

Excel Tutorial to Improve Your Efficiency (2007 Version)

Introduction

My purpose with this Excel tutorial is to illustrate some Excel tips that will dramatically improve your efficiency. I make no attempt to be as encyclopedic as some of the 800-page Excel manuals available. I concentrate on *common* tasks, not every last thing that can be done in Excel. Also, I presume that you have *some* Excel knowledge. For example, I assume you know about rows and columns, values, labels, and formulas, relative and absolute addresses, and other basic Excel elements. If you know virtually nothing about Excel, you probably ought to work through an “Excel for Dummies” book and *then* work through this tutorial.

The style of this tutorial should be easy to follow. Main topics appear in bold black type. Specific direction headings are in **yellow**, and these are followed by detailed directions in **red**. Additional comments about the directions appear in **blue**. Then there are “Try it!” exercises in **green**. These “Try it!” exercises are a key feature of this tutorial. I have embedded numerous sample Excel spreadsheets so that you can try out the directions right away—without switching into Excel. When you double-click on one of these spreadsheets, you launch Excel, and the spreadsheet “comes alive.” The menus and toolbars even change to those for Excel. By clicking outside one of these spreadsheets, you’re back in Word.

The easiest way to maneuver around this tutorial is to switch to outline view. To do so, select Word’s View→Outline menu item, which gives you an extra Outlining menu. Select this menu and then click on the Show Level dropdown. This lets you choose the level of the outline. You’ll get good results by choosing Level 2. This lets you see all first-level and second-level headings. Put your cursor on a heading of interest and then switch back to normal view by selecting Word’s View→Print Layout menu item. (Try it right now. It’s easy!)

Finally, I suggest that you save this file—RIGHT NOW—as **MyXLTutorial.docx** (or some such name) and work with the copy. That way, if you mess anything up as you try the exercises, you can always go back and retrieve the original file (**ExcelTutorial.docx**).

Have fun!

Differences in Excel 2007

If Excel 2007 is your first exposure to Excel, you’ll love it. If you’re used to Excel 2003 or an earlier version, however, you’ll have some unlearning to do. There are two *big* changes in Excel 2007 and a host of smaller changes. The first big change is that worksheets are *much* bigger. You used to have about 65,000 rows and 256 columns. Now you have over a million rows and over 16,000 columns. You’ll hardly ever *use* this much space, but it’s available.

The second big change is the one you’ll notice right away: you no longer have the usual menus and toolbars. Microsoft has completely reorganized the user interface to provide menus and ribbons. Each menu (Home, Insert, Page Layout, etc.) has an associated ribbon that is similar to the old toolbars. For example, if you click on the Formula menu, you get a ribbon with buttons that are useful for working with formulas. Each ribbon has several groups of buttons. For example, the Formulas ribbon has one group called Defined Names with buttons for manipulating range names. There’s only one way to learn these ribbons: practice and experiment. If you’re used to the old Excel, you’ll undoubtedly curse at the new ones a few times when you can’t find something, but you can be assured that they *are* more logically organized than the old versions.

One particular menu item you’ll miss is Tools→Options, for changing various options in Excel. It’s not gone; it’s just in a different place. To get to it, click on the new Office button (in the upper left corner of the screen) and then the Excel Options button. From there, you can

experiment with the various options. While you're at it, note the options along the left when you click on the Office button. For example, that's where you find Save As and Print, two obvious favorites.

You can't change the new ribbons; they are built in and fixed. However, you *can* customize the user interface to some extent. First, Excel gives you a quick access toolbar (QAT) at the top left of the screen. You can put your favorite buttons on this toolbar so that they're always visible and available. The QAT comes with a few favorite buttons, but you can add more. Just click on the dropdown arrow to the right of the QAT and check any of your favorites (or click on More Commands to get more buttons). Second, if you don't mind writing a bit of XML, you can create new ribbons that have your favorite Excel buttons or even new buttons attached to your own macros. (This is somewhat advanced, so it won't be covered in this document.)

With the old menus gone, what about the keyboard shortcuts many of you depend on? As far as I can tell, they still work, or at least most of them still work. For example, you can still press Ctrl-s to save a file or Ctrl-p to print a file. However, in this document, where I've embedded spreadsheets, the keyboard shortcuts for *Excel* won't work. This is evidently because the spreadsheets are embedded in a *Word* document.

One other important change is file extensions. Unless you wrote your own add-ins in earlier versions of Excel, the only file extension you probably ever saw was .xls. You can still save files with this extension, but they'll lose any features new to Excel 2007. To save them and keep new features, you should save them as .xlsx files. Note that if you do so, people with Excel 2003 won't be able to open them unless they download and install a free file format converter from Microsoft. (Google for Microsoft Office Compatibility Pack to find this free download.) One other new file extension is .xlsm. If your Excel file has associated VBA macros, then you should save it as an .xlsm file.

Basic Spreadsheet Tasks

This section illustrates a number of ways to make you more efficient: how to select ranges, how to copy and paste, how to deal with absolute and relative addresses, how to insert and delete rows or columns, and a few others. Most Excel users know how to do these things, but they often do them inefficiently. So even though this material might seem elementary, read on—you might just pick up a few tricks you weren't aware of.

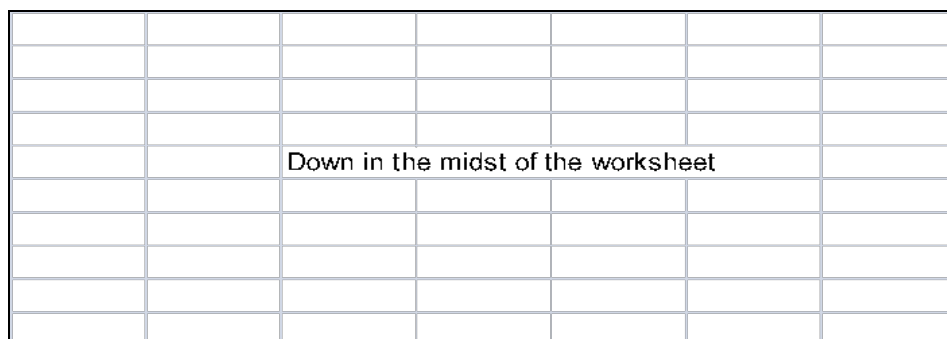
Moving to the Top of the Sheet

Often you want to reorient yourself by going back to the "home" position on the worksheet.

To go to the top left of the sheet (cell A1):

Press Ctrl-Home (both keys at once).

Try it!



Using End-Arrow Key Combinations

To go to the end of a range (top, bottom, left, or right):

Press the End key, then the appropriate arrow key. For example, press End and then right arrow to go to the right edge of a range.

Try it! Starting at a corner (a bordered cell), move around to the other corners.

		8	1	1	1	
		5	5	7	6	
		10	5	3	10	
		10	10	5	9	
		5	4	3	4	
		7	4	5	1	
		2	8	9	10	

The action of an End-arrow combination depends on where you start. It takes you to the last nonblank cell if you start in a *nonblank* cell. (If there aren't any nonblank cells in that direction, it takes you to the far edge of the sheet.) If you start in a *blank* cell, it takes you to the first nonblank cell.

Splitting the Screen

It is often useful to split the screen so that you can see more information.

To split the screen vertically, horizontally, or both:

Click on the narrow “screen splitter” bar just to the right of the bottom scroll bar (for vertical splitting) or just above the right-hand scroll bar (for horizontal splitting) and drag this to the left or down.

Splitting gives you two “panes” (or four if you split in both directions). Once you have these panes, practice scrolling around in any of them, and see how the others react.

Try it! Split the screen either way and then remove the split.

	77	74	52	93	12	
	9	38	4	73	38	
	30	31	37	98	10	
	45	35	3			
	34	23	11			
	80	45	26	64	17	
	33	13	55	31	53	
	64	55	99	5	87	

Selecting a Range

Usually in Excel, you select a range and then do something to it (such as enter a formula in it, format it, delete its contents, and so on). Therefore, it is extremely important to be able to select a range *efficiently*. It's easy if the whole range appears on the screen, but it's a bit trickier if you

can't see the whole range. In the latter case the effect of dragging (the method most users try) can be frustrating—things scroll by too quickly. Try one of the methods below instead.

To select a range that fits on a screen:

Click on one corner of the range and drag to the opposite corner.

Or:

Click on one corner, hold down the Shift key, and click on the opposite corner.

Try it! Select the range B2:D7.

	9	2	10			
	1	3	8			
	8	1	10			
	7	5	4			
	5	1	1			
	5	10	7			

To select a range doesn't fit on a screen:

Click on one corner of the range, say, the upper left corner. Then, holding the Shift key down, use the End-arrow combinations (End and right arrow, then, if necessary, End and down arrow) to get to the opposite corner.

Or:

Split the screen so that one corner shows in one pane and the opposite corner shows in the other pane. Click on one corner, hold the Shift key down, and click on the opposite corner.

Try it! Select the range B2:C100 or the range E2:N5. Try both of the methods suggested above.

	6	3		7	13	9
	7	2		5	6	11
	1	3		1	2	9
	6	4		9	13	1
	4	4				
	2	7				
	2	4				
	2	2				
	8	6				

Selecting Multiple Ranges

Say you want to format more than one range in a certain way (as currency, for example). The quickest way is to select all ranges at once and then format them all at once.

To select more than one range:

Select the first range, press the Ctrl key, select the second range, press the Ctrl key, select the third range, and so on.

For example, to select the ranges B2:D5 and F2:H5, click on B2, hold down the Shift key and click on D5 (so now the first range is selected), hold down the Ctrl key and click on F2, and finally hold down the Shift key and click on H5.

Try it! Select all three numerical ranges shown.

	10	13		7	9	
	1	9		12	7	
	5	4		1	1	
				6	12	
				15	2	
		13	4	10		

Copying and Pasting

Copying and pasting (often formulas) is one of the most frequently done tasks in Excel, and it can be a real time-waster if done inefficiently. Many people do it as follows. They select the range to be copied (often in an inefficient manner), then select the Edit→Copy menu item, then select the paste range (again, often inefficiently), and finally select the Edit→Paste menu item. There are much quicker ways to get the job done!

To copy and paste using keyboard shortcuts:

Select the copy range (using one of the efficient selection methods described above), press **Ctrl-c** (for copy), select the paste range (again, efficiently), and press **Ctrl-v** (for paste). (If you practice pressing Ctrl-c or Ctrl-v with the little finger and index finger of your left hand, it will become automatic in a short time.)

The copy range will still have a dotted line around it. Press the Esc key to get rid of it.

Try it! Copy the formula in cell C2 down through cell C8 using Ctrl-c and Ctrl-v.

	3	3	9			
	4	1				
	2	3				
	2	1				
	5	1				
	4	2				
	3	5				

To copy and paste using toolbar buttons:

Proceed as above, but use the copy and paste buttons (in the left section of the Home ribbon) instead of the Ctrl-c and Ctrl-v key combinations.

Try it! Copy the formula in C2 down through cell C8 using the Copy and Paste buttons.

	1	3	3			
	4	1				
	5	5				
	1	4				
	3	3				
	3	5				
	1	5				

Buttons or key combinations? This is a matter of personal taste, but either is quicker than menu choices!

A frequent task is to enter a formula in one cell and copy it down a row or across a column. There are several *very* efficient ways to do this.

To avoid copying and pasting altogether, use **Ctrl-Enter**:

Starting with the top or left-hand cell, select the range where the results will go. (Use the selection methods described earlier, especially if this range is a long one.) Type in the formula, and press **Ctrl-Enter** instead of Enter.

Try it! Fill up the range C2:C8 with Ctrl-Enter.

	6	2				
	9	1				
	4	9				
	7	6				
	2	3				
	8	5				
	3	9				

Formula should multiply the value in column A by the corresponding value in column B

Pressing Ctrl-Enter enters what you typed in all of the selected cells (adjusted for relative addresses), so in general, it can be a real time saver. For example, it can be used to enter the number 10 in a whole range of cells. Just select the range, type 10, and press Ctrl-Enter.

Try it! Fill up the range B2:D8 with the value 10 by using Ctrl-Enter.

To copy with the drag handle:

Enter the formula in the top or left-hand cell of the intended range. Place the cursor on the “drag handle” at the lower right of this cell (the cursor becomes a plus sign), and drag this handle down or across to copy.

Try it! Copy the formula in C2 down through C8 with the “drag handle”.

10	7	70				
6	9					
4	9					
5	1					
2	8					
8	10					
10	7					

To copy by double-clicking on the drag handle:

Enter the formula in the top or left-hand cell of the intended range. Double-click on the drag handle.

This method uses Excel’s built-in intelligence, but it works only in certain situations. Let’s say you have numbers in the range A3:B100. You want to enter a formula in cell C3 and copy it down to cell C100. Because this is a common thing to do, Excel does it for you if you double-click on the drag handle. It senses the “filled-up” range in column B and guesses that you want another filled-up range right next to it in column C. If there were no adjacent filled-up range, double-clicking on the drag handle wouldn’t work.

Try it! Copy the formula in C2 down through C8 by double-clicking the “drag handle.”

7	8	56				
2	1					
10	2					
4	1					
4	5					
3	8					
7	8					

Copying and Pasting with the Special/Values Option

Often you have a range of cells that contains formulas, and you would like to replace the formulas with the values they produce. Usually, you paste these values onto the copy range, that is, you overwrite the formulas with values. However, you could also select another range for the paste range.

To copy formulas and paste values:

Select the range with formulas, press Ctrl-c to copy, and select the range where you want to paste the values (which could be the same as the copy range). Then (because there is no keyboard equivalent) select the Paste dropdown on the Home ribbon, and select the Paste Values option.

Try it! Copy the range D2:D8 to itself, but paste values.

		1	1	1		
		2	3	6		
		9	5	45		
		7	4	28		
		8	8	64		
		8	10	80		
		7	6	42		

You might want to experiment with the other options in the Paste dropdown. For example, if you have a set of labels entered as a row and you want this same set of labels entered somewhere else as a column, try copying and pasting with the Transpose option.

Moving (Cutting and Pasting)

Often you would like to move information from one place in the sheet to another.

To move (cut and paste):

Select the range to be cut, press **Ctrl-x** (for cutting), select the upper left corner of the paste range, and press **Ctrl-v**. (The little finger-index finger combination in your left hand is also good for pressing **Ctrl-x**.)

As with copying and pasting, ribbon buttons can be used instead of key combinations, but either is more efficient than selecting menu items. Also, note that you need only select the upper left cell of the paste range. Excel knows that the shape of the paste range must be the same as the shape of the cut range.

Try it! Move the range A2:C8 to the range D2:F8. (Watch how relative addresses affect the eventual formulas in column F.)

	2	3	6			
	7	4	28			
	3	6	18			
	3	9	27			
	6	7	42			
	6	6	36			
	1	4	4			

Absolute/Relative References

Absolute and references are indicated in formulas by dollar signs or the lack of them, and they indicate what happens when you copy or move a formula to a range. You typically want some parts of the formula to stay fixed (absolute) and others to change relative to the cell position. This is a crucial concept for efficiency in spreadsheet operations, so you should take some time to understand it thoroughly. Here are two important things to remember: (1) The dollar signs are relevant *only* for the purpose of copying or moving; they have no inherent effect on the formula. For example, the formulas =5*B3 and =5*\$B\$3 in cell C3, say, produce exactly the same result. Their difference is relevant only if you want to copy cell C3 to some range. (2) There is never any need to *type* the dollar signs. This can be done with the F4 key.

To make a cell reference absolute or mixed absolute/relative using the F4 key:

Enter a cell reference such as B3 in a formula. Then press the F4 key.

In fact, pressing the F4 key repeatedly cycles through the possibilities: B3 (neither row nor column fixed), then \$B\$3 (both column B and row 3 fixed), then B\$3 (only row 3 fixed), then \$B3 (only column B fixed), and back again to B3.

Try it! Enter the appropriate formula in cell B7 and copy across to E7. (Scroll to the right to see the correct answer.)

Fixed cost	\$50				
Variable cost	\$2				
Month	Jan	Feb	Mar	Apr	
Units produced	224	194	228	258	
Total cost					

Try it again! Enter one formula with appropriate absolute/relative addressing in cell C5 that can be copied to C5:F9. (Scroll to the right to see the correct answer.)

Table of revenues for various unit prices and units sold					
		Units sold			
		50	100	150	200
Unit price	\$3.25				
	\$3.50				
	\$3.75				
	\$4.00				
	\$4.25				

Showing Formulas in R1C1 Notation

Speaking of relative and absolute addresses in formulas, it's pretty amazing how intelligent spreadsheets are. When you copy a formula in cell C1 such as =A1+B1 down, it automatically changes appropriately: =A2+B2, then =A3+B3, etc. In a sense, these are all the same formula. Each says to add the two values to the left of the current cell. Excel allows you to see this equivalence even more clearly by viewing the formulas in a different format, called R1C1 notation. In this format, *each* of the formulas in column C is written as =RC[-2]+RC[-1]. R stands for row, and C stands for column. The fact that there is nothing next to R means we stay in the *same* row. The numbers in brackets next to C mean to go 2 columns to the left and 1 column to the left. (For columns, negative numbers mean to go to the left, positive to the right. For rows, negative numbers mean to go up, positive down.)

If there is a number next to R or C that is *not* in brackets, it indicates an absolute reference. For example, =R2C[2] placed in cell D5 is equivalent to =F\$2 since the row reference is absolute and the column reference is relative to column D.

The usual way of expressing formulas, such as =C5+D5, is called **A1 format**. The new way discussed here is called **R1C1 format**. You can easily toggle between them.

To toggle between A1 and R1C1 formats:

Select the Office button, select Excel options, select Formulas, and check or uncheck the R1C1 reference style option.

Try it! Toggle between A1 and R1C1 reference style in the following spreadsheet, and for either, examine the formulas inside the border and in the Total row. (Note: As far as I can tell, you'll have to actually open Excel to try this one. When you double-click on the spreadsheet below, the Excel menu appears, but not the Office button. So you evidently can't change the Excel options from within Word. I suppose that makes sense.)

Multiplication table					
		1	2	3	4
1		1	2	3	4
2		2	4	6	8
3		3	6	9	12
4		4	8	12	16
Total		10	20	30	40

I have read one Excel book that advocates the use of R1C1 notation everywhere, reasoning that this notation makes more logical sense. Maybe the author has a point, but the A1 notation is so ingrained in most of us that the transition to R1C1 notation would probably start a revolt among millions of Excel users!

Inserting and Deleting Rows or Columns

Often you want to insert or delete rows or columns. Note that deleting a row or column is *not* the same as clearing the contents of a row or column—making all of its cells blank. Deleting a row or column means wiping it out completely.

To insert one or more blank rows:

Click on a row number and drag down as many rows as you want to insert. Then click on the Insert dropdown and select Insert Sheet Rows. (Interestingly, the Insert dropdown is in the Cells group of the Home ribbon, *not* on the Insert ribbon.)

The rows you insert are inserted above the first row you selected. For example, if you select rows 8 through 11 and then insert, four blank rows will be inserted between the old rows 7 and 8.

Try it! Insert blank rows for the data for Feb, Apr, and May.

Month	Price	Units sold	Revenue			
Jan	\$3.00	100	\$300.00			
Mar	\$3.25	50	\$162.50			
Jun	\$3.50	200	\$700.00			

Columns are inserted in the same way.

Try it! Insert blank columns for sales reps Baker, Miller, and Smith (so that the sales reps are in alphabetical order from left to right).

Sales rep	Allison	Jones	Taylor			
Commission rate	5.4%	6.5%	4.3%			
Sales	\$15,000	\$12,000	\$17,000			
Commission	\$810	\$780	\$731			

To delete one or more rows:

Click on a row number and drag down as many rows as you want to delete. Then select the Delete dropdown and select Delete Sheet Rows. Columns are deleted in exactly the same way.

Try it! The company no longer carries products K322 and R543, so get rid of their rows.

Product Code	Units sold	Unit price				
J645	148	\$15.00				
K322	278	\$17.50				
L254	384	\$25.00				
M332	13	\$30.50				
R543	247	\$22.40				
S654	315	\$35.00				

Filling a Series

Say you want to fill column A, starting in cell A2, with the values 1, 2, and so on up to 1000. There is an easy way.

To fill a column range with a series:

Enter the first value in the first cell (1 in cell A2). With the cursor in the starting cell (A2), select the Fill dropdown and then select the Series option to obtain a dialog box. (The Fill dropdown is in the Editing group of the Home ribbon.) Change the Rows setting to Columns, make sure the

Type setting is Linear, make sure 1 is in the Step Value box, enter the final value (1000) in the Stop Value box, and click on OK.

As you can guess from this dialog box, many other options are possible. Don't be afraid to experiment with them.

Try it! The series of days in column A should go from 1 to 25, in column D it should go from 26 to 50.

Day	Sales		Day	Sales		
	\$227			\$167		
	\$157			\$107		
	\$143			\$255		
	\$129			\$113		
	\$102			\$186		
	\$116			\$124		
	\$269			\$271		
	\$111			\$288		
	\$210			\$273		

Summation Button

The SUM function is used so often to sum across rows or columns that a button (the Σ button) is available to automate the procedure. To illustrate its use, suppose you have a table of numbers in the range B3:E7. You want the row sums to appear in the range F3:F7, and you want the column sums to appear in the range B8:E8. It's easy.

To produce row and column sums with the summation button:

Select the range(s) where you want the sums (F3:F7 and B8:E8—remember how to select multiple ranges?), and click on the summation button.

Note that if you select multiple cells, you get the sums automatically. If you select a single cell (such as when you have a single column of numbers to sum), you'll be shown the sum formula "for your approval" and you'll have to press Enter to actually enter it. Why does Excel do it this way—your guess is as good as mine!

Try it! Use the summation button to fill in the row and column sums.

	51	94	15	7		
	37	6	2	41		
	13	83	29	88		
	73	64	46	32		
	38	11	3	80		

The summation button is in the Editing group of the Home ribbon. If you want a sum, click directly on the Σ button. Alternately, you can click on the AutoSum dropdown for other options, including Average, Max, Min, and others.

Transposing a Range

Often you set up a spreadsheet and then decide that you would rather have a portion of it transposed. That is, you would like to “turn it on its side,” so that rows become columns and vice versa. This is simple with one of Excel’s Paste options.

To transpose a range:

Select a range that you want to transpose and press Ctrl-c to copy it. Then select the upper left cell of the range where you want the transposed version to go, select the Paste dropdown, and select the Transpose option.

Make sure there is enough room for the transposed version. For example, if the original range has 3 rows and 5 columns, the transposed version will have 5 rows and 3 columns. If you select cell D5, say, as the upper left cell for the transposed version, everything in the range D5:F9 will be overwritten by the transposed version.

Try it! Transpose the range A3:D6 to a range with upper left cell F3.

Sales figures						
	Jan	Feb	Mar			
North	2300	3200	4500			
West	4300	3500	4200			
East	5500	6500	6000			

Range Names

Range names are extremely useful for making your formulas more understandable. After all, which formula makes more sense: =B20-B21 or =Revenue-Cost? Efficient use of range names takes some experience, but here are a few useful tips.

To create a range name:

Select a range that you want to name. Then type the desired range name in the upper left “name box” on the screen. (This box is just above the column A heading. It usually shows the cell address, such as E13, where the cursor is.)

You can also select the Formulas menu and use the Define Name option in the Defined Names group to name a range, but typing the range name in the name box is quicker and more intuitive. By the way, range names are not case sensitive, so that Revenue, revenue, and REVENUE can be used interchangeably.

Try it! Name the rectangular range containing the numbers Data.

	71	31	9	69	5	
	15	74	46	84	27	
	14	49	25	38	83	
	40	43	20	75	83	
	28	72	30	92	75	
	41	56	90	89	73	
	28	81	43	81	61	

When you type the range name in the name box, make sure you then press Enter to make the name “stick.” It’s easy to type the name and then click on some other cell *without* pressing Enter. (I’ve done it hundreds of time.) If you do so, you’ll find that the range name was not created.

To delete a range name:

Select the Name Manager in the Defined Names group of the Formulas ribbon. This shows a list of all range names in your workbook. Click on the one you want to delete, and then click on the Delete button.

Suppose a range has name Costs and you want to rename it UnitCosts. If you highlight the range, the name box will show Costs. If you then overwrite this with UnitCosts in the name box, the range will have *two* names, Costs and UnitCosts. There is nothing inherently wrong with this, but if you want only a single name, UnitCosts, you will have to delete the Costs name, as described here.

Try it! The numerical range is currently named Data. Delete this range name and then rename the range Database.

	15	86	50	41	21	
	63	18	10	1	9	
	14	52	16	18	65	
	89	88	28	72	7	
	15	28	57	92	2	
	29	90	86	100	83	
	10	82	100	65	4	

Suppose you have the labels Revenue, Cost, and Profit in cells A20, A21, and A22, and you would like the cells B20, B21, and B22 (which will contain the values of revenue, cost, and profit) to have these range names. Here’s how to do it quickly.

To create range names from adjacent labels:

Select the range consisting of the labels *and* the cells to be named (A20:B22). Then click on the Create from Selection button in the Defined Names group of the Formulas ribbon. In the resulting dialog box, make sure the appropriate option (in this case, Left Column) is checked, and click on OK.

Excel tries (usually successfully) to guess where the labels are that you want to use as range names. If it guesses incorrectly, you can always override its guess.

Try it! Name the ranges A3:A8, B3:B8, and so on according to the labels in row 2.

Month	UnitsSold	UnitPrice	Revenue			
Jan	100	\$1.25	\$125.00			
Feb	150	\$1.25	\$187.50			
Mar	200	\$1.40	\$280.00			
Apr	230	\$1.40	\$322.00			
May	200	\$1.50	\$300.00			
Jun	300	\$1.50	\$450.00			

Sometimes you enter a formula using cell addresses, such as **=B20-B21**. Later, you name B20 as Revenue and B21 as Cost. The formula does *not* change to **=Revenue-Cost** automatically. However, it would be nice to make it change (and hence become more readable). In Excel 2003 and previous versions, you could highlight the formula and use the Insert→Name→Apply menu item to change the formula so that it showed range names. However, this option appears to have been deleted in Excel 2007. You can use the Use in Formula option from the Defined Names group, but this is basically just a “crutch” to help you redefine your formula. Microsoft evidently had its reasons for deleting this functionality, but I’m not what they were. Too bad!

To see a list of all range names and check which ranges they apply to:

Click on the down arrow at the right of the name box, and click on any of the range names you see. That range will then be selected automatically.

Try it! There are five named ranges below. Locate them.

	73		4	28		
	92		29	10		
	31		38	35		
	60		2	40		
	49		45	5		
	22		12	44		
	39		6	21		
	88					
	98	Junk	Junk	Junk	Junk	
	110					

Sometimes it is straightforward to use range names in formulas. For example, if B20 is named Revenue and B21 is named Cost, then entering the formula **=Revenue-Cost** in, say, cell B22 is a natural thing to do. But consider this situation. The range B3:B14 contains revenues for each of 12 months, and its range name is Revenue. Similarly, C3:C14 contains costs, and its range name is Cost. For each month you want that month’s revenue minus cost in the appropriate cell in column D. You will get it correct if you select the range D3:D14, type the formula **=Revenue-Cost**, and press Ctrl-Enter. If you click on any cell in this range, you’ll see the formula **=Revenue-Cost**.

This is pretty amazing. How does Excel know that the formula in D3, for example, is really **=B3-C3**? Let’s just say that it’s smart enough to figure this out. If it confuses *you*, however, you can always enter **=B3-C3** and copy it down. Then you’re safe, but you’ve lost the advantage of range names!

Try it! Enter the formula for all of D3:D14 using range names. (If you like, calculate profits again in column E in the usual way, without range names.)

	Revenue	Cost	Profit			
	\$1,600	\$1,400				
	\$2,000	\$1,800				
	\$2,100	\$1,800				
	\$2,900	\$2,800				
	\$500	\$400				
	\$1,700	\$1,500				
	\$2,000	\$1,900				
	\$2,500	\$2,300				

To paste a list of all range names on a sheet:

To document your spreadsheet, it is often useful to create a list of all of your range names and their corresponding cell addresses. This is easy with Paste Names option.

To paste a list of all range names and corresponding addresses, select a cell with plenty of blank space below it, select the Use in Formula dropdown in the Defined Names group, and click on the Paste Names option.

Try it! Paste a list of all range names, starting in cell D2.

Revenue	\$34,000		Range names used:		
Cost	\$22,000				
Profit	\$12,000				

Basic Excel Functions

There are many useful functions in Excel. You should become familiar with the ones most useful to you. For example, financial analysts should learn the financial functions. But here are a few that everyone should know. (By the way, I capitalize the names of these functions, just for emphasis. However, they are *not* case sensitive. You can enter SUM or sum, for example, with the same result.)

SUM Function

The SUM function is probably the most used Excel function of all. It sums all values in one or more ranges.

To use the SUM function:

Enter the formula =SUM(*range*), where *range* is any range. This sums the numerical values in the range.

Actually, it is possible to include more than one range in a SUM formula, as long as they are separated by commas. (This can also be done with the COUNT, COUNTA, AVERAGE, MAX, and MIN functions discussed below.) For example, =SUM(B5,C10:D12,Revenues) is allowable (where Revenues is the name for some range). The result is the sum of the numerical values in *all* of these ranges combined. Note that if any cell in any of these ranges contains a label rather than a number, it is ignored in the sum.

Try it! Use the SUM function in cell B10 to calculate the total of all costs.

Table of costs for units produced in one month (along side)				
for use in another month (along top)				
	Feb	Mar	Apr	May
Jan	\$4,300	\$3,100	\$3,400	\$5,900
Feb		\$5,500	\$6,400	\$6,500
Mar			\$6,800	\$6,400
Apr				\$5,900
Total cost				

COUNT, COUNTA Functions

The COUNT function counts all of the cells in a range with *numeric* values. The COUNTA functions counts all *nonblank* cells in a range.

To use the COUNT function:

Enter the formula =COUNT(*range*), where *range* is any range. This returns the number of numeric values in the range.

To use the COUNTA function:

Enter the formula =COUNTA(*range*), where *range* is any range. This returns the number of nonblank cells in the range.

For example, if cells A1, A2, and A3 contain Month, 1, and 2, respectively, then =COUNT(A1:A3) returns 2, whereas =COUNTA(A1:A3) returns 3.

Try it! Use the COUNT and COUNTA functions to fill in cells E1 and E2. Note that there are students below the visible portion of the spreadsheet.

Student ID	Exam score	Number enrolled
3416	62	Number who took exam
6125	73	
1535	74	
2323	Absent	
577	77	
9044	57	
8403	67	
5892	90	
4242	77	

AVERAGE Function

The AVERAGE function averages all of the *numeric* cells in a range.

To use the AVERAGE function:

Enter the formula =AVERAGE(*range*) where *range* is any range. This produces the average of the *numeric* values in the range.

Note that the AVERAGE function ignores labels and blank cells in the average. So, for example, if the range C3:C50 includes scores for students on a test, but cells C6 and C32 are blank because these students haven't yet taken the test, then =AVERAGE(C3:C50) averages only the scores for

the students who took the test. (It does *not* automatically average in zeros for the two who didn't take the test.)

Try it! Use the AVERAGE function to calculate the averages in cells B1 and B2. (For B2, you'll have to replicate the exam scores in column C and make some changes.)

Average exam score (for students who took the exam)					
Average exam score (if absent students get zeros)					
Student ID	Exam score				
1533	68				
8031	74				
9859	80				
9106	63				
3535	72				
8192	Absent				

MAX, MIN Functions

The MAX function returns the largest *numeric* value in a range. Similarly, the MIN function returns the smallest numeric value in a range.

To use MAX and MIN functions:

Enter the formula =MAX(*range*) or =MIN(*range*) where *range* is any range. These produce the obvious results: the maximum (or minimum) value in the range.

Try it! Use the MAX and MIN functions to fill in the range B8:C9. For example, you want the values \$2300 and \$3600 in cells B8 and C9.

Sales rep	Allison	Baker	Jones	Miller	Smith	Taylor
Jan sales	\$3,700	\$2,400	\$2,300	\$3,000	\$3,800	\$3,700
Feb sales	\$2,600	\$2,200	\$2,400	\$2,800	\$3,600	\$2,300
	Min sales	Max sales				
Jan						
Feb						

SUMPRODUCT Function

There are many times when you need to sum products of values in two (or possibly more than two) similar-sized ranges. Fortunately, there is an Excel function that sums products quickly.

To use the SUMPRODUCT function

Enter the formula =SUMPRODUCT(*range1,range2*), where *range1* and *range2* are exactly the same size. For example, they might be two column ranges with 10 cells each, or they might be two ranges with 4 rows and 10 columns each. The formula sums the products of the corresponding values from the two ranges.

There can actually be more than two ranges in the SUMPRODUCT formula, separated by commas, as long as all of them have exactly the same size. This is not as common as having only two ranges, but it is sometimes useful

Try it! Sum the products of the two ranges in the following spreadsheet to find the total shipping cost. Enter the result in cell G1. (Scroll to the right for the answer.) By the way, if you are tempted to write the formula *without* the SUMPRODUCT function as the sum of 9 products, as many of my students continue to do, imagine how long your formula would be if there were 10 plants and 50 cities!

Unit shipping costs	City1	City2	City3		Total cost	
Plant1	1.25	1.35	1.55			
Plant2	1.15	1.45	1.25			
Plant3	1.35	1.45	1.15			
Units shipped	City1	City2	City3			
Plant1	155	180	0			
Plant2	250	130	185			
Plant3	0	210	140			

IF Function

IF functions are very useful, and they vary from simple to complex. I'll provide a few examples.

To enter a basic IF function:

Enter the formula =IF(*condition*,*expression1*,*expression2*), where *condition* is any condition that is either true or false, *expression1* is the value of the formula if the condition is true, and *expression2* is the value of the formula if the condition is false.

A simple example is =IF(A1<5,10,"NA"). Note that if either of the expressions is a label (as opposed to a numeric value), it should be enclosed in double quotes.

Try it! Enter appropriate IF formulas in columns C and D. (Scroll to the right to see the correct answer.)

For each product, if the end inventory is less than or equal to 50 units, enough units are ordered to bring stock back up to 200; otherwise, no units of that product are ordered			
Product	End inventory	Order placed (yes or no)?	# of units ordered
1	100		
2	40		
3	20		
4	70		

Sometimes IF functions are nested. For example, there might be three possibilities, depending on whether the value in cell A1 is negative, zero, or positive. A nested IF formula can then be used as follows.

To use nested IF functions:

Enter the formula =IF(*condition1*,*expression1*,IF(*condition2*,*expression2*,*expression3*)). If *condition1* is true, the relevant value is *expression1*. Otherwise, *condition2* is checked. If it is true, the relevant value is *expression2*. Otherwise, the relevant value is *expression3*.

An example is =IF(A1<0,10,IF(A1=0,20,30)). Suppose this formula is entered in cell B2. Then if A1 contains a negative number, B2 contains 10. Otherwise, if A1 contains 0, B2 contains 20. Otherwise (meaning that A1 must contain a positive value), B2 contains 30.

Try it! Use a nested IF function to fill in the grades in column C. (Scroll to the right to see the correct answer.)

Each student gets an A (if score is 90 or above), S for satisfactory (if score if 60 or above but less than 90) or U for unsatisfactory if score is below 60						
Student	Score	Grade				
1	70					
2	95					
3	55					
4	80					
5	60					
6	90					

Sometimes more complex conditions (AND/OR conditions) are useful in IF functions. These are not difficult once you know the syntax.

To use an AND condition in an IF function:

Enter the formula =IF(AND(*condition1,condition2*),*expression1,expression2*). This results in *expression1* if both *condition1* and *condition2* are true. Otherwise, it results in *expression2*.

Note the syntax. The keyword AND is followed by the conditions, separated by a comma and enclosed within parentheses. Also, note that more than two conditions could be included in the AND, all separated by commas.

Try it! Use an IF function with an AND condition to fill in the bordered range. (Scroll to the right to see the correct answer. Note the double quotes for labels.)

Investor sells stock only if its price has gone up three consecutive days (including the current day)						
Day	Price change	Sell (yes or no)?				
1	Up					
2	Down					
3	Up					
4	Up					
5	Up					
6	Down					

To use an OR condition in an IF function:

Enter the formula =IF(OR(*condition1,condition2*),*expression1,expression2*). This results in *expression1* if either *condition1* or *condition2* is true (or if both are true). Otherwise, it results in *expression2*.

Again, more than two conditions could be included in the OR.

Try it! Use an IF function with an OR condition to fill in the bonuses in column F. (Scroll to the right to see the correct answer.)

Any student who scores at least 95 on <i>any</i> of the exams gets a bonus which is 1% of their total score.					
Student	Exam 1	Exam 2	Exam 3	Exam 4	Bonus
1	87	83	83	80	
2	77	72	74	97	
3	80	95	79	75	
4	82	87	96	88	
5	78	94	81	79	
6	75	83	80	72	

VLOOKUP, HLOOKUP Functions

Lookup tables are useful when you want to compare a particular value to a set of values, and depending on where your value falls, assign a given “answer.” For example, you might have a tax table that shows, for any gross adjusted income, what the corresponding tax is. There are two versions of lookup tables, vertical (VLOOKUP) and horizontal (HLOOKUP). Since they are virtually identical except that vertical goes down whereas horizontal goes across, I’ll only discuss the VLOOKUP function. Besides, VLOOKUP appears to be used much more frequently than HLOOKUP.

The VLOOKUP function takes three arguments plus an optional fourth argument: (1) the value to be compared, (2) a lookup table, with the values to be compared against always in the leftmost column, and (3) the column number of the lookup table where you find the “answer”; and (4) TRUE or FALSE (which is TRUE by default if omitted). Because the VLOOKUP function is often copied down a column, it is usually necessary to make the second argument an absolute reference, and this is accomplished most easily by giving the lookup table a range name such as LookupTable. (Range names are always treated as absolute references.)

The most common use of a lookup table is when the values in the first column (the comparison column) are sorted in ascending order. (Then the fourth argument can be omitted.) Let’s say you want to assign letter grades to students based on a straight scale: below 60, an F; at least 60 but below 70, a D; at least 70 but below 80, a C; at least 80 but below 90, a B; and 90 or above, an A. The spreadsheet sample below shows how you would set this up. The comparison column in the lookup table starts at 0 (the lowest grade possible), then records the cutoff scores 60 through 90. The lookup table in the range E2:F6 is range-named LookupTable.

The formula in cell C2 is =VLOOKUP(B2,LookupTable,2), and this is copied down column C. This formula compares the value in B2 (67) to the values in column E and chooses the largest value less than or equal to it. This is 60. Then because the last argument in the VLOOKUP function is 2, the score reported in C2 comes from the *second* column of the lookup table next to 60, namely, D.

Student	Score	Grade	Lookup table	
1	67	D	0	F
2	72	C	60	D
3	77	C	70	C
4	70	C	80	B
5	66	D	90	A
6	81	B		
7	93	A		
8	59	F		
9	90	A		

To use a VLOOKUP function:

Create a lookup table with at least two columns, where the values in the first column are in ascending order, and (for best results) give the table range a range name. Then enter the formula =VLOOKUP(value,lookup table,column #), as described above.

Try it! Enter a lookup table in columns E and F, and VLOOKUP functions in column C. Assume there is a quantity discount pricing scheme: for orders less than 300 units, the unit price is \$3; for orders of at least 300 units but less than 400, the unit price is \$2.50; for orders of 400 units or more, the unit price is \$2. (Scroll to the right to see the correct answer.)

Order #	Units sold	Total cost				
1	373					
2	475					
3	459					
4	441					
5	238					
6	349					
7	344					
8	203					

There are times where the first column of the lookup table are *not* sorted in ascending order. This is still allowable, but then you need to include the fourth argument with value FALSE in the VLOOKUP function. In this case, VLOOKUP will look for an exact match and will return an error if doesn't find an exact match.

Try it! Use a VLOOKUP function in column C to find the gradepoints for each student. Note that the grades in column E of the lookup table are in the "natural" order, but they are not in Excel's A-Z sort order. Therefore, FALSE must be entered as the fourth argument in the VLOOKUP function.

Student	Grade	Gradepoints		Grade	Gradepoints	
Adams	B			A	4	
Davis	A-			A-	3.7	
Edwards	C+			B+	3.3	
Johnson	B-			B	3	
Myers	B+			B-	2.7	
Smith	A			C+	2.3	
Thomson	C			C	2	
				C-	1.7	

Paste Function (fx) Button

If you haven't used this button, located just to the left of the formula bar, you should give it a try. I like to call it the *function wizard*. It not only lists all of the functions available in Excel (by category), but it also leads you through the use of them. As an example, suppose you know there is an Excel function that calculates payments on a loan, but you're not sure what its name is or how to use it. You could proceed as follows.

To use the function wizard:

Select a blank cell where you want the function to go. Press the f_x button and click on the category that seems most appropriate (Financial in this case). Scan through the list for a likely candidate and select it (try PMT). At this point you can get help, or you can press the OK button and enter the appropriate arguments for the function (interest rate, term, and principal, the latter expressed as a negative number).

Try it! Use the function wizard to help you determine the function in cell B6. Use the range names in cells B3 through B5 for improved readability. (Scroll to the right to see the correct formula.)

Payments for Mr. Jones, who just bought a new car	
Amount financed	\$15,000
Annual interest rate	8.90%
Term (number of months financed)	36
Monthly payment	

Data Tables

Data tables, also called what-if tables, allow you to see very quickly how one or more outputs change as one or two key inputs change. There are two types of data tables: one-way tables and two-way tables. A one-way table has one input and *any* number of outputs. A two-way table has two inputs but only *one* output. I'll demonstrate both types.

To illustrate, suppose Mr. Jones buys a new car for \$20,000, makes a \$5,000 down payment, and finances the remaining amount over the next 36 months at an 8.5% annual interest rate. There are at least two outputs that might be of interest: the monthly payment and the total interest paid through the duration of the loan. These are affected by at least two inputs: the amount of the down payment and the annual interest rate.

Let's first look at a simple one-way data table, which illustrates how a *single* output, monthly payment, varies as the annual interest rate varies. This is shown in the following spreadsheet.

Price of car	\$20,000	Data table	
Down payment	\$5,000	Interest rate	Monthly payment
Amount financed	\$15,000		\$473.51
Annual interest rate	8.50%	8.00%	\$470.05
Number of payments	36	8.25%	\$471.78
		8.50%	\$473.51
Monthly payment	\$473.51	8.75%	\$475.25
Total interest paid	\$2,046.47	9.00%	\$477.00

To create the above table (which I already did):

Enter a link for the output in cell E3. Because the monthly payment was calculated with the PMT function in cell B7, simply enter **=B7** in cell E3. Then, starting in cell D4, enter any sequence of interest rates. Select the entire table—the range D3:E8. Next, select the Data menu, and in the Data Tools group, select Data Table from the What-If Analysis dropdown. Finally, enter B4 as the column input cell. There is no row input cell, so leave it blank.

It is important to understand what happens when you do this, so read this paragraph several times until it sinks in. Excel takes each interest rate in column D, substitutes it into the column input cell you designated (cell B4), recalculates the formula in cell E3 (the one I colored gray for emphasis) with this new interest rate, and records the answer in the data table. You use a *column* input cell because the possible inputs (interest rates) are listed in a column.

It is also possible to capture *multiple* outputs in a one-way data table. An example appears below, where the single input is still the interest rate, but there are two outputs: monthly payment and total interest paid. This table is formed exactly as before except that the table range is now D3:F8.

Price of car	\$20,000		Data table		
Down payment	\$5,000		Interest	Monthly payment	Total interest
Amount financed	\$15,000			\$473.51	\$2,046.47
Annual interest rate	8.50%		8.00%	\$470.05	\$1,921.64
Number of payments	36		8.25%	\$471.78	\$1,983.98
			8.50%	\$473.51	\$2,046.47
Monthly payment	\$473.51		8.75%	\$475.25	\$2,109.09
Total interest paid	\$2,046.47		9.00%	\$477.00	\$2,171.86

Try it! Create a one-way data table that shows the monthly payments and the total interest paid for each term (value in cell B5) from 12 to 48 in increments of 12. (Scroll to the right for the correct answer.)

Price of car	\$20,000				
Down payment	\$5,000				
Amount financed	\$15,000				
Annual interest rate	8.50%				
Number of payments	36				
Monthly payment	\$473.51				
Total interest paid	\$2,046.47				

Two-way tables allow you to vary *two* inputs, one along a row and one along a column, and capture a *single* output in the body of the table. The following spreadsheet illustrates this, where the annual interest rate and the amount of the down payment both vary, and the single output is the monthly payment.

Price of car	\$20,000		Data table of monthly payment			
Down payment	\$5,000		\$473.51	\$4,000	\$5,000	\$6,000
Amount financed	\$15,000		8.00%	\$501.38	\$470.05	\$438.71
Annual interest rate	8.50%		8.25%	\$503.23	\$471.78	\$440.33
Number of payments	36		8.50%	\$505.08	\$473.51	\$441.95
			8.75%	\$506.94	\$475.25	\$443.57
Monthly payment	\$473.51		9.00%	\$508.80	\$477.00	\$445.20
Total interest paid	\$2,046.47					

To create the above table (which I already did for you):

Enter the formula **=B7** for the single output in the upper left corner, cell D2, of the data table. (Again, I colored this cell gray for emphasis.) Enter any sequence of down payments to the right of this and any sequence of interest rates below this. Next, select the Data menu, and in the Data Tools group, select Data Table from the What-If Analysis dropdown. Finally, enter B2 as the row input cell, and enter B4 as the column input cell.

Note that B2 is the *row* input cell because various down payments are entered in a row (E2:G2). Similarly, B4 is the *column* input cell because interest rates are entered in a column. Excel substitutes each down payment into cell B2, each interest rate into cell B4, calculates the formula in cell D2, and records the answer (monthly payment) in the body of the table.

Try it! Create a two-way data table that shows the total interest paid for each term from 12 to 48 in increments of 12, and each down payment from \$3000 to \$5000 in increments of \$1000. Put down payments along the side, terms along the top. (Scroll to the right for the correct answer.)

Price of car	\$20,000				
Down payment	\$5,000				
Amount financed	\$15,000				
Annual interest rate	8.50%				
Number of payments	36				
Monthly payment	\$473.51				
Total interest paid	\$2,046.47				

Working with Text

Most people use Excel functions to calculate numbers. However, Excel also provides several very useful functions for manipulating text. These can be *real* time savers. This section illustrates some of the most common ways to manipulate text in Excel.

Parsing Text with Text Functions

Suppose, for example, that someone has given you a spreadsheet such as the one below. Each cell in column A has a person's name: last name, then a comma and a space, and then first name. Your job is to *parse* (that is, *separate*) these names as indicated for the name in cells B1 and C1. That is, column B should have all of the first names and column C should have all of the last names. Unfortunately, imagine that there are 5000 names in column A. What do you do? If you just start typing, you'll be at for a *long* time (and you'll undoubtedly make mistakes).

Jones, Bob	Bob	Jones			
Davis, Stephen					
Thompson, Andy					
Wilson, John					
Fredericks, Kathy					
Williams, Karen					
Smith, Tom					
Jennings, Peter					
Benson, Ted					
Samson, Jason					

Fortunately, there is a *much* better way. Logically, we observe that every name has a comma and a space. The first name is what comes after the comma and space, and the last name is what

comes before it. For the name Jones, Bob, there are 10 characters (including the comma and space, and the comma and space are characters 6 and 7. So the first name is the rightmost 3 characters and the last name is the leftmost 5 characters. In any parsing operation, this is the first and probably the most crucial step: find a pattern.

FIND, LEN, RIGHT, LEFT, and MID Functions

To do the parsing, we use Excel's FIND, LEN, RIGHT, LEFT, and MID functions. (Actually, we don't need the MID function for this particular example, but it is necessary for many similar examples.)

To use the FIND function:

The FIND function has the syntax =FIND(*subtext*,*text*,*n*), where *text* is the original text, *subtext* is a piece of text to be found in *text*, and *n* is an optional integer argument. This function returns the position of the *first* occurrence of *subtext*, if any. If there are no occurrences, it returns an error. If the argument *n* is missing, it starts the search at character 1; otherwise, it starts the search at character *n*.

Here are a few examples:

- The formula =FIND(“,”;“Jones, Bob”) returns 6, because the comma is the sixth character in Jones, Bob.
- The formula =FIND(“nes”;“Jones, Bob”) returns 3, because the subtext “nes” begins in the third character of Jones, Bob.
- The formula =FIND(“o”;“Jones, Bob”) returns 2, because the *first* occurrence of the letter “o” occurs in the second character of Jones, Bob. However, the formula =FIND(“o”, “Jones, Bob”,3) returns 9, because the first occurrence of “o” starting from character 3 is character 9.
- The formula =FIND(“t”;“Jones, Bob”) returns an error, because the letter “t” isn't found in Jones, Bob.

To use the LEN function:

Enter the formula =LEN(*text*). This returns the number of characters in *text*.

To use the RIGHT function:

Enter the formula =RIGHT(*text*, *n*), where *n* is an integer. This returns the rightmost *n* characters in *text*. (If *n* is greater than LEN(*text*), it return the entire *text*.)

To use the LEFT function:

Enter the formula =LEFT(*text*, *n*), where *n* is an integer. This returns the leftmost *n* characters in *text*. (If *n* is greater than LEN(*text*), it return the entire *text*.)

To use the MID function:

Enter the formula =MID(*text*,*s*,*n*), where *s* is an integer and *n* is an optional integer. If *n* is missing, this returns all of *text* from character *s* on. (*s* stands for start.) Otherwise, it returns *n* characters of *text*, starting with character *s*.

Here are some examples:

- The function =MID(“Freddie”,5) returns “die”.
- The function =MID(“Freddie”,5,2) returns “di”.

- The function =MID("Jones, Bob",FIND(",","Jones, Bob")+2) returns "Bob". (See why?)

Between these five functions—FIND, LEN, RIGHT, LEFT, and MID—you can do some powerful text manipulation in Excel, and this can save hours of tedious typing. You will be the hero at your office!

We now solve the above name-parsing problem in a series of steps, as indicated below. In cell B1, we use LEN to get the length of the text in cell A1. In cell C1, we use FIND to find the position of the comma. In cell D1, we add 1 to find the position of the space after the comma. We then find the first name with the RIGHT function, using the rightmost (10-7) characters. Similarly, we find the last name with the LEFT function, using the leftmost (6-1) characters. The beauty of this method is that we have to do it only once, in row 1. Then we can copy all of the formulas down for the rest of the names, even if there are 5000 of them! (Try it below. Copy the formulas down.)

Jones, Bob	10	6	7	Bob	Jones	
Davis, Stephen						
Thompson, Andy						
Wilson, John						
Fredericks, Kathy						
Williams, Karen						
Smith, Tom						
Jennings, Peter						
Benson, Ted						
Samson, Jason						

Try it! This is a bit more difficult than the name-parsing exercise above. Now each name is a first name, then a middle initial and a period, and then a last name. The goal is to have three columns with the first name, the middle initial (no period), and the last name. (Scroll to the right for a possible solution. This solution finds each piece of the name in a single formula. You might prefer to break it up into several formulas.)

Bob E. Jones						
Stephen C. Davis						
Andy T. Thompson						
John F. Wilson						
Kathy C. Fredericks						
Karen D. Williams						
Tom T. Smith						
Peter F. Jennings						
Ted R. Benson						
Jason E. Samson						

Concatenating Text

The opposite of parsing text is concatenating text. To parse text, we start with a long piece of text and separate it into pieces. To concatenate, we start with two or more pieces of text and "string them together." For example, we might parse "Bob Jones" into "Bob" and "Jones". In the other direction, we might concatenate "Bob" and "Jones" into "Bob Jones". The key to concatenating in Excel is the ampersand symbol, &.

To concatenate the text in cells A1 and B1, for example, enter the formula =A1&B1 in cell C1. Literal text can also be included, such as =A1&" "&B1. In this case, if A1 contains "Jones" and B1 contains "Bob", then C1 will contain "Jones, Bob". This concatenates the last name, a literal comma and space, and the first name.

Try it! Concatenate the first names, middle initials, and last names in columns A, B, and C so that full names of the form Jones, Bob E. appear in column D.

Bob	E	Jones				
Stephen	C	Davis				
Andy	T	Thompson				
John	F	Wilson				
Kathy	C	Fredericks				
Karen	D	Williams				
Tom	T	Smith				
Peter	F	Jennings				
Ted	R	Benson				
Jason	E	Samson				

Working with Dates

We tend to take dates and times for granted, but they appear in many, if not most, real-world business spreadsheets. They are actually quite complex to work with, and a complete treatment of dates and times could fill a long chapter of a book. I can't hope to cover all of the intricacies of Excel-based dates and times in this tutorial, but I will illustrate some of the most useful methods for working with them.

How Dates Are Stored

The first thing to realize is that a date-time value that appears in a cell as something like 3/14/2004 4:30 PM is really stored as a *number*. Excel refers to this number as the *serial* value. The serial value corresponding to this particular date-time happens to be 38060.688. The value to the left of the decimal is the number of days since a “base” date (January 1, 1900), and the decimal part indicates the fraction of time through the day, starting right after midnight. In this example 4:30 PM is 68.8% of the way from midnight to the next midnight. If the decimal part is omitted, then it is assumed to be midnight. In this case, you usually format the value so that only the date, not the time, appears.

If you happen to know the numbering system for dates (that is, you somehow know that 38060 corresponds to 3/14/2004), you could enter the number 38060 in a cell and then *format* it to look like a date. Usually, however, you enter a date like 3/14/2004 in a cell. In fact, if you enter the date in one of several formats Excel recognizes (including 3/14/2004, March 14, 2004, 3-14-04, and a few others), Excel will know that it is a date, but underneath the hood, it will store the value as 38060.

The same goes for date-times, such as 3/14/2004 4:30 PM, or just times, such as 4:30 PM. If you enter either of these in a format that Excel recognizes as a date-time or just a time, Excel will recognize that these as date-times or times, but it will store them as numbers (68030.688 or 0.688).

By the way, if you enter a date-time in a format that Excel doesn't recognize, you're entirely out of luck. It will simply be stored as a label, with no inherent “date-time” meaning. And Excel can be fussy. For example, it doesn't understand 3/14/2004 4:30PM, simply because there is no space between 4:30 and PM! (I wish I were kidding, but I'm not.) On the other hand, Excel is smart about dates. It will recognize that 9/31/1999 and 2/29/2003 are not dates—regardless of how they are formatted.

Try it! The following spreadsheet includes some date-times, dates (no times specified), and times (no dates specified), formatted in several allowable ways. Format these as numbers

with 3 decimals to see the serial values. Then reformat them as date-times (in formats you can choose).

3/14/04 1:30 PM						
12/31/1901						
3:50						
12:00 PM						
29-Feb-96						
9/24/46 14:45						

How Excel Deals with the Y2K Problem

Do you remember the Y2K problem that worried everyone in late 1999? The problem was that a large number of dates, in all kinds of systems, listed 2-digit years, such as 85 instead of 1985. How would systems know whether 05 means 1905 or 2005? Fortunately, the predicted disaster didn't happen, but there is still a Y2K problem because of the ambiguity it can create. Excel handles it (at least for now) by interpreting a 2-digit year less than 30 as being in the 21st century, and all others as being in the 20th century. For example, it interprets 3/14/17 as March 14, 2017 and 5/20/72 as May 20, 1972. Of course, the easiest way to avoid the ambiguity is to use 4-digit years whenever you enter dates. This is the main lesson of Y2K!

Try it! In the spreadsheet above, enter a date such as 3/14/29 (the last two digits less than 30). Then enter another date such as 5/20/30 (the last two digits greater than or equal to 30). Watch how Excel transforms the year automatically. Do you see how this could be a problem if you really meant 3/14/29 to be the birthday of a person born in 1929?

NOW, TODAY Functions

Excel has two useful functions for entering the current date, or the current date and time, automatically, TODAY and NOW.

To use the TODAY function:

Enter the formula =TODAY() in any blank cell. (The blank parentheses *are* necessary.) It returns the current date. If you open this spreadsheet tomorrow, it will list tomorrow's date, that is, it will update each day.

To use the NOW function:

Enter the formula =NOW() in any blank cell. (Again, the blank parentheses *are* necessary.) It returns the current date and time.

Try it! Enter the TODAY and NOW functions in column B of the following spreadsheet. Then format result of NOW as time only—no date. After you've worked a few other exercises below, come back to this spreadsheet and see whether the time has changed. (You might have to force a recalculation to see the change. You can do this easily by pressing the F9 key.)

Today's date						
Current time						

YEAR, MONTH, DAY, WEEKDAY Functions

If a cell contains a date, it is easy to strip off the year, month, day, or even weekday with the YEAR, MONTH, DAY, and WEEKDAY functions. In all of the following examples, *datecell* refers to a cell that contains a date.

To use the YEAR function:

Enter the formula =YEAR(*datecell*). This returns a 4-digit year such as 1993.

To use the MONTH function:

Enter the formula =MONTH(*datecell*). This returns a month from 1 (January) to 12 (December).

To use the DAY function:

Enter the formula =DAY(*datecell*). This returns the day of the month, from 1 to the number of days in the month.

To use the WEEKDAY function:

Enter the formula =WEEKDAY(*datecell*). This returns a number from 1 to 7, where 1 refers to Sunday, 2 to Monday, and so on. If you want to get the name of the day, such as Thursday, use a VLOOKUP function, as indicated in the exercise below.

Try it! For each date in column A, return the year, month, month name, day, weekday, and weekday name in columns B–G. For the month and weekday names, use the lookup tables given. (Scroll to the right to see them.) Do you recognize these “infamous” dates? Try one more: your own birthdate.

Date	Year	Month	Month name	Day	Weekday	Weekday name
9/11/2001						
1/28/1986						
11/22/1963						
8/6/1945						
12/7/1941						
4/14/1912						

DATEDIF Function

Because dates are stored as numbers, it is possible to perform arithmetic on them. One particular arithmetic operation makes sense: subtraction. If you subtract one date from another, you find the number of days between them. You can do this by simple subtraction of cells, or you can get more control with the DATEDIF function. If you use simple subtraction, you automatically get

the difference in *days*. For example, the formula =A1-A2, where A1 contains 11/25/2004 and A2 contains 11/20/2004, gives 5. The DATEDIF function provides more options. (Evidently, DATEDIF has been in Excel for years, but its documentation in online help has been sporadic. See the interesting article at <http://www.cpearson.com/excel/datedif.htm>.)

To use the DATEDIF function:

Enter the formula =DATEDIF(*earlierdate, laterdate, interval*), where *interval* can be “y”, “m”, “d”, or a few other options not covered here. If *interval* is “y” (quotes required), this returns the number of years between the two dates. Similarly, if *interval* is “m” or “d”, it returns the number of months or number of days between the two dates.

Try it! Use simple subtraction of cells to find the number of days between the following two dates. Then use DATEDIF with the three options described above to get the number of days, months, and years between the two dates. (Note: You’ll have to reformat the answer in cell B6 as a number. Excel wants to format it as a date for some odd reason. Scroll to the right for the answers.)

Dates of last two physical exams					
First exam	4/5/1999				
Next exam	7/15/2003				
Days (subtraction)					
Days (DATEDIF)					
Months (DATEDIF)					
Years (DATEDIF)					

DATE, DATEVALUE Functions

The last two date functions I’ll mention are DATE and DATEVALUE. The DATE function takes a date, such as 3/14/2004, and return the serial value, in this case 68030. I see no obvious uses for this function. (After all, you can see the serial value if you want by formatting the date cell as a number. Besides, you probably have little interest in the actual number anyway.) I’ve mentioned it simply because you might see it in online help and think it’s important.

The DATEVALUE function, on the other hand, can be a real life-saver. It takes a date stored as *text* and returns the corresponding serial value, which you could then format as a date. Why is this important? Suppose you have a lot of data stored in some type of legacy system in your business, and you are somehow able to “dump” it into Excel. The chances are that all dates will be imported as text. That is, a cell value will look something like 5/20/2001, but Excel won’t recognize it as a date. Without the DATEVALUE function, you would have retype all of the dates! However, DATEVALUE saves you this mind-numbing work.

To use the DATEVALUE function:

Enter the formula =DATEVALUE(*datetext*), where *datetext* refers to a date stored as text. Then format the result as a date.

Try it! Create dates in column B corresponding to the *text* values in column A that happen to look like dates. (I typed an apostrophe in front of each value in column A. That’s why they are interpreted as text, not as dates.)

Dates imported from a legacy system						
Nov 25, 1997						
Nov 26, 1997						
Nov 29, 1997						
Dec 1, 1997						
Dec 3, 1997						

Statistical Functions

There are many statistical functions built into Excel. I've already discussed the simplest statistical function, the AVERAGE function. (Statisticians often call this the *mean*, or the *sample mean*, but it's simply the average of a set of numbers.) In this section I'll discuss a few other statistical functions.

MEDIAN, MODE Functions

The median is the "middle" value in a set of numbers, in the sense that half are above it and half are below it. More precisely, when the values are sorted in increasing order, it is the middle value if there are an odd number of values, and it is the average of the two middle values if there are an even number of values. Fortunately, Excel takes care of the details with the MEDIAN function.

To use the MEDIAN function:

Enter the formula =MEDIAN(*datarange*), where *datarange* is any range that contains a set of values.

Note that if *datarange* contains labels or blank cells, these are ignored when calculating the median. This is true for the other statistical function listed below. (Of course, it is also true for the AVERAGE function.)

Try it! Find the median of the numbers in column A (which extend well below row 10). Is it the same as the average? Why or why not? (Scroll to the right to see the answers.)

33	Median:				
3	Average:				
0	Mode:				
1					
4					
14					
5					
41					
3					
5					

The mode is the most likely value, the one that occurs most often. If no value occurs more than once or twice, then the mode isn't very useful. However, if a clothing retailer records the size of each dress sold, then the mode is quite useful. It is the best-selling dress size. Excel's MODE function calculates the mode.

To use the MODE function:

Enter the formula =MODE(*datarange*), where *datarange* is any range that contains a set of values.

Try it! Calculate the mode of the set of values in the above spreadsheet. How many times does this value occur? (Again, answers are to the right.)

PERCENTILE, QUARTILE Functions

We often like to rank the values in a data set in some way. For example, if you can the SAT exam to get into college, you essentially learn your ranking by a percentile. If you are at the 85th percentile, you know that 85% of all people scored lower than you, and only 15% scored higher. Excel has two useful functions for finding this type of information: PERCENTILE and QUARTILE.

The QUARTILE function is probably the easiest to understand. Imagine that you have 1000 scores. You sort them from low to high and then divide them into 4 sets of 250 scores each, where the first set contains the smallest scores, the second set contains the next smallest scores, and so on. The first, second, and third quartiles are the breakpoints between these sets. For example, 25% of the scores are below the first quartile, and 75% are above it. (The second quartile is the median.)

To use the QUARTILE function:

Enter the formula =QUARTILE(*datarange*,*k*), where *k* is 1, 2, or 3. For example, if you enter 3 for *k*, you get the third quartile.

Try it! Find all three quartiles for the following data set (which extends well below row 10). Enter the formulas for these quartiles, labeled Q1, Q2, and Q3, in column D. (Scroll to the right to see the answers.)

		Quartiles			Percentiles
86.23		Q1			P01
89.69		Q2			P05
76.12		Q3			P10
129.72					P25
82.41					P50
106.6					P75
80.99					P90
73.26					P95
112.7					P99
46.21					

The PERCENTILE function takes any percentage (expressed as a decimal number) and returns the value such that the given percentage of all values is below this value. For example, if we use the percentage 10% (0.10) and the PERCENTILE function returns 45, then we know that 10% of all values in the data set are below 45 and 90% are above it.

To use the PERCENTILE function:

Enter the formula =PERCENTILE(*datarange*,*pct*), where *pct* is any percentage expressed as a decimal, such as 0.10 for 10%.

Try it! Find the requested percentiles in the above spreadsheet, and place your formulas in column G. For example, P10 indicates the 10th percentile. Do you see the relationship between the quartiles and certain percentiles? (Again, answers to the right.)

Note that the QUARTILE and PERCENTILE functions are typically used on data sets with a large number of observations. If they are used on a very small data set, you might be surprised at the results. In fact, you might even disagree with them. For example, if the data set has the 5

values 10, 20, 30, 40, and 50, Excel's PERCENTILE function indicates that the 10th percentile is 14. This is certainly not a very intuitive answer, and you might disagree with it, but we might ask why someone would want the 10th percentile of a 5-value data set in the first place!

RANK, LARGE, SMALL Functions

If we simply want to rank a set of values from 1 to N , where N is the number of values, Excel offers the RANK function. We can either rank them from highest to lowest or vice versa. Of course, an alternative is to sort the numbers. If for any reason we don't want to sort, then we can use the RANK function.

To use the RANK function:

Enter the formula `=RANK(value,datarange,order)`, where *value* is a single value in the *datarange*, and *order* is an optional argument. If *order* is missing (or 0), rank 1 corresponds to the largest value; if *order* is 1 (or any other nonzero value), rank 1 corresponds to the smallest value.

Typically, we start with a data set in a column. Then we use the RANK function to rank the first value and copy it down (with *datarange* absolute) to rank all of the others. If two or more values are tied, they share a rank.

Try it! Rank the values in A1:A10 from largest to smallest in column B. Then rank them from smallest to largest in column C. How does RANK deal with the tied 55's? (Scroll to the right for answers.)

55						
66						
76						
34						
55						
68						
75						
44						
38						
58						

The LARGE and SMALL functions are less well-known Excel functions, but they can be useful. Let's say we have a data set with 250 values and we want the 10 largest values. Then we can find these with the LARGE function, using arguments 1 through 10. Similarly, we can use the SMALL function to find the 10 smallest values.

To use the LARGE function:

Enter the formula `=LARGE(datarange,k)`, where k is any integer from 1 to the number of data values. This returns the k th largest value.

To use the SMALL function:

Enter the formula `=SMALL(datarange,k)`, where k is any integer from 1 to the number of data values. This returns the k th smallest value.

Typically, we would enter a sequence of values of k in some column, enter the LARGE or SMALL formula for the first of them, and then copy down.

Try it! For the data set in column A, which extends well below row 10, find the 10 largest values and 10 smallest values in columns D and E. Use the values of k in column C. (Scroll to the right for the answers.)

679.02		1				
1208.42		2				
1635.72		3				
729.63		4				
1036.08		5				
716.62		6				
972.92		7				
1175.68		8				
1271.38		9				
1048.93		10				

STDEV, VAR Functions

Excel has a number of other statistical functions, including those for measuring the amount of variability in a set of numbers. The most widely used measures of variability are *variance* and *standard deviation*, and Excel implements these with the VAR and STDEV functions. (These functions calculate the *sample* variance and standard deviation. There are similar functions, VARP and STDEVP, for calculating the *population* variance and standard deviation. For large data sets—over 100 values, say—the sample and population measures are practically the same.)

The variance is essentially the average of the squared deviations from the average, and the standard deviation is the square root of the variance. Each measures variability, but standard deviation is easier to understand because it is expressed in the same units as the data—dollars, for example. Besides, many data sets have the property that about 68% of the data are within 1 standard deviation of the average, about 95% are within 2 standard deviations of the average, and virtually all of the data are within 3 standard deviations of the average.

To use the VAR function:

Enter the formula =VAR(*datarange*).

This is equivalent to =(STDEV(*datarange*))^2, the square of the standard deviation

To use the STDEV function:

Enter the formula =STDEV(*datarange*).

This is equivalent to =SQRT(VAR(*datarange*)), the square root of the variance.

Try it! Find the average, variance, and standard deviation of the data in column A. Then use COUNTIF to find the percentage of values that are within 1 standard deviation of the average. Is it about 68%? (Scroll to the right to see the answers.)

1034.44		Average				
877.35		Variance				
794.42		Standard deviation				
828.45						
765.9						
1088.95						
786.05						
1523.02						
920.48						
777.23						

Financial Functions

Excel has a wide assortment of financial functions that are used by thousands of financial analysts every day. I'll illustrate the most common of these. If you're going to be a financial analyst, you should become familiar with the others!

PMT Function

The PMT function is usually used to find the monthly payment for a car loan or a home mortgage loan. The inputs are typically an annual interest rate, a term (number of months financed), and the amount borrowed (the principal). The PMT function finds the amount you have to pay each month of the term. Part of this payment is principal and part is interest. At the end of the term, you will have paid just enough to pay off the entire loan.

To use the PMT function:

Enter the formula `=PMT(interestrate,term,principal)`. For technical reasons, if you want PMT to return a *positive* value, you should enter the principal as a *negative* number. The interest rate argument should be the monthly rate (assuming you're paying monthly), which is the annual rate divided by 12.

Try it! Suppose you take out a \$30,000 loan for a new car when the annual interest rate is 6.75% and the term of the loan is 36 months. Calculate your monthly payment in cell B4. Then use a data table (the perfect place for one!) to see how your monthly payment varies for terms of 24, 36, 48, or 60 months. (Scroll to the right to see the answers.)

Principal	\$30,000		Data table for payment versus term			
Annual interest	6.75%					
Term	36					
Payment						

NPV, XNPV Functions

Much of the theory of finance is about the time value of money. Basically, a dollar earned in the future is less valuable than a dollar earned today because the dollar earned today can earn interest. To account for this, we *discount* future inflows or outflows to get their *present values* now. If we have a future stream of inflows and/or outflows, then the sum of all of their present values is called the *net present value*, usually abbreviated NPV. If this stream is incurred at regular time intervals—at the end of each year, say—then we can use Excel's handy NPV function to calculate the NPV of the stream.

To use the NPV function:

Enter the formula `=NPV(rate,cashstream)`. Here, *rate* is the discount rate, which is essentially the rate the company believes it can make on its money, and *cashstream* is a stream of cash inflows or outflows that occur at the ends of successive periods, starting at the end of period 1.

Note that if there is a cash inflow or outflow right away, at the beginning of period 1, it should be entered *outside* the NPV function. The reason is that it doesn't need to be discounted.

Try it! Assume a company pays \$100,000 at the beginning of year 1 to get into an investment. It then receives the cash inflows at the ends of years 1 through 5 shown below. What is the NPV of this investment (inflows minus outflows) with the given discount rate? Does the NPV increase or decrease as the discount rate increases? Why? (Scroll down to see the answer.)

Discount rate	12%					
Initial payment	\$100,000					
Cash inflows at ends of years						
Year	1	2	3	4	5	
Inflow	\$20,000	\$35,000	\$50,000	\$45,000	\$35,000	
NPV						

If a company incurs cash inflows or outflows at irregular times, such as January 15, then May 30, then July 1, and so on, the NPV function cannot be used. You could go back to your finance book to see how to discount future payments directly, but there is an easier way, using the little-known XNPV function. This function is actually part of the Analysis ToolPak that ships with Excel. However, you might not have it loaded. To check, click on the Office button, then Excel Options, then Add-Ins, and then Go. If the Analysis ToolPak item isn't checked, check it. Then you can use the XNPV function.

To use the XNPV function:

Enter the formula `=XNPV(discountrate,cashvalues,dates)`, where *discountrate* is the same as with NPV and *cashvalues* and *dates* are a stream of pairs of cash values and the dates when they are incurred.

Typically, the first cash value will be an outflow and the rest will be inflows. In this case, the initial outflow (investment) *should* be included in the XNPV function, and the NPV will be discounted back to the date of this initial payment. This payment should be entered as a negative number.

Try it! First make sure the Analysis ToolPak is loaded. Then find the net present value of the stream of cash inflows shown below, where the first is really a payment at the beginning of 2005. (Scroll down to see the answer.)

Discount rate	12%				
Stream of cash flows					
Date	1-Jan-05	25-May-05	15-Sep-05	15-Jan-06	30-Jun-06
Cash	(\$85,000)	\$22,500	\$31,000	\$65,500	\$55,000
NPV					

IRR Function

Companies often have investment opportunities where they pay initially and then get returns in the future. When discussing NPV, I illustrated in the "Try it!" exercise that the NPV of such an investment decreases as the discount rate increases. In fact, for large enough discount rates, the

NPV will typically become negative, meaning that the future returns are not enough to offset the initial cost. The discount rate at which NPV changes from positive to negative is called the internal rate of return, or IRR. Specifically, the IRR is the discount rate at which NPV equals 0. Companies are interested in IRR for the following reason. They typically have a *hurdle rate* that they use to discount potential investments. If the NPV of an investment is positive, discounted at the hurdle rate, it is worth pursuing. If it is negative, it is not worth pursuing. Stated equivalently, the investment is worth pursuing only if its IRR is greater than the company's hurdle rate.

Fortunately, Excel has an IRR function that calculates an investment's IRR—a tough calculation. As with the NPV function, we assume the investment is structured so that there is an initial cash payment at the beginning of year 1 and then regular cash returns at the ends of years 1, 2, and so on.

To use the IRR function:

Enter the formula =IRR(*cashstream*,*rateguess*), where *cashstream* is a stream of cash flows, where the first (the initial payment) should be negative, and *rateguess* is an initial guess for the IRR.

This seems strange. Why should you have to guess at the answer? It is because Excel calculates the IRR iteratively, starting with your guess. Your actual guess shouldn't make any difference in the final answer except in unusual cases.

Try it! Find the IRR for the following investment, using an initial guess of 15%. What does this IRR say about the attractiveness of the investment? (Scroll down to see the answer.)

Cash flows (initial value is a payment)					
(\$100,000)	\$35,000	\$45,000	\$40,000	\$25,000	\$10,000
IRR					

Lookup and Reference Functions

I have heard an Excel expert make the claim that the VLOOKUP function, discussed earlier, is the most useful Excel function of all. Whether or not you believe this claim, the VLOOKUP and HLOOKUP functions are just two of some very useful *reference* functions in Excel's arsenal. If you click on the *f_x* button for online help on Excel functions, you will see a Lookup and Reference category of functions that includes VLOOKUP, HLOOKUP, and several others. I will illustrate a few of these others here. Although they are arguably not as useful as VLOOKUP, they can definitely come in handy in certain situations.

INDEX Function

The INDEX function is useful for finding the value in a particular cell of a rectangular range. You access this value by the *index* of the cell you want. Specifically, you specify a row index and a column index. For example, if the range has 10 rows and 20 columns, the index 4,8 indicates row 4 and column 8 of the range.

To use the INDEX function:

Enter the formula =INDEX(*range,r_index,c_index*), where *r_index* is an integer (for row) and *c_index* is an integer (for column).

For example, the formula =INDEX(D11:F20,4,2) refers to the value in cell E14, the cell in the fourth row and second column of the range.

If *range* is a single-column range, then the *c_index* argument can be omitted. In that case, =INDEX(*range,r_index*) returns the value in row *r_index* of the range. For example, the formula =INDEX(A11:A20,3) refers to the value in cell A13, the third cell in the range. Similarly, if *range* is a single-row range, only *c_index* needs to be included. For example, the formula =INDEX(B10:E10,3) refers to the value in cell D10, the third cell in the range.

Try it! Find the unit shipping cost from Plant2 to City3 with a formula in cell B6. Then enter a formula in cell B10 that finds the unit shipping cost from the plant to the city specified in cells B8 and B9. This formula should respond appropriately to whatever plant and city indexes are entered in cells B8 and B9. (Scroll to the right for the answers.)

Unit shipping costs	City1	City2	City3			
Plant1	1.25	1.35	1.55			
Plant2	1.15	1.45	1.25			
Plant3	1.35	1.45	1.15			
Unit shipping from Plant2 to City3						
Plant	1					
City	2					
Unit shipping cost						

MATCH Function

The MATCH function is handy for finding a cell in a range that matches a given value. It is often used in the following situation. Suppose you have decision variable such as order quantity that needs to be chosen so as to maximize some profit. You enter some formulas that link order quantity to the profit. Then you create a data table that finds the profit for a number of possible order quantities. The MATCH function lets you locate the cell in the data table with the smallest profit.

To use the MATCH function:

Enter the formula =MATCH(*value,range,FALSE*). This returns the index of the cell in *range* that matches *value*. For example, if the match occurs in the third cell of the range, this returns 3.

The last argument, FALSE, indicates that we want an *exact* match. (Actually, FALSE can be replaced by 0 with the same effect.) If there is no exact match, the formula returns an error. The MATCH function can also look for an inexact match by using third argument TRUE. I'll let you look this up in online help.

The following spreadsheet indicates how MATCH can be used in conjunction with INDEX. You can imagine that a profit model has led to the table shown, where each order quantity listed leads to the corresponding profit. For this small example, it is obvious that the largest expected profit is \$5,640, which corresponds to an order quantity of 300. However, if inputs (not shown) that drive the profit model change, the best order quantity and the corresponding profit could change. We want formulas in cells B9 and B10 to show and the best profit and best order quantity regardless of where they appear in the table.

Table of profit versus order quantity					
Order quantity	Profit				
200	\$5,325				
250	\$5,430				
300	\$5,640				
350	\$5,565				
400	\$5,235				
Best profit	\$5,640				
Best order quantity	300				

The formula in cell B10 is =INDEX(A3:A7,MATCH(B9,B3:B7,0)). The whole purpose of the MATCH function here is to provide the row index for the INDEX function. That is, we know the best order quantity is one of the values in the range A3:A7, and we need to know which one. So we find a match for the *maximum* value in B3:B7 to the values in B3:B7. In this case, the maximum is the third profit, so the formula is equivalent to =INDEX(A3:A7,3).

Try it! The table below shows the profit for each order quantity and each unit stockout cost (along the top). Enter formulas in the range B9:E9 to find the best order quantity for each unit stockout cost. (Scroll down to see the answer.)

Profit for various order quantities (column A) and unit stockout costs (row 3)					
	\$2	\$3	\$4	\$5	
200	\$5,235	\$5,175	\$5,025	\$4,935	
250	\$5,465	\$5,285	\$5,190	\$5,085	
300	\$5,355	\$5,325	\$5,265	\$5,100	
350	\$5,215	\$5,195	\$5,130	\$5,125	
400	\$5,145	\$5,085	\$5,015	\$5,005	
Best order quantities					

OFFSET function

The OFFSET function allows you to reference a range (or single cell) relative to another cell. It is hard to appreciate unless you see some examples, so I will present two below.

To use the OFFSET function:

Enter the formula =OFFSET(*cell,r_offset,c_offset,height,width*). Here, *r_offset* and *c_offset* are integers that can be positive, negative, or zero, and *height* and *width* are optional positive integers. If either *height* and *width* are missing, they default to 1. This formula returns a reference to a range that has *height* rows and *width* columns. To find its upper left cell, start at *cell*, move *r_offset* rows down (if positive) or up (if negative), and move *c_offset* columns to the right (if positive) or the left (if negative).

For example the formula =OFFSET(A1,2,3,4,1) returns a reference to the range D3:D6. It is a range with 4 rows and 1 column, and its upper left cell, D3, is offset from cell A1 by going 2 rows down and 3 columns to the right. As another example, the formula =OFFSET(F4,0,-3) refers to a single cell (because the last two arguments are missing), and this single cell is 3 columns to the left of F4, namely, C4.

A good example of the OFFSET function appears in the spreadsheet below. Our company sells to a retailer, and the retailer pays a certain number of months later, as indicated by the payment delay (in months) in cell B1. The sales occur as indicated in row 4, but the receipts from the retailer occur 2 months later, as indicated in row 5.

Payment delay	2								
Month	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun
Sales	3200	4600	5500	2500	3300	4200	2900	4500	2750
Receipts				4600	5500	2500	3300	4200	2900

Without the OFFSET function, we could simply put links in the Receipts row that point to sales 2 months earlier. For example, the formula for receipts in January would be =C4. But what if the retailer decides to delay payments by 3 months instead of 2? Then we would need to fix the links in the Receipts row. However, a clever use of OFFSET avoids this updating of links. We use the formula =OFFSET(E4,0,-\$B\$1) in cell E5 for January and then copy it across row 5 for the other months. This formula says to start in the January sales cell and go a certain number of cells to the left (because of the minus sign), where this number is specified in cell B1. To see how it works, try changing the value in cell B1 to 3 or 1. (What would go wrong if you changed this delay to 4 months? How would you fix it?)

Try it! Suppose that a manufacturing company buys raw materials from a supplier. If the cost in any month is \$x, the company pays 40% of this cost 1 month from now and the other 60% 2 months from now. Use the OFFSET function to calculate the payments made in January through August, based on the costs through July. These calculated payments should change automatically if you change the inputs in cells B1:B2 and D1:D2. For example, make sure your formulas react correctly if you enter 2 and 3 in cells D1 and D2. (Scroll down to see the answer.)

Percent paid	40%	when:	1	month from now		
	60%		2	months from now		
Month	Oct	Nov	Dec	Jan	Feb	Mar
Cost	\$1,950	\$2,900	\$4,150	\$2,500	\$3,350	\$2,850
Payment						

Another great use of the OFFSET function is to create a *dynamic* range name, one that expands or contracts depending on the number of data values in a range. Consider the following spreadsheet, where monthly sales values are entered in column B, and the total of all sales values is calculated in cell E1. Every month, an extra sales value and its month label are appended to the list in columns A and B. We can make the formula in cell E1 automatically adjust to the appended values by using a dynamic range name.

Month	Sales		Total	\$17,030		
Jan-05	\$2,450					
Feb-05	\$3,140					
Mar-05	\$3,250					
Apr-05	\$4,510					
May-05	\$3,680					

To do this, I selected the Define Name dropdown from the Defined Names group in the Formulas ribbon. In the Name box at the top of the resulting dialog box, I entered Sales as the range name. In the Refers To box at the bottom, I entered the formula =**OFFSET(B1,1,0,COUNT(B:B),1)**. Then I clicked on OK. Finally, I entered the formula =SUM(Sales) in cell E1. Note that COUNT(B:B) refers to the count of all numeric cells in column B. So this OFFSET function refers to a range that starts 1 cell down from cell B1 and has as many rows as there are numeric values in column B. To see how it adjusts, enter a sales value for June in cell B7 and watch how the total changes automatically.

Note that COUNT(B:B) counts *all* of the numeric values in column B, so if there were some other numbers down below sales, they could mess up the logic in the OFFSET function. For example, suppose there were a numeric value in cell B100. Then, given the data above, COUNT(B:B) would return 6, so at that point, the Sales range name would refer to the range B2:B7, that is, 6 cells starting with B2. The moral is that if you want to use dynamic range names, it's best not to include "junk" in the affected columns or rows.

Try it! The following spreadsheet has sales by month and by salesperson. We want to sum sales over all months and salespeople in cell B2. Create a dynamic range name that updates when new months are added *and* when salesperson columns are added or deleted. Then use this range name to sum sales in cell B2. (Scroll across to see answer.)

Total sales						
			Smith	Jones	Wilson	Donalds
	Jan-05	\$6,700	\$7,400	\$5,800	\$9,000	
	Feb-05	\$5,800	\$8,900	\$5,500	\$7,900	
	Mar-05	\$5,000	\$7,900	\$5,900	\$8,300	
	Apr-05	\$6,700	\$5,800	\$6,000	\$9,700	
	May-05	\$9,400	\$7,800	\$6,700	\$8,400	
	Jun-05	\$6,200	\$5,800	\$9,400	\$8,500	
	Jul-05	\$8,700	\$5,300	\$6,900	\$7,500	

A great use for dynamic range names is with pivot tables. Suppose you want to base a pivot table on a data set that is likely to expand (or even contract) as time goes on. You can create a dynamic range name such as MyData for the data set. Then when the pivot table wizard asks for the data range, enter MyData. If your data set expands or contracts in the future, all you need to do is click on the Refresh pivot table button (!) to update the pivot table with the newest data. If you show this trick to your colleagues at work, you'll be a hero!