is a keynote speaker, mentalist, and professional improvisational comedian.

Microsoft Press

Chapter at a glance

Customize

Customize the Excel 2013 program window,
page 13
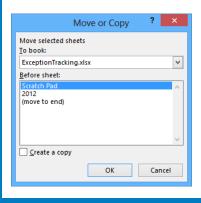
Create

Create workbooks,
page 25
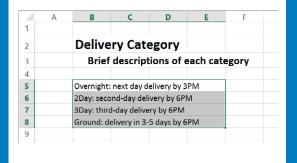
Modify

Modify workbooks,
page 30

Merge

Merge and unmerge cells,
page 39

Getting started with Excel 2013

1

IN THIS CHAPTER, YOU WILL LEARN HOW TO

- Identify the different Excel 2013 programs.

- Identify new features of Excel 2013.

- Customize the Excel 2013 program window.

- Create workbooks.

- Modify workbooks.

- Modify worksheets.

- Merge and unmerge cells.

When you create a Microsoft Excel 2013 workbook, the program presents a blank work-book that contains one worksheet. You can add or delete worksheets, hide worksheets within the workbook without deleting them, and change the order of your worksheets within the workbook. You can also copy a worksheet to another workbook or move the worksheet without leaving a copy of the worksheet in the first workbook. If you and your colleagues work with a large number of documents, you can define property values to make your workbooks easier to find when you and your colleagues attempt to locate them by searching in File Explorer or by using Windows 8 Search.

TIP In Windows 8, File Explorer has replaced Windows Explorer. Throughout this book, this browsing utility is referred to by its Windows 8 name. If your computer is running Windows 7, use Windows Explorer instead.

You can also make Excel easier to use by customizing the Excel program window to fit your work style. If you have several workbooks open at the same time, you can move between the workbook windows quickly. However, if you switch between workbooks frequently, you might find it easier to resize the workbooks so they don't take up the entire Excel window. If you do this, you can switch to the workbook that you want to modify by clicking the title bar of the workbook you want.

The Microsoft Office User Experience team has enhanced your ability to customize the Excel user interface. If you find that you use a command frequently, you can add it to the Quick Access Toolbar so it's never more than one click away. If you use a set of commands frequently, you can create a custom ribbon tab so they appear in one place. You can also hide, display, or change the order of the tabs on the ribbon.

In this chapter, you'll get an overview of the different Excel programs that are available and discover features that are available in Excel 2013. You'll also create and modify workbooks and worksheets, make workbooks easier to find, and customize the Excel 2013 program window.

PRACTICE FILES To complete the exercises in this chapter, you need the practice files contained in the Chapter01 practice file folder. For more information, see "Download the practice files" in this book's Introduction.

Identifying the different Excel 2013 programs

The Microsoft Office 2013 suite includes programs that give you the ability to create and manage every type of file you need to work effectively at home, business, or school. The programs include Microsoft Word 2013, Excel 2013, Outlook 2013, PowerPoint 2013, Access 2013, InfoPath 2013, Lync 2013, OneNote 2013, and Publisher 2013. You can purchase the programs as part of a package that includes multiple programs or purchase most of the programs individually.

With the Office 2013 programs, you can find the tools you need quickly and, because they were designed as an integrated package, you'll find that most of the skills you learn in one program transfer readily to the others. That flexibility extends well beyond your personal computer. In addition to the traditional desktop Excel program, you can also use Excel 2013 on devices with ARM chips and over the web. The following describes the different Excel 2013 programs that are available to you:

- **Microsoft Excel 2013 desktop edition** This program is installed directly on your computer. It includes all of the capabilities built into Excel 2013. You can purchase the desktop edition as part of an Office program suite, as a separate program, or as part of the Office 365 subscription package that lets you install the desktop versions of Office programs over the Internet.

TIP Office 365 is a cloud-based subscription licensing solution. Some of the Office 365 subscription levels provide access to the full Excel 2013 program, Excel Web App, or both.

- **Microsoft Excel 2013 RT** Microsoft developed an edition of Windows 8 for devices powered by an ARM processor. Devices running this edition of Windows 8, called Windows RT, come with an edition of Office 2013 named Microsoft Office 2013 RT. The Office 2013 RT program suite includes Excel, OneNote, PowerPoint, and Word.

  Excel 2013 RT takes advantage of ARM devices' touch screen capabilities by including Touch Mode. When you enable Touch Mode, the Excel 2013 RT interface changes slightly to make it easier to work with the program by tapping the screen with your finger or a stylus and by providing an on-screen keyboard through which you can enter data. You can also work with Excel 2013 RT by using a physical keyboard, a mouse, and your device's track pad.

  TIP Excel 2013 RT includes almost all of the functionality found in the Excel 2013 full desktop program; the main difference is that Excel 2013 RT does not support macros. If you open a macro-enabled workbook in Excel 2013 RT, the macros will be disabled.

- **Microsoft Excel 2013 Web App** Information workers require their data to be available to them at all times, not just when they're using their personal computers. To provide mobile workers with access to their data, Microsoft developed Office Web Apps, which include online versions of Excel, Word, PowerPoint, and OneNote. Office Web Apps are available as part of an Office 365 subscription or for free as part of the Microsoft SkyDrive cloud service.

  You can use Excel Web App to edit files stored in your SkyDrive account or on a Microsoft SharePoint site. Excel Web App displays your Excel 2010 and Excel 2013 files as they appear in the desktop program and includes all of the functions you use to summarize your data. You can also view and manipulate (but not create) PivotTables, add charts, and format your data to communicate its meaning clearly.

  Excel Web App also includes the capabilities to share your workbooks online, to embed them as part of another webpage, and to create web-accessible surveys that save user responses directly to an Excel workbook in your SkyDrive account.

  After you open a file by using Excel Web App, you can choose to continue editing the file in your browser (such as Windows Internet Explorer 10) or open the file in the desktop program. When you open the file in your desktop program, any changes you save are written to the version of the file on your SkyDrive account. This practice means that you will always have access to the most recent version of your file, regardless of where and how you access it.

- **Microsoft Excel Mobile** If you have a Windows Phone 8 device, you can use Excel Mobile to view and manipulate your workbooks. You can create formulas, change the formatting of worksheet cells, sort and filter your data, and summarize your data by using charts. You can also connect your phone to your SkyDrive account, so all of those files will be available even if you don't have a notebook or other computer to work with at the moment.

Identifying new features of Excel 2013

Excel 2013 includes all of the most useful capabilities included in previous versions of the program. If you've used an earlier version of Excel, you probably want to know about the new features introduced in Excel 2013. The following sections summarize the most important changes from Excel 2010, Excel 2007, and Excel 2003.

If you are upgrading from Excel 2010

For users of Excel 2010, you'll find that Excel 2013 extends the program's existing capabilities and adds some very useful new ones. The features introduced in Excel 2013 include:

- **Windows 8 functionality** Excel 2013, like all Office 2013 programs, takes full advantage of the capabilities of the Windows 8 operating system. When it is running on a computer running Windows 8, Excel embodies the new presentation elements and enables you to use a touch interface to interact with your data.

- **A window for each workbook** Every workbook now has its own program window.

- **New functions** More than 50 new functions are available, which you can use to summarize your data, handle errors in your formulas, and bring in data from online resources.

- **Flash Fill** If your data is in list form, you can combine, extract, or format the data in a cell. When you continue the operation, Excel detects your pattern and offers to extend it for every row in the list.

- **Quick Analysis Lens** Clicking the Quick Analysis action button, which appears next to a selected cell range, displays different ways to visually represent your data. Clicking an icon creates the analysis instantly.

- **Recommended PivotTable** PivotTables create interactive and flexible data summaries. You can have Excel recommend a series of PivotTables to create from your data, click the one you want, and keep working.

- **Recommended Charts** As with Recommended PivotTables, Excel recommends the most suitable charts based on patterns in your data. You can display the suggested charts, click the one you want, and modify it so it's perfect.

- **Chart formatting control** You can fine-tune your charts quickly and easily. Change the title, layout, or other elements of your charts from a new and interactive interface.

- **Chart animations** When you change the underlying data in a chart, Excel updates your chart and highlights the change by using an animation.

- **Cloud capability** You can now share workbooks stored online or post part of a workbook to your social network by posting a link to the file.

- **Online presentation capability** You can share your workbook and collaborate in real time with others as part of a Microsoft Lync conversation or meeting. You can also allow others to take control of your workbook during the conversation or meeting.

If you are upgrading from Excel 2007

In addition to the features added in Excel 2013, the Excel programming team introduced the following features in Excel 2010:

- **Manage Excel files and settings in the Backstage view** When the User Experience and Excel teams focused on the Excel 2007 user interface, they discovered that several workbook management tasks that contained content-related tasks were sprinkled among the ribbon tabs. The Excel team moved all of the workbook management tasks to the Backstage view, which users can access by clicking the File tab.

- **Preview data by using Paste Preview** With this feature, you can preview how your data will appear in the worksheet before you commit to the paste.

- **Customize the Excel 2010 user interface** The ability to make simple modifications to the Quick Access Toolbar has been broadened to include many more options for changing the ribbon interface. You can hide or display built-in ribbon tabs, change the order of built-in ribbon tabs, add custom groups to a ribbon tab, and create custom ribbon tabs, which can also contain custom groups.

- **Summarize data by using more accurate functions** In earlier versions of Excel, the program contained statistical, scientific, engineering, and financial functions that would return inaccurate results in some relatively rare circumstances. The Excel programming team identified the functions that returned inaccurate results and collaborated with academic and industry analysts to improve the functions' accuracy.

- **Summarize data by using sparklines** In his book *Beautiful Evidence* (Graphics Press), Edward Tufte describes sparklines as "intense, simple, wordlike graphics." Sparklines take the form of small charts that summarize data in a single cell. These small but powerful additions to Excel 2010 and Excel 2013 enhance the program's reporting and summary capabilities.

- **Filter PivotTable data by using slicers** Slicers visually indicate which values appear in a PivotTable and which are hidden. They are particularly useful when you are presenting data to an audience that contains visual thinkers who might not be skilled at working with numerical values.

- **Filter PivotTable data by using search filters** Excel 2007 introduced several new ways to filter PivotTables. These filtering capabilities have been extended with the introduction of search filters. With a search filter, you begin entering a sequence of characters that occur in the term (or terms) by which you want to filter. As you enter these characters, the filter list of the PivotTable field displays only those terms that reflect the values entered into the search filter box.

- **Visualize data by using improved conditional formats** The Excel programming team greatly extended the capabilities of the data bar and icon set conditional formats introduced in Excel 2007. The team also enabled you to create conditional formats that refer to cells on worksheets other than the one on which you're defining the format.

- **Create and display math equations** With the updated equation designer, you can create any equation you require. The editor has several common equations built in, such as the quadratic formula and the Pythagorean theorem, but it also contains numerous templates that you can use to create custom equations quickly.

- **Edit pictures within Excel 2010** One very helpful capability is the ability to remove the background elements of an image. Removing an image's background enables you to create a composite image in which the foreground elements are placed in front of another background. For example, you can focus on a flower's bloom and remove most of the leaves and stem from the photo. After you isolate the foreground image, you can place the bloom in front of another background.

If you are upgrading from Excel 2003

In addition to the changes in Excel 2010 and Excel 2013, users upgrading from Excel 2003 will notice several more significant changes:

- **The ribbon** Unlike in previous versions of Excel, in which you hunted through a complex toolbar and menu system to find the commands you wanted, you can use the ribbon user interface to find everything you need at the top of the program window.

- **Larger data collection capability** The larger worksheet includes more than 1 million rows and 16,000 columns.

- **New file format** The Excel file format (.xlsx) uses XML and file compression techniques to reduce the size of a typical file by 50 percent.

- **Expanded cell and worksheet formatting** Vast improvements have been made to the color management and formatting options found in previous versions of the program. You can have as many different colors in a workbook as you like, for example, and you can assign a design theme to a workbook.

- **Excel tables** These enable you to enter and summarize your data efficiently. If you want to enter data in a new table row, all you have to do is enter the data in the row below the table. When you press Tab or Enter after entering the last cell's values, Excel expands the table to include your new data. You can also have Excel display a Totals row, which summarizes your table's data by using a function that you specify.

- **Improved charting** With the charting engine, you can create more attractive charts.

- **Formula AutoComplete** When you enter formulas into an Excel worksheet cell, the program displays a list of options from which you can choose for each formula element, greatly accelerating formula entry.

- **Additional formulas** With the added formulas, such as AVERAGEIFS, users can summarize data conditionally.

- **Conditional formatting** With conditional formats, users can create data bars and color scales, assign icon sets to values, assign multiple conditional formats to a cell, and assign more than three conditional formatting rules to a cell.

Working with the ribbon

As with all Office 2013 programs, the Excel ribbon is dynamic, meaning that as its width changes, its buttons adapt to the available space. As a result, a button might be large or small, it might or might not have a label, or it might even be an entry in a list.

For example, when sufficient horizontal space is available, the buttons on the Home tab are spread out, and the available commands in each group are visible.

If you decrease the horizontal space available to the ribbon, small button labels disappear and entire groups of buttons might hide under one button that represents the entire group. Clicking the group button displays a list of the commands available in that group.
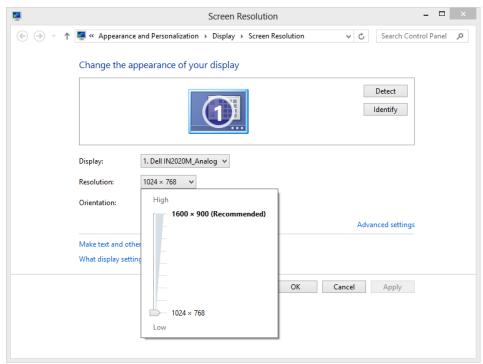
When the ribbon becomes too narrow to display all the groups, a scroll arrow appears at its right end. Clicking the scroll arrow displays the hidden groups.

The width of the ribbon depends on three factors:

- **Program window width** Maximizing the program window provides the most space for the ribbon. To maximize the window, click the **Maximize** button, drag the borders of a nonmaximized window, or drag the window to the top of the screen.
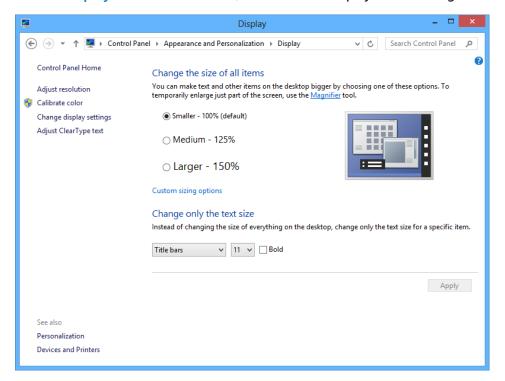
- **Screen resolution** Screen resolution is the size of your screen display expressed as pixels wide × pixels high. Your screen resolution options are dependent on the display adapter installed in your computer, and on your monitor. Common screen resolutions range from 800 × 600 to 2560 × 1600. The greater the number of pixels wide (the first number), the greater the number of buttons that can be shown on the ribbon.

To change your screen resolution:

1 Display the **Screen Resolution** control panel item in one of the following ways:

   - Right-click the Windows desktop, and then click **Screen Resolution**.

   - Enter screen resolution in Windows 8 Search, and then click **Adjust screen resolution** in the **Settings** results.

   - Open the **Display** control panel item, and then click **Adjust resolution**.

2 On the **Screen Resolution** page, click the **Resolution** arrow, click or drag to select the screen resolution you want, and then click **Apply** or **OK**.

- **The magnification of your screen display** If you change the screen magnification setting in Windows, text and user interface elements are larger and therefore more legible, but fewer elements fit on the screen. You can set the magnification from 100 to 500 percent.

You can change the screen magnification from the Display page of the Appearance And Personalization control panel item. You can display the Display page directly from Control Panel or by using one of the following methods:

- Right-click the Windows desktop, click **Personalize**, and then in the lower-left corner of the **Personalization** window, click **Display**.

- Enter display in Windows 8 Search, and then click **Display** in the **Settings** results.

To change the screen magnification to 125 percent or 150 percent, click that option on the Display page. To select another magnification, click the Custom Sizing Options link and then, in the Custom Sizing Options dialog box, click the magnification you want in the drop-down list or drag the ruler to change the magnification even more.

After you click OK in the Custom Sizing Options dialog box, the custom magnification is shown on the Display page along with any warnings about possible problems with selecting that magnification. Click Apply on the Display page to apply the selected magnification.

Customizing the Excel 2013 program window

How you use Excel 2013 depends on your personal working style and the type of data collections you manage. The Excel product team interviews customers, observes how differing organizations use the program, and sets up the user interface so that many users won't need to change it to work effectively. If you do want to change the program window, including the user interface, you can. You can change how Excel displays your worksheets; zoom in on worksheet data; add frequently used commands to the Quick Access Toolbar; hide, display, and reorder ribbon tabs; and create custom tabs to make groups of commands readily accessible.

Zooming in on a worksheet

One way to make Excel easier to work with is to change the program's zoom level. Just as you can "zoom in" with a camera to increase the size of an object in the camera's viewer, you can use the zoom setting to change the size of objects within the Excel program window. For example, if Peter Villadsen, the Consolidated Messenger European Distribution Center Manager, displayed a worksheet that summarized his distribution center's package volume by month, he could click the View tab and then, in the Zoom group, click the Zoom button to open the Zoom dialog box. The Zoom dialog box contains controls that he can use to select a preset magnification level or to enter a custom magnification level. He could also use the Zoom control in the lower-right corner of the Excel window.

Clicking the Zoom In control increases the size of items in the program window by 10 percent, whereas clicking the Zoom Out control decreases the size of items in the program window by 10 percent. If you want more fine-grained control of your zoom level, you can use the slider control to select a specific zoom level or click the magnification level indicator, which indicates the zoom percentage, and use the Zoom dialog box to set a custom magnification level.
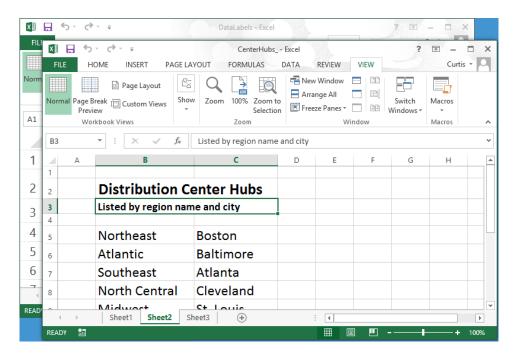
The Zoom group on the View tab contains the Zoom To Selection button, which fills the program window with the contents of any selected cells, up to the program's maximum zoom level of 400 percent.

TIP The minimum zoom level in Excel is 10 percent.

Arranging multiple workbook windows

As you work with Excel, you will probably need to have more than one workbook open at a time. For example, you could open a workbook that contains customer contact information and copy it into another workbook to be used as the source data for a mass mailing you create in Word. When you have multiple workbooks open simultaneously, you can switch between them by clicking the View tab and then, in the Window group, clicking the Switch Windows button and clicking the name of the workbook you want to view.
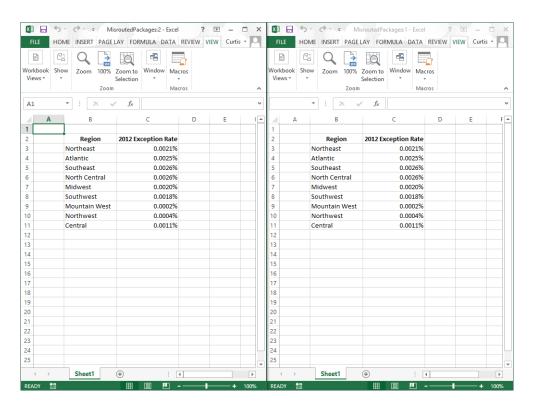
You can arrange your workbooks on the desktop so that most of the active workbook is shown but the others are easily accessible. To do so, click the View tab and then, in the Window group, click the Arrange All button. Then, in the Arrange Windows dialog box, click Cascade.

Many Excel workbooks contain formulas on one worksheet that derive their value from data on another worksheet, which means you need to change between two worksheets every time you want to test how modifying your data changes the formula's result. However, an easier way to approach this is to display two copies of the same workbook simultaneously,

displaying the worksheet that contains the data in the original window and displaying the worksheet with the formula in the new window. When you change the data in either copy of the workbook, Excel updates the other copy. To display two copies of the same workbook, open the workbook and then, on the View tab, in the Window group, click New Window to opens a second copy of the workbook. To display the workbooks side by side, on the View tab, click Arrange All. Then, in the Arrange Windows dialog box, click Vertical and then click OK.

If the original workbook's name is MisroutedPackages, Excel displays the name MisroutedPackages:1 on the original workbook's title bar and MisroutedPackages:2 on the second workbook's title bar.

TROUBLESHOOTING If the controls in the Window group on the View tab don't affect your workbooks as you expect, you might have a program, such as SkyDrive for PC, open in the background that prevents those capabilities from functioning.

Adapting exercise steps

The screen shots shown in this book were captured at a screen resolution of 1024 × 768, at 100-percent magnification. If your settings are different, the ribbon on your screen might not look the same as the one shown in this book. As a result, exercise instructions that involve the ribbon might require a little adaptation. This book's instructions use this format:

- On the **Insert** tab, in the **Illustrations** group, click the **Chart** button.

If the command is in a list, the instructions use this format:

- On the **Home** tab, in the **Editing** group, click the **Find** arrow and then, in the **Find** list, click **Go To**.

If your display settings cause a button to appear differently on your screen than it does in this book, you can easily adapt the steps to locate the command. First click the specified tab, and then locate the specified group. If a group has been collapsed into a group list or under a group button, click the list or button to display the group's commands. If you can't immediately identify the button you want, point to likely candidates to display their names in ScreenTips.

This book provides instructions based on traditional keyboard and mouse input methods. If you're using Excel on a touch-enabled device, you might be giving commands by tapping with your finger or with a stylus. If so, substitute a tapping action any time the instructions ask you to click a user interface element. Also note that when the instructions ask you to enter information in Excel, you can do so by typing on a keyboard, tapping in the entry field under discussion to display and use the on-screen keyboard, or even speaking aloud, depending on your computer setup and your personal preferences.

In this exercise, you'll change a worksheet's zoom level, zoom to maximize the display of a selected cell range, switch between workbooks, and arrange all open workbooks on your screen.
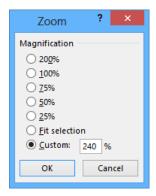
➡ SET UP **You need the PackageCounts and MisroutedPackages workbooks located in the Chapter01 practice file folder to complete this exercise. Open both workbooks, and then follow the steps.**

1 In the **MisroutedPackages** workbook, in the lower-right corner of the Excel window, click the **Zoom In** control five times to change the worksheet's zoom level to **150%**.

2   Select cells **B2:C11**.

3   On the **View** tab, in the **Zoom** group, click the **Zoom to Selection** button to display the selected cells so that they fill the program window.
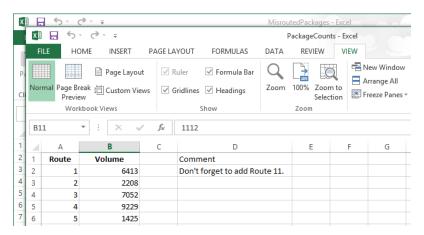
◢	A	B	C
2		**Region**	**2012 Exception Rate**
3		Northeast	0.0021%
4		Atlantic	0.0025%
5		Southeast	0.0026%
6		North Central	0.0026%
7		Midwest	0.0020%
8		Southwest	0.0018%
9		Mountain West	0.0002%
10		Northwest	0.0004%
11		Central	0.0011%

Sheet1   ⊕

4   On the **View** tab, in the **Zoom** group, click the **Zoom** button to open the **Zoom** dialog box.

```
Zoom                    ?  ×
┌─ Magnification ──────────────┐
│  ○ 200%                       │
│  ○ 100%                       │
│  ○ 75%                        │
│  ○ 50%                        │
│  ○ 25%                        │
│  ○ Fit selection              │
│  ● Custom:   240  %           │
│                               │
│    [ OK ]      [ Cancel ]     │
└───────────────────────────────┘
```

5   Click **100%**, and then click **OK** to return the worksheet to its default zoom level.

6   On the **View** tab, in the **Window** group, click the **Switch Windows** button, and then click **PackageCounts** to display the **PackageCounts** workbook.

7   On the **View** tab, in the **Window** group, click the **Arrange All** button to open the **Arrange Windows** dialog box.

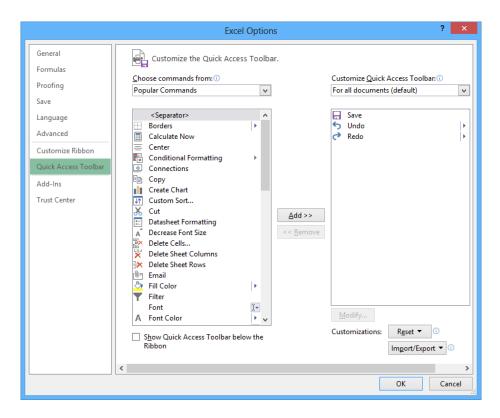8   Click **Cascade**, and then click **OK** to cascade the open workbook windows.

❌ CLEAN UP **Close the PackageCounts and MisroutedPackages workbooks, saving your changes if you want to.**

Adding buttons to the Quick Access Toolbar

As you continue to work with Excel 2013, you might discover that you use certain commands much more frequently than others. If your workbooks draw data from external sources, for example, you might find yourself displaying the Data tab and then, in the Connections group, clicking the Refresh All button much more often than the program's designers might have expected. You can make any button accessible with one click by adding the button to the Quick Access Toolbar, located just above the ribbon in the upper-left corner of the Excel program window.

To add a button to the Quick Access Toolbar, click the File tab to display the Backstage view, and then click Options in the left pane. In the Excel Options dialog box, display the Customize The Quick Access Toolbar page. This page contains two panes. The pane on the left lists all of the controls that are available within a specified category, and the pane on the right lists the controls currently displayed on the Quick Access Toolbar. To add a command to the Quick Access Toolbar, in the Choose Commands From list, click the category that contains the control you want to add. Excel displays the available commands in the

pane below the Choose Commands From field. Click the control you want, and then click the Add button.

You can change a button's position on the Quick Access Toolbar by clicking its name in the right pane and then clicking either the Move Up or Move Down button at the right edge of the dialog box. To remove a button from the Quick Access Toolbar, click the button's name in the right pane, and then click the Remove button. When you're done making your changes, click the OK button. If you prefer not to save your changes, click the Cancel button. If you saved your changes but want to return the Quick Access Toolbar to its original state, click the Reset button and then click either Reset Only Quick Access Toolbar, which removes any changes you made to the Quick Access Toolbar, or Reset All Customizations, which returns the entire ribbon interface to its original state.
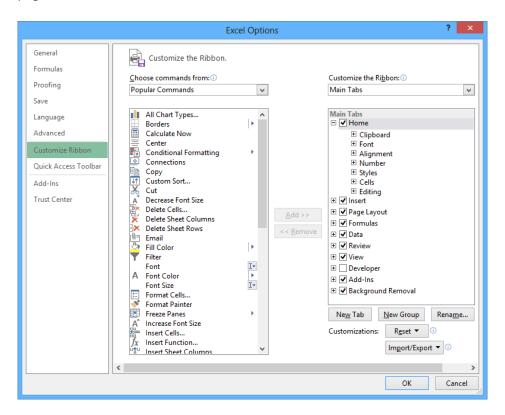
You can also choose whether your Quick Access Toolbar changes affect all your workbooks or just the active workbook. To control how Excel applies your change, in the Customize Quick Access Toolbar list, click either For All Documents to apply the change to all of your workbooks or For Workbook to apply the change to the active workbook only.

If you'd like to export your Quick Access Toolbar customizations to a file that can be used to apply those changes to another Excel 2013 installation, click the Import/Export button and then click Export All Customizations. Use the controls in the dialog box that opens to save your file. When you're ready to apply saved customizations to Excel, click the Import/Export button, click Import Customization File, select the file in the File Open dialog box, and click Open.

Customizing the ribbon

Excel enhances your ability to customize the entire ribbon by enabling you to hide and display ribbon tabs, reorder tabs displayed on the ribbon, customize existing tabs (including tool tabs, which appear when specific items are selected), and create custom tabs.

To begin customizing the ribbon, display the Backstage view and then click Options. In the Excel Options dialog box, click Customize Ribbon to display the Customize The Ribbon page.
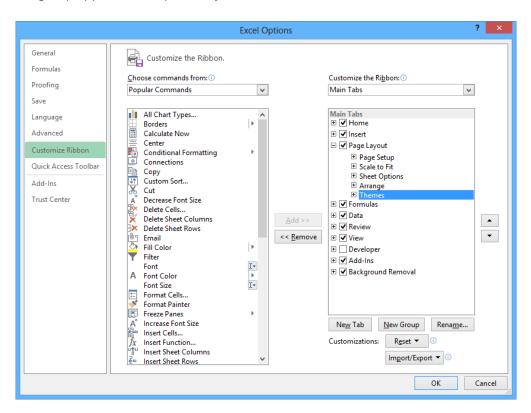
To select which tabs appear in the tabs pane on the right side of the screen, click the Customize The Ribbon field's arrow and then click either Main Tabs, which displays the tabs that can appear on the standard ribbon; Tool Tabs, which displays the tabs that appear when you click an item such as a drawing object or PivotTable; or All Tabs.

TIP The procedures taught in this section apply to both the main tabs and the tool tabs.

Each tab's name has a check box next to it. If a tab's check box is selected, then that tab appears on the ribbon. You can hide a tab by clearing the check box and bring the tab back by selecting the check box. You can also change the order in which the tabs are displayed on the ribbon. To do so, click the name of the tab you want to move and then click the Move Up or Move Down arrow to reposition the selected tab.

Just as you can change the order of the tabs on the ribbon, you can change the order in which groups of commands appear on a tab. For example, the Page Layout tab contains five groups: Themes, Page Setup, Scale To Fit, Sheet Options, and Arrange. If you use the Themes group less frequently than the other groups, you could move the group to the right end of the tab by clicking the group's name and then clicking the Move Down button until the group appears in the position you want.

To remove a group from a built-in tab, click the name of the group in the right pane and click the Remove button. If you remove a group from a built-in tab and later decide you want to put it back on the tab, display the tab in the right pane. Then, click the Choose Commands From field's arrow and click Main Tabs. With the tab displayed, in the left pane, click the expand control (which looks like a plus sign) next to the name of the tab that contains the group you want to add back. You can now click the name of the group in the left pane and click the Add button to put the group back on the selected tab.

The built-in tabs are designed efficiently, so adding new command groups might crowd the other items on the tab and make those controls harder to find. Rather than adding controls to an existing tab, you can create a custom tab and then add groups and commands to it. To create a custom tab, click the New Tab button on the Customize The Ribbon page of the Excel Options dialog box. When you do, a new tab named New Tab (Custom), which contains a group named New Group (Custom), appears in the tab list.

You can add an existing group to your new tab by clicking the Choose Commands From field's arrow, selecting a collection of commands, clicking the group you want to add, and then clicking the Add button. You can also add individual commands to your tab by clicking a command in the command list and clicking the Add button. To add a command to your tab's custom group, click the new group in the right tab list, click the command in the left list, and then click the Add button. If you want to add another custom group to your new tab, click the new tab, or any of the groups within that tab, and then click New Group.

TIP You can change the order of the groups and commands on your custom ribbon tabs by using the techniques described earlier in this section.

The New Tab (Custom) name doesn't tell you anything about the commands on your new tab, so you can rename it to reflect its contents. To rename any tab on the ribbon, display the Customize The Ribbon page of the Excel Options dialog box, click the tab you want to modify, and then click the Rename button. Enter the tab's new name in the Rename dialog box, and click OK. To rename any group on the ribbon, click the name of the group, and then click Rename. When you do, the Rename dialog box appears. Enter a new name for the group in the Display Name box and click OK.

If you'd like to export your ribbon customizations to a file that can be used to apply those changes to another Excel 2013 installation, click the Import/Export button and then click Export All Customizations. Use the controls in the dialog box that opens to save your file. When you're ready to apply saved customizations to Excel, click the Import/Export button, click Import Customization File, select the file in the File Open dialog box, and click Open.

When you're done customizing the ribbon, click the OK button to save your changes or click Cancel to keep the user interface as it was before you started this round of changes. You can also change a tab, or the entire ribbon, back to the state it was in when you installed Excel. To restore a single tab, click the tab you want to restore, click the Reset button, and then click Reset Only Selected Ribbon Tab. To restore the entire ribbon, including the Quick Access Toolbar, click the Reset button and then click Reset All Customizations.

Maximizing usable space in the program window

You can increase the amount of space available inside the program window by hiding the ribbon, the formula bar, or the row and column labels.

To hide the ribbon, double-click the active tab label. The tab labels remain visible at the top of the program window, but the tab content is hidden. To temporarily redisplay the ribbon, click the tab label you want. Then click any button on the tab, or click away from the tab, to rehide it. To permanently redisplay the ribbon, double-click any tab label.

KEYBOARD SHORTCUT Press Ctrl+F1 to hide and unhide the ribbon. For a complete list of keyboard shortcuts, see "Keyboard shortcuts" at the end of this book.
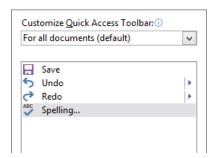
To hide the formula bar, clear the Formula Bar check box in the Show/Hide group on the View tab. To hide the row and column labels, clear the Headings check box in the Show/Hide group on the View tab.

In this exercise, you'll add a button to the Quick Access Toolbar and customize the ribbon.

→ SET UP **You need the PackageCounts workbook located in the Chapter01 practice file folder to complete this exercise. Open the workbook, and then follow the steps.**

1   Click the **File** tab to display the **Backstage** view, and then click **Options** to open the **Excel Options** dialog box.

2   Click **Quick Access Toolbar** to display the **Customize The Quick Access Toolbar** page.

3   Click the **Choose commands from** arrow, and then in the list, click **Review Tab** to display the commands in the **Review Tab** category in the command list.

4   Click the **Spelling** command, and then click **Add** to add the **Spelling** command to the **Quick Access Toolbar**.

5   Click **Customize Ribbon** to display the **Customize The Ribbon** page of the **Excel Options** dialog box.

6   If necessary, click the **Customize the Ribbon** box's arrow and click **Main Tabs**. In the right tab list, click the **Review** tab and then click the **Move Up** button three times to move the **Review** tab between the **Insert** and **Page Layout** tabs.

7   Click the **New Tab** button to create a tab named **New Tab (Custom)**, which appears below the most recently active tab in the **Main Tabs** list.

8   Click the **New Tab (Custom)** tab name, click the **Rename** button, enter **My Commands** in the **Display Name** box, and click **OK** to change the new tab's name to **My Commands**.

9   Click the **New Group (Custom)** group's name and then click the **Rename** button. In the **Rename** dialog box, click the icon that looks like a paint palette (second row, fourth from the right). Then, in the **Display name** box, enter **Formatting**, and click **OK** to change the new group's name to **Formatting**.

10   In the right tab list, click the **My Commands** tab name. Then, on the left side of the dialog box, click the **Choose Commands From** box's arrow and click **Main Tabs** to display that group of tabs in the left tab list.

11   In the left tab list, click the **Home** tab's expand control, click the **Styles** group's name, and then click the **Add** button to add the **Styles** group to the **My Commands** tab.

12   In the left tab list, below the **Home** tab, click the **Number** group's expand control to display the commands in the **Number** group.

13   In the right tab list, click the **Formatting** group you created earlier. Then, in the left tab list, click the **Number Format** item and click the **Add** button to add the **Number Format** item to the **Formatting** custom group.

14   Click **OK** to save your ribbon customizations, and then click the **My Commands** tab on the ribbon to display the contents of the new tab.
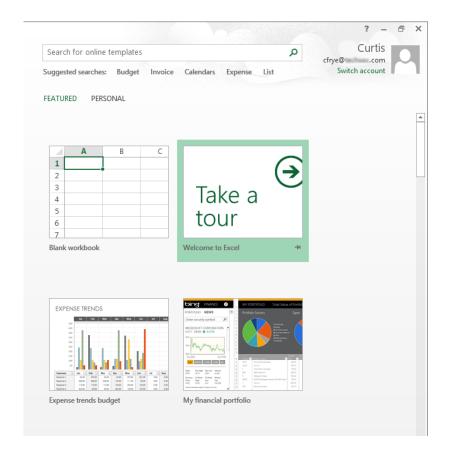

IMPORTANT The remaining exercises in this book assume that you are using Excel 2013 as it was installed on your computer. After you complete this exercise, you should reset the ribbon to its original configuration so that the instructions in the remaining exercises in the book are consistent with your copy of Excel.

✖ CLEAN UP **Close all open workbooks, saving your changes if you want to.**

Creating workbooks

Every time you want to gather and store data that isn't closely related to any of your other existing data, you should create a new workbook. The default workbook in Excel has one worksheet, although you can add more worksheets or delete existing worksheets if you want. Creating a workbook is a straightforward process—you just display the Backstage view, click New, and click the tile that represents the type of workbook you want.

KEYBOARD SHORTCUT Press Ctrl+N to create a blank workbook.

When you start Excel, the program displays the Start experience. With the Start experience, you can select which type of workbook to create. You can create a blank workbook by clicking the Blank Workbook tile or click one of the built-in templates available in Excel. You can then begin to enter data into the worksheet's cells. You could also open an existing workbook and work with its contents. In this book's exercises, you'll work with workbooks created for Consolidated Messenger, a fictional global shipping company. After you make changes to a workbook, you can save it to preserve your work.

KEYBOARD SHORTCUT Press Ctrl+S to save a workbook.

TIP Readers frequently ask, "How often should I save my files?" It is good practice to save your changes every half hour or even every five minutes, but the best time to save a file is whenever you make a change that you would hate to have to make again.
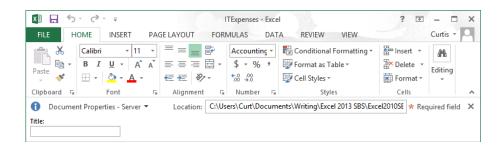
When you save a file, you overwrite the previous copy of the file. If you have made changes that you want to save, but you also want to keep a copy of the file as it was when you saved it previously, you can use the Save As command to specify a name for the new file. To open the Save As dialog box, in the Backstage view, click Save As.

KEYBOARD SHORTCUT Press F12 to open the Save As dialog box.

You can also use the controls in the Save As dialog box to specify a different format for the new file and a different location in which to save the new version of the file. For example, Lori Penor, the chief operating officer of Consolidated Messenger, might want to save an Excel file that tracks consulting expenses as an Excel 2003 file if she needs to share the file with a consulting firm that uses Excel 2003.

After you create a file, you can add information to make the file easier to find when you search by using File Explorer or Windows 8 Search to search for it. Each category of information, or property, stores specific information about your file. In Windows, you can search for files based on the file's author or title, or by keywords associated with the file. A file that tracks the postal code destinations of all packages sent from a vendor might have the keywords *postal*, *destination*, and *origin* associated with it.

To set values for your workbook's built-in properties, you can display the Backstage view, click Info, click Properties, and then click Show Document Panel to display the Document Properties panel below the ribbon. The standard version of the Document Properties panel has fields for the file's author, title, subject, keywords, category, and status, and any comments about the file.

You can also create custom properties by clicking the arrow located to the right of the Document Properties label, and clicking Advanced Properties to open the Properties dialog box. On the Custom page of the Properties dialog box, you can click one of the existing custom categories or create your own by entering a new property name in the Name field, clicking the Type arrow and selecting a data type (for example, Text, Date, Number, or Yes/No), selecting or entering a value in the Value field, and then clicking Add. If you want to delete an existing custom property, point to the Properties list, click the property you want to get rid of, and click Delete. After you finish making your changes, click the OK button. To hide the Document Properties panel, click the Close button in the upper-right corner of the panel.

When you're done modifying a workbook, you should save your changes and then, to close the file, display the Backstage view and then click Close. You can also click the Close button in the upper-right corner of the workbook window.
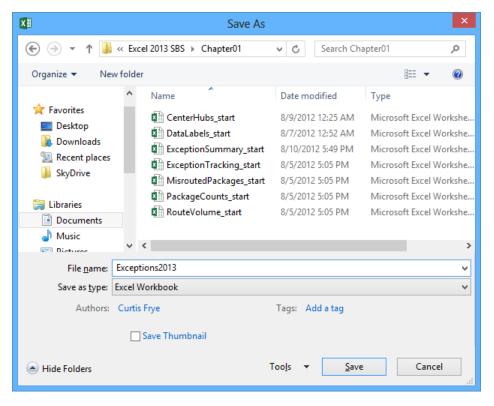
KEYBOARD SHORTCUT Press Ctrl+W to close a workbook.

In this exercise, you'll close an open workbook, create a new workbook, save the workbook with a new name, assign values to the workbook's standard properties, and create a custom property.
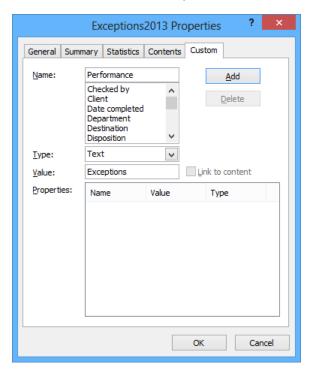
➡ SET UP **You need the ExceptionSummary workbook located in the Chapter01 practice file folder to complete this exercise. Open the workbook, and then follow the steps.**

1   Click the **File** tab to display the **Backstage** view, and then click **Close** to close the **ExceptionSummary** workbook.

2   Display the **Backstage** view, and then click **New** to display the **New** page.

3   Click **Blank workbook**, and then click **Create** to open a new, blank workbook.

4   Display the **Backstage** view, click **Save As**, click **Computer**, and then click **Browse** to open the **Save As** dialog box.

5   Use the navigation controls to display the **Chapter01** folder. In the **File name** field, enter **Exceptions2013**.

6   Click the **Save** button to save your work and close the **Save As** dialog box.

7   Display the **Backstage** view, click **Info**, click **Properties**, and then click **Show Document Panel** to display the **Document Properties** panel.

8   In the **Keywords** field, enter **exceptions, regional, percentage**.

9   In the **Category** field, enter **performance**.

10   Click the arrow at the right end of the **Document Properties** button, and then click **Advanced Properties** to open the **Exceptions2013 Properties** dialog box.

11   Click the **Custom** tab to display the **Custom** page.

12   In the **Name** field, enter **Performance**.

13   In the **Value** field, enter **Exceptions**.
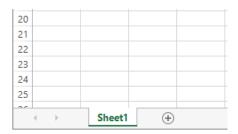

14   Click the **Add** button, and then click **OK** to save the properties and close the **Exceptions2013 Properties** dialog box.

❌ CLEAN UP **Close the Exceptions2013 workbook, saving your changes if you want to.**

Modifying workbooks

Most of the time, you create a workbook to record information about a particular activity, such as the number of packages that a regional distribution center handles or the average time a driver takes to complete all deliveries on a route. Each worksheet within that workbook should represent a subdivision of that activity. To display a particular worksheet, click the worksheet's tab on the tab bar (just below the grid of cells).
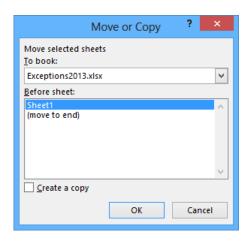
In the case of Consolidated Messenger, the workbook used to track daily package volumes could have a separate worksheet for each regional distribution center. As mentioned earlier,

new Excel workbooks contain one worksheet; because Consolidated Messenger uses nine regional distribution centers, you would need to create eight new worksheets. To create a new worksheet, click the New Sheet button (which looks like a plus sign in a circle) at the right edge of the tab bar.

When you create a worksheet, Excel assigns it a generic name such as *Sheet2*, *Sheet3*, or *Sheet4*. After you decide what type of data you want to store on a worksheet, you should change the default worksheet name to something more descriptive. For example, you could change the name of Sheet1 in the regional distribution center tracking workbook to *Northeast*. When you want to change a worksheet's name, double-click the worksheet's tab on the tab bar to highlight the worksheet name, enter the new name, and press Enter.

Another way to work with more than one worksheet is to copy a worksheet from another workbook to the current workbook. One circumstance in which you might consider copying worksheets to the current workbook is if you have a list of your current employees in another workbook. You can copy worksheets from another workbook by right-clicking the tab of the sheet you want to copy and, on the shortcut menu, clicking Move Or Copy to open the Move Or Copy dialog box.

TIP When you select the Create A Copy check box, Excel leaves the copied worksheet in its original workbook, whereas clearing the check box causes Excel to delete the worksheet from its original workbook.

After the worksheet is in the target workbook, you can change the worksheets' order to make the data easier to locate within the workbook. To change a worksheet's location in the workbook, you drag its sheet tab to the location you want on the tab bar. If you want to remove a worksheet from the tab bar without deleting the worksheet, you can do so by right-clicking the worksheet's tab on the tab bar and clicking Hide on the shortcut menu. When you want Excel to redisplay the worksheet, right-click any visible sheet tab and then click Unhide. In the Unhide dialog box, click the name of the sheet you want to display, and click OK.

To differentiate a worksheet from others, or to visually indicate groups or categories of worksheets in a multiple-worksheet workbook, you can change the color of a worksheet tab. To do so, right-click the tab, point to Tab Color, and then click the color you want.

TIP If you copy a worksheet to another workbook, and the destination workbook has the same Office Theme applied as the active workbook, the worksheet retains its tab color. If the destination workbook has another theme applied, the worksheet's tab color changes to reflect that theme. For more information about Office themes, see Chapter 4, "Changing workbook appearance."

If you determine that you no longer need a particular worksheet, such as one you created to store some figures temporarily, you can delete the worksheet quickly. To do so, right-click its sheet tab, and then click Delete.
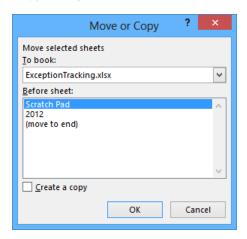
In this exercise, you'll insert and rename a worksheet, change a worksheet's position in a workbook, hide and unhide a worksheet, copy a worksheet to another workbook, change a worksheet's tab color, and delete a worksheet.
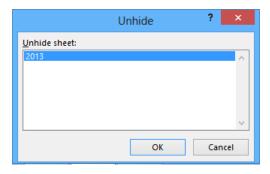
➡️ SET UP **You need the ExceptionTracking workbook located in the Chapter01 practice file folder to complete this exercise. Open the workbook, and then follow the steps.**

1   On the tab bar, click the **New Sheet** button to create a new worksheet.

2   Right-click the new worksheet's sheet tab, and then click **Rename** to highlight the new worksheet's name.

3   Enter **2013**, and then press **Enter**.

4   On the tab bar, double-click the **Sheet1** sheet tab to highlight the worksheet's name.

5   Enter **2012**, and then press **Enter**.

6   Right-click the **2013** sheet tab, point to **Tab Color**, and then, in the **Standard Colors** palette, click the green swatch to change the **2013** sheet tab to green.

7   On the tab bar, drag the **2012** sheet tab to the right of the **Scratch Pad** sheet tab.

8   Right-click the **2013** sheet tab, and then click **Hide** to remove the **2013** sheet tab from the tab bar.

9   Right-click the **2012** sheet tab, and then click **Move or Copy** to open the **Move or Copy** dialog box.

> **Move or Copy**   ?   ✕
>
> Move selected sheets
> To book:
>
> ExceptionTracking.xlsx   ⌄
>
> Before sheet:
>
> Scratch Pad
> 2012
> (move to end)
>
> ☐ Create a copy
>
> OK   Cancel

10 Click the **To book** arrow, and then in the list, click **(new book)**.

11 Select the **Create a copy** check box.

12 Click **OK** to create a new workbook and copy the selected worksheet into it.

13 On the **Quick Access Toolbar**, click **Save** to open the **Save As** dialog box.

14 In the **File name** field, enter **2012 Archive**, and then press **Enter** to save the workbook.

15 On the **View** tab, click the **Switch Windows** button, and then click **ExceptionTracking** to display the **ExceptionTracking** workbook.

16 On the tab bar, right-click the **Scratch Pad** sheet tab, and then click **Delete**. In the dialog box that opens, click **Delete** to confirm the operation.

17 Right-click the **2012** sheet tab, and then click **Unhide** to open the **Unhide** dialog box.

18   Click **2013**, and then click **OK** to close the **Unhide** dialog box and display the **2013** worksheet in the workbook.

✖ CLEAN UP **Close the ExceptionTracking workbook and the 2012 Archive workbook, saving your changes if you want to.**

Modifying worksheets

After you put up the signposts that make your data easy to find, you can take other steps to make the data in your workbooks easier to work with. For example, you can change the width of a column or the height of a row in a worksheet by dragging the column's right border or the row's bottom border to the position you want. Increasing a column's width or a row's height increases the space between cell contents, making your data easier to read and work with.

TIP You can apply the same change to more than one row or column by selecting the rows or columns you want to change and then dragging the border of one of the selected rows or columns to the location you want. When you release the mouse button, all the selected rows or columns change to the new height or width.

Inserting rows, columns, and cells

Modifying column width and row height can make a workbook's contents easier to work with, but you can also insert a row or column between cells that contain data to make your data easier to read. Adding space between the edge of a worksheet and cells that contain data, or perhaps between a label and the data to which it refers, makes the workbook's contents less crowded. You insert rows by clicking a cell and clicking the Home tab on the ribbon. Then, in the Cells group, in the Insert list, click Insert Sheet Rows. Excel inserts a row above the row that contains the active cell. You insert a column in much the same way, by choosing Insert Sheet Columns from the Insert list. When you do this, Excel inserts a column to the left of the active cell.

When you insert a row, column, or cell in a worksheet that has had formatting applied, the Insert Options button appears. When you click the Insert Options button, Excel displays a list of choices you can make about how the inserted row or column should be formatted, as described in the following table.

Option	Action
Format Same As Above	Applies the formatting of the row above the inserted row to the new row
Format Same As Below	Applies the formatting of the row below the inserted row to the new row
Format Same As Left	Applies the formatting of the column to the left of the inserted column to the new column
Format Same As Right	Applies the formatting of the column to the right of the inserted column to the new column
Clear Formatting	Applies the default format to the new row or column

If you want to delete a row or column, right-click the row or column head and then, on the shortcut menu that appears, click Delete. You can temporarily hide rows or columns by selecting those rows or columns and then, on the Home tab, in the Cells group, clicking the Format button, pointing to Hide & Unhide, and then clicking either Hide Rows or Hide Columns. The rows or columns you selected disappear, but they aren't gone for good as they would be if you'd used Delete. Instead, they have just been removed from the display until you call them back. To return the hidden rows to the display, select the row or column headers on either side of the hidden rows or columns. Then, on the Home tab, in the Cells group, click the Format button, point to Hide & Unhide, and then click either Unhide Rows or Unhide Columns.

Just as you can insert rows or columns, you can insert individual cells into a worksheet. To insert a cell, click the cell that is currently in the position where you want the new cell to appear. On the Home tab, in the Cells group, in the Insert list, click Insert Cells to open the Insert dialog box. In the Insert dialog box, you can choose whether to shift the cells surrounding the inserted cell down (if your data is arranged as a column) or to the right (if your data is arranged as a row). When you click OK, the new cell appears, and the contents of affected cells shift down or to the right, as appropriate. Similarly, if you want to delete a block of cells, select the cells, and on the Home tab, in the Cells group, in the Delete list, click Delete Cells to open the Delete dialog box—complete with options that you can use to choose how to shift the position of the cells around the deleted cells.

TIP The Insert dialog box also includes options you can click to insert a new row or column; the Delete dialog box has similar options for deleting an entire row or column.

If you want to move the data in a group of cells to another location in your worksheet, select the cells you want to move and point to the selection's border. When the pointer changes to a four-pointed arrow, you can drag the selected cells to the desired location on the worksheet. If the destination cells contain data, Excel displays a dialog box asking whether you want to overwrite the destination cells' contents. If you want to replace the existing values, click OK. If you don't want to overwrite the existing values, click Cancel and insert the required number of cells to accommodate the data you want to move.

In this exercise, you'll insert a column and row into a worksheet, specify insert options, hide a column, insert a cell into a worksheet, delete a cell from a worksheet, and move a group of cells within the worksheet.
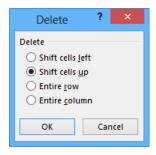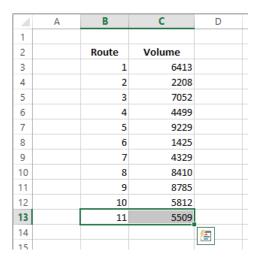
1   On the **May 12** worksheet, select cell **A1**.

2   On the **Home** tab, in the **Cells** group, click the **Insert** arrow, and then in the list, click **Insert Sheet Columns** to create a new column **A**.

3   In the **Insert** list, click **Insert Sheet Rows** to create a new row **1**.

4   Click the **Insert Options** button that appears below the lower-right corner of the selected cell, and then click **Clear Formatting** to remove the formatting from the new row **1**.

5   Right-click the column header of column **E**, and then click **Hide** to remove column **E** from the display.

	A	B	C	D	F
1					
2		Route	Volume		
3		1	6413		
4		2	2208		
5		3	7052		
6		4	9229		
7		5	1425		
8		6	4329		
9		7	8410		
10		8	8785		
11		9	5812		
12		10	1112		
13					

6   On the tab bar, click the **May 13** sheet tab to display the worksheet of the same name.

7   Click cell **B6**.

8   On the **Home** tab, in the **Cells** group, click the **Delete** arrow, and then in the list, click **Delete Cells** to open the **Delete** dialog box.

9   If necessary, click **Shift cells up**, and then click **OK**. Excel deletes cell **B6**, moving the cells below it up to fill in the gap.

10   Click cell **C6**.

11   In the **Cells** group, in the **Insert** list, click **Insert Cells** to open the **Insert** dialog box.

12   If necessary, click **Shift cells down**, and then click **OK** to close the **Insert** dialog box, create a new cell **C6**, and move cells **C6:C11** down to accommodate the inserted cell.

13   In cell **C6**, enter **4499**, and then press **Enter**.

14   Select cells **E13:F13**.

15   Point to the border of the selected cells. When the pointer changes to a four-pointed arrow, drag the selected cells to cells **B13:C13**. The dragged cells replace cells **B13:C13**.

	A	B	C	D
1				
2		Route	Volume	
3		1	6413	
4		2	2208	
5		3	7052	
6		4	4499	
7		5	9229	
8		6	1425	
9		7	4329	
10		8	8410	
11		9	8785	
12		10	5812	
13		11	5509	
14				
15				

✖ CLEAN UP **Close the RouteVolume workbook, saving your changes if you want to.**

Merging and unmerging cells

Most Excel worksheets contain data about a specific subject, such as packages shipped, revenue, or operating costs. One of the best ways to communicate the contents of a worksheet is to use a label. For example, you might create a list of Consolidated Messenger's delivery regions with the city where the region's distribution hub is located.

	A	B	C
1			
2		**Distribution Center Hubs**	
3		Listed by region name and city	
4			
5		Northeast	Boston
6		Atlantic	Baltimore
7		Southeast	Atlanta
8		North Central	Cleveland
9		Midwest	St. Louis
10		Southwest	Albuquerque
11		Mountain West	Denver
12		Northwest	Portland
13		Central	Omaha
14			

The text *Distribution Center Hubs* appears to span two cells, B2 and C2, but is in fact contained within cell B2. If you select cell B2, Excel highlights the cell's border, which obscures the text. If you want to combine cells B2 and C2 into a single cell, you can do so by merging the cells into a single cell.

To merge two or more cells, you select the cells, click the Home tab, click Merge & Center, and then click Merge Cells. Now when you click cell B2, the selection border extends along the entire merged cell without blocking the text.

	A	B	C
1			
2		**Distribution Center Hubs**	
3		Listed by region name and city	
4			
5		Northeast	Boston
6		Atlantic	Baltimore
7		Southeast	Atlanta
8		North Central	Cleveland
9		Midwest	St. Louis
10		Southwest	Albuquerque
11		Mountain West	Denver
12		Northwest	Portland
13		Central	Omaha
14			

IMPORTANT When you merge two or more cells, Excel retains just the text in the range's top left cell. All other text is deleted.

When you click the Merge & Center button, a list of options appears. In addition to merging cells, you can click Merge & Center to combine the selected cells into a single cell and center the text within the merged cell. You should strongly consider using the Merge & Center option for label text, such as above a list of data where the title spans more than one column.

You can also merge the cells in multiple rows at the same time. For example, suppose your list has a main heading and a subheading. You can merge the cells in the two rows that contain headings by clicking the Home tab, clicking Merge & Center, and then clicking Merge Across.

IMPORTANT When you select the header cells, click the Home tab, click Merge & Center, and then click either Merge & Center or Merge Cells, Excel deletes any text that's not in the top-left cell of the selected range.

If you want to split merged cells into their individual cells, click the Home tab, click Merge & Center, and then click Unmerge Cells.

In this exercise, you will merge cells, unmerge cells, merge and center cells, and use Merge Across to combine cells in several rows into one merged cell per row.

1 Select cells **B2:D2**.

2 Click the **Home** tab, click the **Merge & Center** arrow, and then click **Merge Cells** to merge the cells into a single cell.

3 Select cells **B3:F3**.

4 Click the **Home** tab, click the **Merge & Center** arrow, and then click **Merge & Center** to merge the cells into a single cell and center its contents.

5 Select cells **B5:E8**.

6 Click the **Home** tab, click the **Merge & Center** arrow, and then click **Merge Across** to merge the cells in each row into a single cell. The operation creates four merged cells.
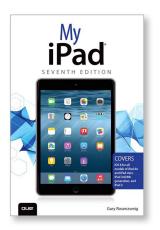
	A	B	C	D	E	F
1						
2		**Delivery Category**				
3		**Brief descriptions of each category**				
4						
5		Overnight: next day delivery by 3PM				
6		2Day: second-day delivery by 6PM				
7		3Day: third-day delivery by 6PM				
8		Ground: delivery in 3-5 days by 6PM				
9						

7 Select cell **B2**.

8 Click the **Home** tab, click the **Merge & Center** arrow, and then click **Unmerge Cells** to split the merged cell into its original cells.

9 Select cells **B2:D2**.

10 Click the **Home** tab, click the **Merge & Center** arrow, and then click **Merge & Center** to merge the cells into a single cell and center its contents.
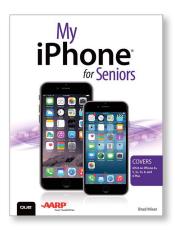
Key points

- Save your work whenever you do something you'd hate to have to do again.

- Assigning values to a workbook's properties makes it easier to find your workbook by searching in File Explorer or by using Windows 8 Search.

- Be sure to give your worksheets descriptive names.

- If you want to use a worksheet's data in another workbook, you can send a copy of the worksheet to that other workbook without deleting the original worksheet.

- You can delete a worksheet you no longer need, but you can also hide a worksheet in a workbook. When you need the data on the worksheet, you can unhide it.

- You can save yourself a lot of cutting and pasting by inserting and deleting worksheet cells, columns, and rows.

- By merging cells, you can add data labels that span multiple columns.

- Customize your Excel 2013 program window by changing how it displays your workbooks, zooming in on data, adding frequently used buttons to the Quick Access Toolbar, and rearranging or customizing the ribbon to meet your needs.
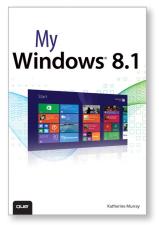
More Best-Selling **My** Books!

My iPad®
SEVENTH EDITION
COVERS iOS 8 for all models of iPad Air and iPad mini, iPad 3rd/4th generation, and iPad 2
Gary Rosenzweig

My iPhone®
EIGHTH EDITION
COVERS iOS 8 on iPhone 4s, 5, 5c, 5s, 6, and 6 Plus
Brad Miser

My iPhone® for Seniors
COVERS iOS 8 on iPhone 4s, 5, 5c, 5s, 6, and 6 Plus
AARP Real Possibilities
Brad Miser

My Office® for iPad®
COVERS Excel, Word, PowerPoint, and OneNote for iPad
Paul McFedries

My Windows® 8.1
Katherine Murray

My GoPro® Hero Camera
COVERS All models of GoPro HERO, HERO3, HERO3+, and HERO4
Jason R. Rich

My Samsung Galaxy S5
Steve Schwartz with Craig James Johnston

My Samsung Galaxy Note® 4
Craig James Johnston
Guy Hart-Davis

Learning to use your smartphone, tablet, camera, game, or software has never been easier with the full-color My Series. You'll find simple, step-by-step instructions from our team of experienced authors. The organized, task-based format allows you to quickly and easily find exactly what you want to achieve.

Visit quepublishing.com/mybooks to learn more.

QUE