

## PART I: INTRODUCTION TO USING THE SPREADSHEET

## CHAPTER 1

## HOW TO USE EXCEL

First things first: this introductory chapter is intended for readers who have no prior experience with Excel, and only provides the minimum information necessary to use the rest of this book. Emphatically, this chapter is *not* meant to replace a spreadsheet manual; if it were, that part alone would occupy more space than that of this entire workbook. Instead, during and after using this workbook, you may be tempted to consult an Excel manual (of which there will be several in your local library and bookstore) to learn what else it can do for you – but that is up to you.

Second: this book is not intended to be read, but instead to be used while you sit at the computer keyboard, trying out whatever is described in the text. Learning to use a spreadsheet is somewhat like learning to swim, to ride a bicycle, or to paint: *you can only learn it by doing it*. So set aside a block of time (one or two hours should do for this chapter, unless you are really new to computers, in which case you might want to reserve several such sessions in order to get acquainted), make yourself comfortable, turn on the computer, and try things out as they are described in, say, the first three sections of this chapter. (If it confuses you on your first try, and there is nobody at hand to help you along, stop, do something else, and come back to it later, or the next day, but don't give up.) Then try the next sections.

In order to run Excel (or any other spreadsheet program), your computer will need an **operating system**. Here we will assume that you have Windows as the operating system on your personal computer, and that you have a compatible version of Excel. Although there are relatively minor differences between the various versions of Excel, they fall roughly into three categories. Excel versions 1 through 4 did not use VBA as their macro language, and the macros described and used in this book will therefore not run on them. The second category includes Excel 5 and Excel 95 (also called Excel version 7; there never was a version 6), which use VBA with readily accessible modules. Excel 97, Excel 98 (for the Mac), and Excel 2000 make up the third category, which has macro modules that are hidden from sight. The instructions given

in this book are specifically for the second and third categories, starting with Excel 5. While they were mostly tested in Excel 97, all versions more recent than Excel 4 will do fine for most of the spreadsheet exercises in this book. Because Excel is *backward compatible*, you can run older software in a more recent version, but not necessarily the other way around.

When you have a Macintosh, your operating system will be different, but Excel will be very similar. After all, both IBM and Mac versions of Excel were written by Microsoft. With relatively minor modifications, mostly reflecting differences between the IBM and Mac keyboards, all exercises in this book will run on the Mac, provided you have Excel version 5 or later.

In either case, whether you use an IBM-compatible PC or a Macintosh, use at least Excel version 5, because earlier versions lacked some of the more useful features of Excel that will be exploited in this book. If you have Excel 4 or earlier, it is time to upgrade.

When you are already familiar with earlier versions of Windows and Excel, you may want to use this chapter as a refresher, or scan the text quickly and then go directly to the next chapter. When you are already familiar with Windows 95 or Windows 98, and with Excel 95 or 97, you may skip this chapter altogether.

## 1.1 Starting Windows

Windows is a so-called *graphical* user interface, in which many programs, files, and instructions are shown pictorially, and in which many operations can be performed by 'pointing and clicking', an approach pioneered in the early 1970s by the Xerox Corporation, and long familiar to Macintosh users. The pointing device is usually a **mouse** or a **trackball**; for many instructions, equivalent typed commands can be used as well. We will use 'mouse' as the generic term for whatever pointing device you may have. There are often several ways to let the computer know what you want it to do. Here we will usually emphasize how to do it with the mouse, because most users find that the easiest.

In what follows we will assume that Windows and Excel have been installed in their complete, standard forms. For some applications we will also use the Solver and the Analysis Toolpak. These come with Excel, but (depending on the initial installation) may have to be loaded as an add-in.

When you start Windows, your monitor will show a screen (the **desktop**) which typically displays, on its left side, a number of pictures (**icons**), each with its own explanatory label. The bottom icon is labeled 'Start', and acts as the *on* switch of Windows. (There is no simple *off* switch, since Windows requires a more elaborate turn-off routine, which rather illogically begins with the Start button, and via the Shut Down command leads you to the Shut Down Windows dialog box, where you can choose between several options.)

## 1.2 A first look at the spreadsheet


3

Icons, such as the start label, are also called **buttons**, as if you could actually push them. Move the mouse so that the sharp point of the arrow on the screen, the **pointer**, indeed 'points to' (i.e., is *inside*) the start button, and press the left mouse button once. (Left and right depend, of course, on the orientation of the mouse. By 'left' we mean the left button when the two or three mouse buttons are pointing away from you, so that you can hold the body of the mouse with your thumb and index finger, or with the palm of your hand, while your index finger, middle finger, and ring finger can play with the buttons.) To briefly depress the left mouse button we will call to **click** the mouse; when you need to do this twice in quick succession we will call it **double clicking**, whereas briefly depressing the right mouse button we will call **right clicking**.

As soon as you have clicked the start button, a dialog box will pop up above it, showing you a number of choices. Manipulate the mouse so that the arrow points to 'Programs', which will now be highlighted, and click. A second dialog box will pop up next to the first to show you the various programs available. One of these will be Excel; click on it to start the spreadsheet. Alternatively, click on the Excel icon if the desktop shows it.

## 1.2 A first look at the spreadsheet

After displaying the Excel logo, the monitor screen will show you a rather busy screen, as illustrated in Fig. 1.2-1. The actual screen you will see may have more bars, or fewer, depending on how the screen has been configured. Please ignore such details for the moment; few if any of the instructions to follow will depend on such local variations.

 At the top of the screen is the **title bar**. In its right-hand corner are three icon buttons, to *minimize* the screen to near-zero size, to *restore* it to medium or full size, and to *close* it. To the left on the same bar you will find the Excel logo and the name of the **file** you use, where 'file' is the generic name for any unit in which you may want to store your work. Below the title bar is the **menu bar** (with such menu headings as **F**ile, **E**dit, **V**iew, **I**nsert, etc.). This is usually followed by a **standard bar** with icons (pictograms showing an empty sheet, an opening file folder, a diskette, a printer, etc.) and a **formula bar**. At this point, the latter will show two windows, of which the larger one will be empty.

Starting from the bottom of the screen and moving upwards, we usually first encounter the **task bar**, which has the Start button in its left corner. Next to the start button you will find the name of the Workbook you are using. When you have not yet given it a name, Excel will just call it Book1, Book2, etc. Above the task bar is the **status bar**, which may be largely empty for now.

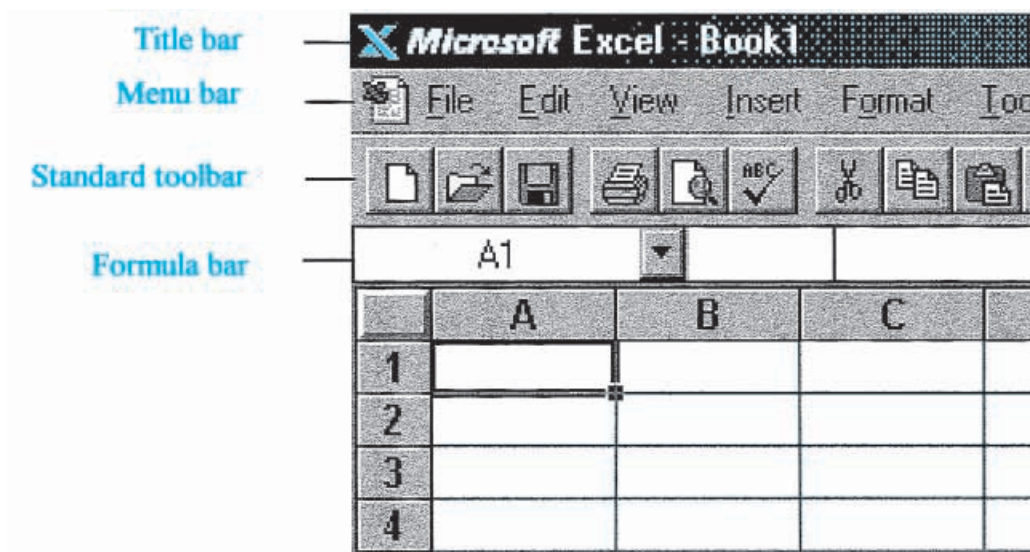


Fig. 1.2-1: The left top corner of the spreadsheet.

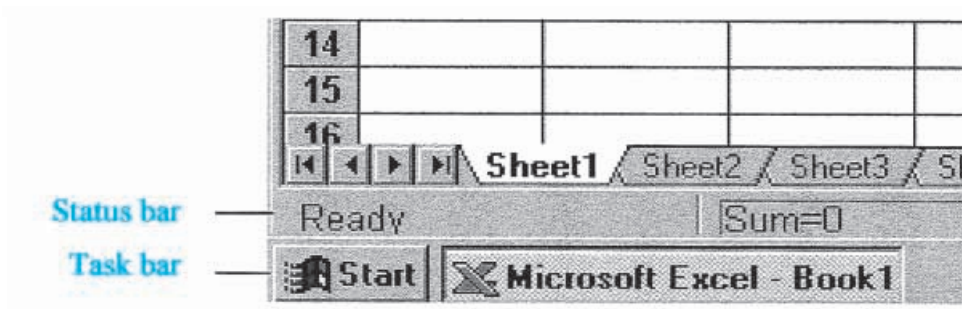


Fig. 1.2-2: The left bottom corner of the spreadsheet.

What we have described so far is the frame *around* the actual spreadsheet. Now we come to the spreadsheet itself, which is called a workbook, and is organized in different pages.

Above the status bar you will find a tab, in Fig. 1.2-2 labeled Sheet1, which identifies which *page* of the work book is open. Here, then, you see the general organization of individual spreadsheet pages into workbooks. You can have as many pages in your workbook as you wish (by adding or removing sheets), and again as many different workbooks as you desire. For the exercises in the present text, you may want to use a new sheet for each exercise, and a new workbook for each chapter, and label them accordingly.

In the region between the formula bar and the status bar you will find the actual working part of the spreadsheet page. It starts at the top with a sequence of rectangles, each containing one letter of the alphabet on a gray background. It ends, at the bottom, with a bar containing a series of tabs;

## 1.2 A first look at the spreadsheet

5

one such **tab**, such as the one labeled Sheet1 in Fig. 1.2-2, will have a white background, indicating the currently open (or 'active') sheet, while the others will be gray. In between these is a rectangular array of blank cells. Each such **cell** can be identified by its (vertical) **column** and its (horizontal) **row**. Columns are labeled by the letters shown just above row 1 of the spreadsheet, while rows are labeled by the numbers shown to the left of column A. The cell at the top left of the spreadsheet is labeled A1, the one below it A2, the one next to A2 is B2, etc. One cell will be singled out by a heavy black border; that is the highlighted, **active** cell in which the spreadsheet anticipates your next action. The **address** of the active cell is displayed in the left-most window of the formula bar; in Fig. 1.2-1 it is cell A1.

To **activate** another cell, move the mouse so that the pointer, which should now have the shape of a hollow **cross**, is within that cell, then click. The corresponding cell coordinates will show on the left-most window of the formula bar. When you move the mouse pointer to another cell and click again, that cell will now become the active one. Note that the left-most window in the formula bar will track the coordinates of the active cell. Play with moving the active cell around in order to get a feel for manipulating the mouse.

A cell can also be specified by typing its coordinates. The simplest way to do so is by using the **function key** labeled F5. (The function keys are usually located above the regular alphabet and number keys, and labeled F1 through F10 or F12. On some keyboards they are found to the left of the alphabet keys.) A **dialog box** will appear, and you just type the coordinates of the cell, say, D11, and deposit this by depressing the large 'enter' key (to the right of the regular alphabet keys). Another way, initially perhaps more convenient for those used to DOS-based spreadsheets, is to use the keystroke sequence Alt + e Alt + g. Here Alt + e denotes that you depress Alt and then, *while keeping Alt down*, also depress e; follow this by Alt + g. Alt specifies that you want to select an item from the menu bar, e selects the Edit command, and g the Go to command, where the underlining indicates the letter to be used: e in Edit, g in Go, o in Format, etc. As a gesture to prior users of Lotus 1-2-3 or QuattroPro, you can even use the slant instead of the Alternate key: / + e / + g. Any of the above methods will produce the dialog box in which to type the cell coordinates.

Below we will usually indicate how to accomplish something by using the mouse. For those more comfortable with using the keyboard rather than the mouse, keystrokes to accomplish the same goals are often available. There is no need to memorize these commands: just look for the underlined letters to find the corresponding letter code. Using keystrokes is often faster than pointing-and-shooting with a mouse, especially when you use a track ball.

Note that, inside the cell area of the spreadsheet, the mouse pointer usually shows as a cross. Select a cell, then move the pointer away from it and back again. You will see that, near the border of the active cell, the pointer changes its shape and becomes an **arrow**. When the pointer shows

as an arrow, you can depress the left mouse button and, while *keeping it down*, move the pointer in the cell area. You will see that this will **drag** the cell *by its border*. By releasing the mouse button you can deposit the cell in a new location; the formula bar will then show its new coordinates.

Practice activating a set of neighboring, contiguous cells; such cell **blocks** or **arrays** are often needed in calculations. Move your mouse pointer to a particular cell, say cell F8, and click to activate it. You can now move the pointer away, the cell remains active as shown by its heavy border; also, the formula bar shows it as the active cell regardless of where you move the mouse pointer, as long as you don't click. Return the pointer to cell F8, and depress the left mouse button *without* releasing it, then (while still keeping the cell button down) move the mouse pointer away from cell F8 and slowly move it in a small circle around cell F8. You are now outlining a cell **block**; its size is clear from the reverse color used to highlight it (it will show as black on a white background, except for the cell with which you started, in this example F8, which will remain white, and which we will call the **anchor cell**). The size of the block will show in the formula bar in terms of rows and columns, e.g., 3R × 2C will denote a block three rows high and two columns wide. By releasing the mouse button you activate the entire block, while the formula bar will return to showing the location of the anchor cell. You can then move away from it; the active block will remain. After you have selected the cell block, go back to it, grab its border (when the pointer is an arrow) and move the entire block around! To deposit the block in a new location, just release the mouse button. To abolish a block, release the mouse button to deposit it, then move the pointer to another cell and click on it.

To activate a block of cells from the keyboard, use F5 (or Alt + e Alt + g), then specify the block by the coordinates of its upper left cell and of its lower right cell, separated by a colon, as in D4:E9, and deposit it with the enter key.

There is yet another way to activate a block, starting from a single active cell. Again move the mouse pointer outside the active cell, but now approach the small square in the right bottom corner of the border around the active cell; this little square is the cell **handle**. The mouse pointer will change into a plus sign when it points to the cell handle; you can then drag the cell *by its handle* (rather than by its border) and make either columns or rows. Again, fix your choice by releasing the mouse button. You can drag it again to make a block out of a row or column. Practice these maneuvers to familiarize yourself with the mouse, and see how the pointer changes from a hollow cross (when you point at the middle of the cell) to an arrow (at its border) to a plus sign (at its handle). Below we will specifically indicate when to use the cell border or the cell handle; if nothing is specified, go to the center of the cell and use its standard pointer, the hollow cross.

## 1.3 A simple spreadsheet and graph

The spreadsheet is designed to facilitate making calculations, especially repeated calculations that would quickly become tedious and boring if you had to do them by hand, or even on a pocket calculator. (Unlike humans, computers do not tire of repetition.) The spreadsheet is also very useful when you have computations that would be too difficult for a pocket calculator. The tabular format, resembling an accountant's ledger, helps us to organize the calculations, while the so-called 'double precision' of the spreadsheet keeps round-off errors in check. When you change one number somewhere in a spreadsheet, the computer automatically recalculates all cells that depend on the one you have just changed. (There are a few exceptions to this statement: special functions and macros do not update automatically. We will alert you when we come to them.) The spreadsheet also makes it very easy to construct graphs. We will demonstrate this now that you know how to move around in the spreadsheet.

Among the things we can place inside a cell are a **number**, a **label** such as a column heading, or a **formula**. A cell can hold only one of these items at a time. Activate cell A1 by clicking on it, then type the letter x, followed by depressing the 'enter' key, or by moving the mouse pointer to a different cell and by then clicking on that other cell. Either method will deposit the typed letter.

Activate cell A3; to do this, either move the mouse pointer to cell A3 and click, or use the down arrow to get there. In cell A3 deposit the number 0. (As with the letter x, nothing will happen until you deposit it, using the Enter key. This lets the computer know that this is all you want to enter, rather than, say, 0.3 or 0.0670089.) Be careful to distinguish between the number 0 and the letter O; they are close neighbors on the keyboard but they are completely different symbols to the computer. Similarly, don't confuse the number 1, the lowercase letter L, and the capital I.

In cell A4 deposit the number 1. The letter x in A1 will usually show as left-justified (i.e., placed in the left corner of its cell), whereas the numbers 0 and 1 will usually be right-justified. (We hedge our bets with the 'will usually be' because all these features can easily be changed, as they may well have been on the computer you are using.) Return to cell A3, then activate both cells (by depressing the left mouse button while pointing to A3, keeping it down while moving to cell A4, then releasing the button). Both cells should now be active, as shown by their shared border.

Now comes a neat trick: grab both cells by their common handle (the little square at the right-hand bottom of their common border), drag the handle down to cell A11, and release the mouse button. With this simple procedure you have made a whole column of numbers, each one bigger by 1 than that in the cell above it!

	A	B	C
1	x	sine	
2			
3	0	0	
4	1	0.707107	
5	2	1	
6	3	0.707107	
7	4	1.23E-16	
8	5	-0.70711	
9	6	-1	
10	7	-0.70711	
11			

Fig. 1.3-1: Detail of the spreadsheet with its two columns and column headings.

Had you started with, say, the number 7 in cell A3, and 4.6 in cell A4, column A would have shown 7, 4.6, 2.2,  $-0.2$ ,  $-2.6$ , and so on, each successive cell differing from its predecessor by  $4.6 - 7 = -2.4$ . In other words, this method of making a column generates constant increments or decrements, in arithmetic progression. Try this, with different values in A3 and A4. Then go back to deposit the series ranging from 0 to 7 with an increment of 1 or, in mathematical notation, the series 0 (1) 7. Incidentally, there are many other ways to fill a column, some of which we will encounter later.

In column B we will now calculate a sine wave. Activate cell B1 and deposit the heading 'sine'. Move to cell B3 and deposit the formula  $=\sin(a3*\pi()/4)$ . The equal sign identifies this as a formula rather than as text; the asterisk indicates a multiplication. The spreadsheet uses the notation  $\pi()$  to denote the value of  $\pi$ ; the brackets alert the computer that this is a function. Excel instructions do not distinguish between lower case and capitals, but the formula bar always displays them as capitals, which are more clearly legible. By now your spreadsheet should look like that depicted in Fig. 1.3-1.

If you were to extend the columns to row 11, the value shown in cell B11 might baffle you, since it may not quite be 0 but a small number close to it, reflecting computer round-off error. But don't worry: the error will usually be below 1 part in  $10^{15}$ .

There is a more convenient way to generate the second column. After you have entered the instruction  $=\sin(a3*\pi()/4)$  in cell B3, grab its handle (at which point the mouse arrow will show as a plus sign) and double-click. This will copy the instruction down as far as the column to its immediate right contains data! This is a very useful method, especially for long columns.



### 1.3 A simple spreadsheet and graph

9

When there are no data to its immediate right, the column to its immediate left will do. When both are absent, the trick will not work.

Finally we will make a graph of this sine wave. Doing so is slightly different in Excel 97, Excel 98 for the Mac, or Excel 2000 on the one hand, and Excel 95 or Excel 5 on the other. We will here describe the procedure for each of these two versions.

Bring the mouse pointer to cell A3, click on it, drag the pointer (while keeping the mouse button depressed) to cell B11, then let go of the mouse button. This will activate (and highlight) the rectangular area from cell A3 through B11 (in spreadsheet parlance: A3:B11) containing the data to be graphed. Alternatively, you can highlight cell A3, then depress the Shift key, and while keeping this key down depress End, ↓, End, and finally →. (The sequence Shift + End, Shift + ↓ will highlight the column A3:A11, while Shift + End, Shift + → will include column B. As with double-clicking on the cell handle to copy an instruction, Shift + End looks for contiguous data.)

#### 1.3a Making a graph in Excel 97 or a more recent version

*If this is your first reading, and you use Excel 95 or Excel 5, skip the following, and continue with section 1.3b.*

In Excel 97 or a more recent version, go with the mouse pointer to the menu bar, click on **I**nsert, and in the resulting drop-down submenu click on **C**hart. Or achieve the same result with the keystrokes Alt + i, Alt + h. Either method will produce a dialog box labeled Chart Wizard – Step 1 of 4 – Chart Type.

In the list of **C**hart types, click on XY (Scatter); do *not* select the Line plot, which in Excel means something quite different from what a scientist might expect. The line plot can give you very misleading graphs because it presumes that the x-values are always equidistant.

As soon as you have selected the XY plot, the right-hand side of the dialog box will show five Chart sub-types: loose points, points connected by smooth or straight lines, or just smooth or straight lines. For now, pick the points connected by smooth lines – you can always change it later. (This is a general property of working with Windows Excel: you need not agonize over a choice, because there are almost always opportunities to change it later. So the best strategy is: when in doubt, pick *something*, move on, and worry about the details later.) Click on the Next > button.

Step 2 of the Chart Wizard shows the **D**ata range selected. Also, under the Series tab, it shows which column will be used for **X**-values, and which for **Y**-values. The default (i.e., the assumption the spreadsheet makes in case you do not overrule it) is to use the left-most column of the selected block for X-values, so you need not take any action here, just press on with Next >. But it is handy to know that you can here, in step 2 of the Chart Wizard, change the assignments for X and Y.

Step 3 lets you enter a Chart title and axes labels. Click on the Chart title window, and enter Sine wave. Then click in the Value (X) Axis window, and enter angle. Finally, click in the Value (Y) Axis window, and enter sine. A picture will show you what your graph is going to look like.

There are other things you can specify at this point, such as the axes, grid-lines, legends, and data labels, but we will forgo them here in order to keep things simple for now, and to illustrate later how to modify the end product. So, on to the Next >.

Step 4 defines the chart location, either As a new sheet, or As object in a spreadsheet page. Select the latter, and Finish. This will place the graph on the spreadsheet.

Now click on the graph, preferably inside its outer frame near its left edge, where the computer cannot misinterpret your command. This will adorn the graph with eight black handles, which allow you to change its size and location. First, locate the mouse pointer on the graph, depress the mouse button, and while keeping it down move the graph to any place you like on the spreadsheet, preferably somewhere where it does not block data from view. To release, simply release the mouse button. Note that the graph as it were floats on the page, and does not obliterate the underlying information. To fit the graph in the cell grid, depress the Alt key, then (while keeping Alt depressed) bring the mouse pointer to a handle in the middle of the side of the graph, where the pointer should change into a two-sided arrow, and pull that pointer toward a cell boundary. Repeat with the other sides. For greater efficiency you can combine this for two adjacent sides by pulling or pushing on two opposing corners.

In the final result, click on the little rectangular box to the right of the graph, then press Delete.

If you want to remove the gray background (which seldom prints well) just click somewhere in the plot area (where the label shows Plot Area), right-click, highlight Format Plot Area, and under Area either select None or, in the choice of colors, click on white. Exit with OK.

If you want to get rid of the horizontal grid lines, point to them (the label will identify Value (Y) Axis Major Gridlines), right-click, and select Clear.

To change the range of the x-scale, point to the axis (the label will show Value (X) Axis), right-click, select Format Axis, and under the Scale tab pick the scale properties you want. And, while you're at it, please note that you can also change the font, size, color, position, and alignment of the numbers of the x-axis. Ditto for the numbers on the vertical axis.

To change the type of graph itself, point at the curve, right-click, and select Format Data Series. Then for the Line pick the Style, Color, and Weight you like, and for Marker the Style, Foreground and Background color, and Size.

And so it goes: you can point at virtually every detail of the graph, and modify it to your taste. Figure 1.3-2 shows you what you might have wrought.