Disclaimer

The version of Excel in the Glade Manual is from Microsoft Office 2010. Microsoft regularly produces new versions for multiple platforms. There are now so many different versions that it is impossible to capture all the variations in one manual. The elements are essentially all still there in whatever version you are using, but accessing them has often changed. You will have to translate what is presented in this manual to your version of Excel. That will be generally not that difficult.

Chapter 0:
Introduction to the Microsoft Excel Spreadsheet

Objectives

This chapter introduces you to the Microsoft Excel spreadsheet. You should gain an understanding of the following topics:

- The range of uses of a spreadsheet
- The basic elements of an Excel spreadsheet
- Manipulation of an existing spreadsheet
- Use of formulas to calculate values

Introduction

This chapter has two parts. In the first part, you will be introduced to the basics of launching, using and saving a spreadsheet. In the second part, you will see some prepared spreadsheet models which illustrate the fundamental ideas of a spreadsheet and the range of uses to which it can be put. You will modify these in some way to explore what you can do with them. In each case, you should try to understand certain principles demonstrated by the models. These demonstration models show you what you will be able to do with Excel by the end of the course, so don't worry about how to construct them just yet. The remaining chapters of the Glade Manual will cover all of that.

Starting Microsoft Excel

A window resembling Figure 0.1 should appear after Excel has launched. You won't be using the new spreadsheet that appears but it is worthwhile examining some of the features of the Excel application before closing this new worksheet.

The most obvious feature, in Figure 0.1, is the array of rows and columns that constitute the worksheet. Notice that the columns are labeled with letters across the top of the worksheet (A through to K are visible) and that the rows are labeled with numbers down the left-hand side (1 through to 23 are visible). These column and row labels form a primitive method of naming each cell on the worksheet – B6 for example is the name of a cell in the B column six rows down. Such names can be used in formulas, although we'll see better ways to name cells later on in this chapter.

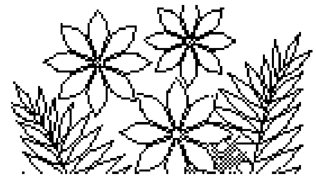

Figure 0.1 the Excel 2010 window at launch

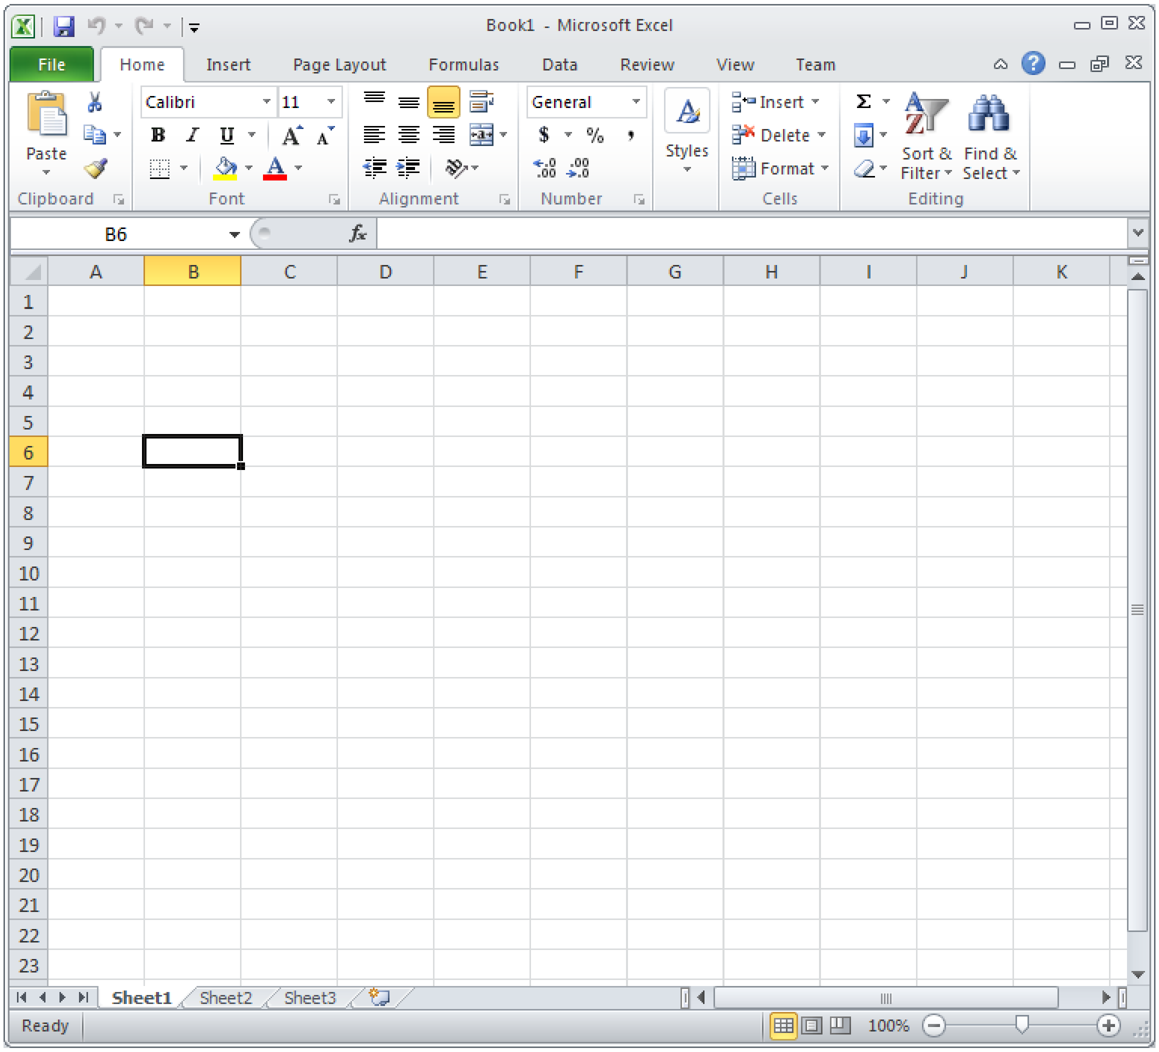

Entering Data into Cells

The second feature you should note is that one of the cells is highlighted with a black border. You can't see it in Fig. 0.1 but on your computer screen it should be cell B6. Point to any other cell with your mouse and click once. This changes the highlighted cell. Notice that the label in the area just above the worksheet on the left of the window contains the name of the selected cell:

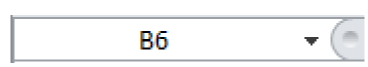

Type anything you like (your name perhaps) into the cell that you selected – i.e. just start typing and you will see whatever you type appear in the cell. At the same time whatever you type will appear across the top of the worksheet – in an area that is called the formula bar:

When you have finished typing press the Enter key and observe what happens to the cells on the worksheet. Do the characters you typed exceed the cell width? Experiment with entering large and small amounts of text in different cells.

Next select a different cell and this time type a number and press Enter. What do you observe about the placement of numbers as opposed to the small amount of text in a cell?

Other Worksheets in the Workbook

Next observe at the bottom of the screen that there are tabs labeled **Sheet1**, **Sheet2** etc. Currently **Sheet1** is highlighted but if you now click on one of the other sheets you will go to another worksheet of cells. This other worksheet is currently empty. Click on **Sheet1** to go back to the one on which you typed some data. There can be many worksheets in a spreadsheet model, with formulas linking them together.

Notice too the arrows next to the sheet tabs, which allow you to move forwards or backwards through the various sheets.

The Excel 2010 Ribbon Interface

With Excel 2007 Microsoft replaced the Command Menu and Toolbar with a "Ribbon" interface. Click on a tab on the top line (Home Insert Page Layout *etc.*), and a set of commands appears in the form of a ribbon below:

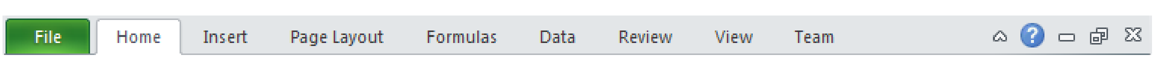

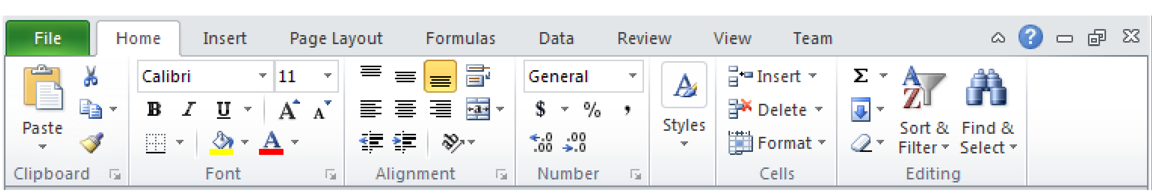

Closing the Workbook

You don't want to use the workbook you have been playing around with. This has just been a brief tour of the Excel application interface. You should now close this workbook so that you can begin to examine the demonstration models. There are two ways to close a workbook – you can either choose the Close command in the File tab or you can click on the **x** button in the top right hand corner of the workbook window. If you chose the latter method be careful to click on the **x** in the workbook window, which is the lower set of window controls. Do one of these now.

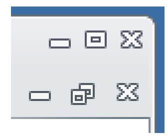

You'll find that a dialogue window appears asking if you want to save the workbook - you should click on Don't Save:

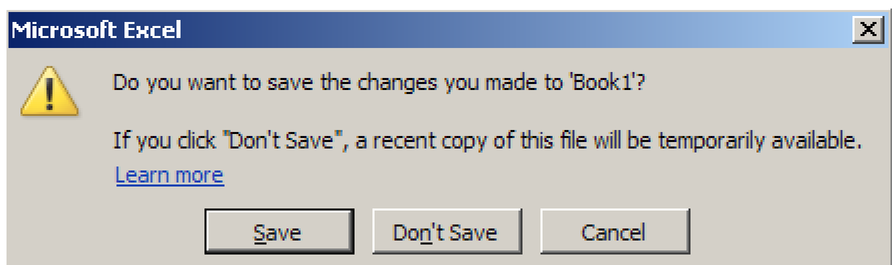

Exploring Some Spreadsheet Models

Download files from the course website

Locate the following Demonstration Models in the “Support Files” for Chapter 0. Save the files to your computer for future use.

Opening an Existing Spreadsheet

There are potentially two ways of opening an existing spreadsheet. One is to use the Open command from the File tab. Another is to locate the file and double-click on it. This time you will have to use the first way.

Demonstration Model 1

The model you have just opened contains two worksheets. The names have been changed from Sheet1 and Sheet2 to Comments and Summing; and the Comments worksheet is the one displayed.

The Comments Worksheet

This worksheet contains a brief description of the purpose and features of the spreadsheet model. It illustrates a documentation principle that you should follow in all of the spreadsheet work you do. Documentation of your work is crucial – not only for other people who may need to understand or modify it, but also for yourself when you return to it after a few weeks or months and find that you have forgotten what you did.

The Summing Worksheet

This worksheet is shown in Fig. 0.2. It shows three methods of adding a row of numbers. The first thing you should do is examine the contents of the cells by selecting certain ones. For example, select the cell B1. It looks like it might have the characters “ **`a l`**” – i.e. part of “**`a list`**” in the cell. However, you can see from the formula bar that cell `B1` is actually empty. Select cell A1 and you will see the entire first row displayed in the formula bar.

Figure 0.2 the Summing worksheet in demonstration model 1

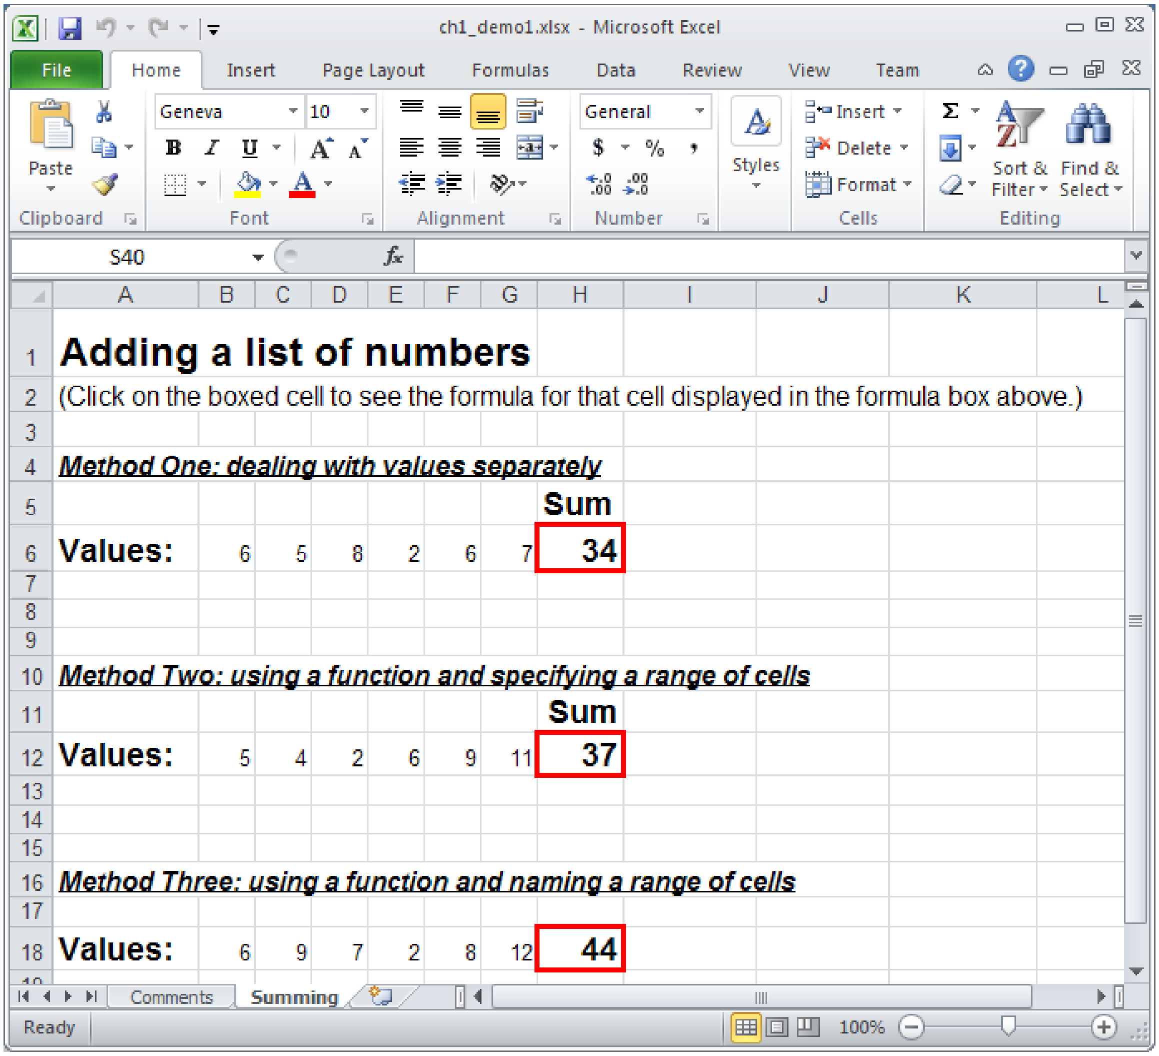

Select cell B6 and you will see that it contains the number 6. Select cell H6, however, and you will see not the number 34 displayed in the formula bar but rather a formula, which has the value 34 as its result from adding the values in the row. Examine the formulas that calculate the other sum values and observe their differences.

Changing Values on the Spreadsheet

Select one of the numbers in a row and type a new value, pressing Enter when you've done. Notice that the sum value is instantly changed to reflect the new value. This interactive feature of a spreadsheet is one of the reasons spreadsheet programs are such valuable tools in so many problem solving applications.

Inserting a Column (or Row)

Add a new column to the spreadsheet so that a new number can be included in the row that is being summed. To do this, select one of the columns C, D, E, F, or G by clicking on the letter in the top row of the worksheet. Then from the Home tab, click on Insert within the Cells Group. A new empty column should appear to the left of whichever column you had selected. A row may be inserted in the same manner – that is, by selecting a whole row by clicking on its number on the far left (you'll select the row below where you want the new one to appear) and from Home, click on Insert within Cells.

The Generality of Formulas

Now select the new cell in one of the rows of numbers and type a new value that you would like to be included in the sum. Do this for each of the three rows of numbers observing in which cases the value for the sum also changes. One of the formulas gives the wrong answer when this new value is inserted, and hence you should draw the lesson that this type of formula should be avoided. Delete the new column you inserted once you have done this.

Naming a Range of Cells

Click on the little downward pointing arrow to the right of the name bar (on the left of the window just above the column headings A, B, C etc:

The name bar: 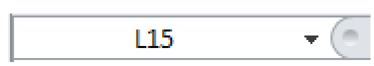

You'll see a single name, Values, appear. Click on that name and observe that the third row of numbers is highlighted. This row of numbers has been defined to have the name Values in this model. It is possible to name single cells or groups of cells like this. The aim is to define easily remembered names for data in your worksheets. These names can then be used in formulas so that they are easier to read, rather than having only cell references such as B13 or whatever.

Click on the Sum value for this row and observe that the formula that calculates this value uses the name Values.

Functions

The formulas for the sum in methods two and three also use a function – the SUM function. Functions take the general form of

SomeFunctionName(someArgumentList)

You'll use various functions extensively in later chapters. For now notice that the SUM function has as its argument list a single value expressed as either a range of cells (B12:G12 for example) or the single name, list, representing a range of cells. This is the most common, although not the only way, of using the SUM function. Other functions might have more than one argument with the arguments separate by commas.

Demonstration Model 2

Close the previous spreadsheet model (you don't need to save it) and open the one called Demonstration Model 2. Read the Comments worksheet carefully before you switch to the second worksheet called DataSheet. The DataSheet worksheet is shown in Fig. 0.3.

Figure 0.3 – the DataSheet worksheet from demonstration model 2

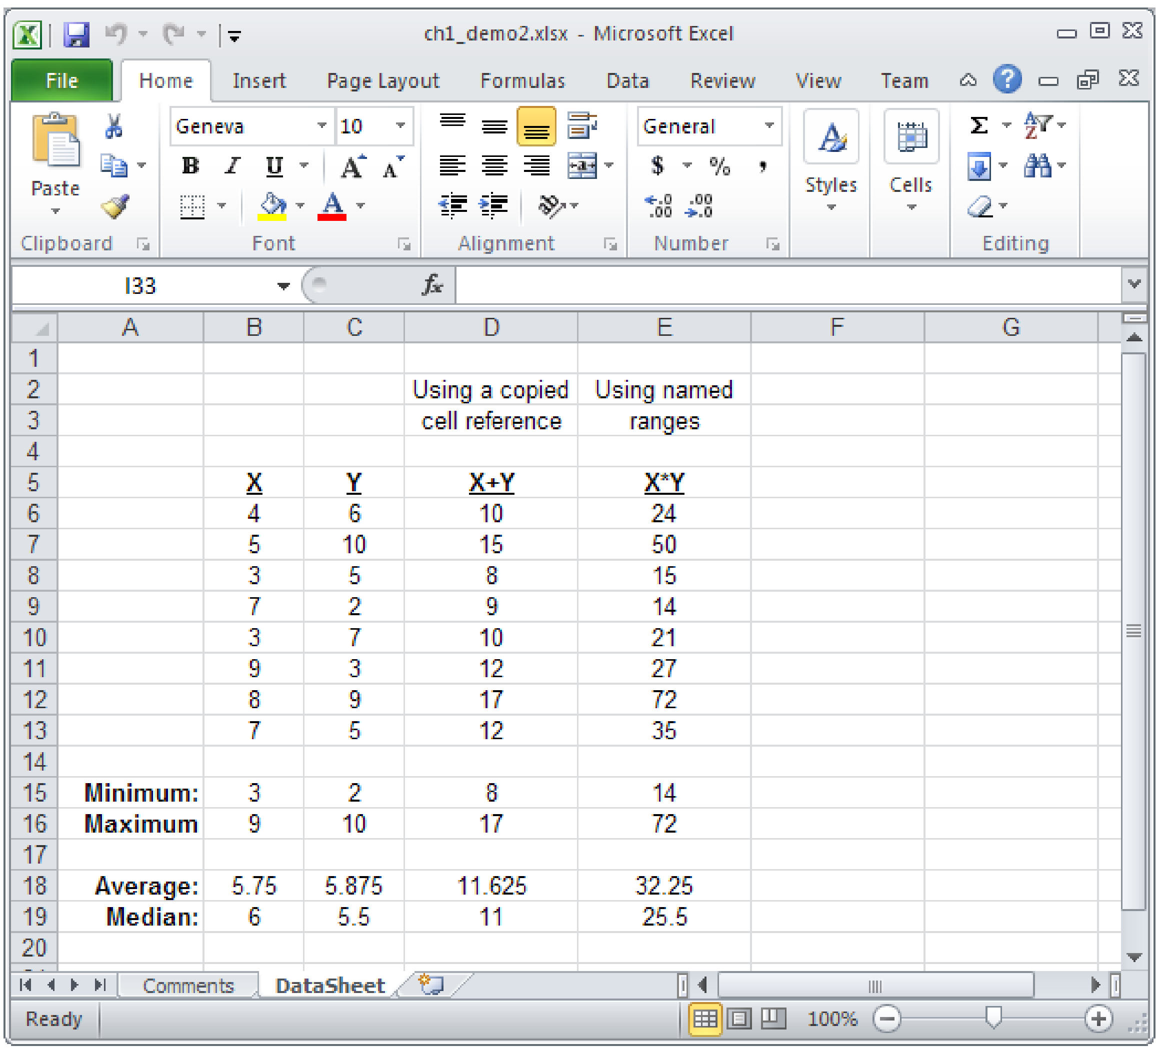

	A	B	C	D	E	F	G
1							
2				Using a copied	Using named		
3				cell reference	ranges		
4							
5		X	Y	X+Y	X*Y		
6		4	6	10	24		
7		5	10	15	50		
8		3	5	8	15		
9		7	2	9	14		
10		3	7	10	21		
11		9	3	12	27		
12		8	9	17	72		
13		7	5	12	35		
14							
15	Minimum:	3	2	8	14		
16	Maximum	9	10	17	72		
17							
18	Average:	5.75	5.875	11.625	32.25		
19	Median:	6	5.5	11	25.5		
20							

This worksheet demonstrates a number of new methods of making calculations, as well as a few new functions. The raw data on this worksheet – that is the data that is entered by hand and therefore not calculated by formulas – is contained in the columns labeled X and Y. Names have been created for these two columns. Click on the arrow in the Name Box and select first X and then Y to see what they refer to.

Copying Formulas – the Fill Command

The data in the other columns labeled X+Y and X*Y is created very easily, even though apparently a lot of formulas are involved. Select the cell D6 containing the value 10 and you'll see that it is calculated by the formula B6+C6. The other cells in this column are essentially calculated using the same formula except that the row numbers change. Examine a few of them. These formulas are created very easily by first typing the top formula, then selecting that cell and all of the others in the column (which are empty at this point) to which you want the formula to apply, and then choosing from Home tab the Fill/Down command within Editing group. This command copies the top formula into all of the selected cells, changing the row numbers as appropriate. The other columns of calculated data are created in essentially the same way – one formula was typed for the top cell, and then filled down the column.

Using Named Ranges in a Formula

Examine the formulas for the first X*Y column. The names X and Y refer to the whole columns (as you observed in a previous section) and here the same formula is filled down the X*Y column to calculate all of these values. Although the names refer to the whole of the X and Y columns the formula "knows" to use just the two X and Y values on the same row in order to calculate any particular X*Y value.

Some New Functions

This worksheet also introduces some additional functions in order to familiarize you with their use. You've seen the SUM function before – and now you can see the MIN, MAX, AVERAGE and MEDIAN functions, all using a column of data as their argument. The column is either expressed as a range of cell references or as a named range. Close the worksheet before you move on to the next example (there's no need to save it).

Demonstration Model 3

Open Demonstration Model 3. This model implements a slightly more complex formula than the previous examples, and also demonstrates the use of the graphing capabilities of a spreadsheet.

The model calculates the distance an arrow will travel if fired with an initial velocity and an initial angle, and assuming the ground is perfectly flat. The initial velocity is treated as a parameter – that is, it is held constant while the initial angle is varied and the distance is calculated for each angle.

The model has three worksheets – the Comments worksheet (read it carefully), the Parameters worksheet (Figure 0.4) and the Trajectory worksheet (Figure 0.5). Besides the initial velocity the Parameters worksheet also contains the other constant quantity – the gravitational constant. (Notice that the units are included in a cell adjacent to the value.)

Figure 0.4 – the Parameters worksheet in the Flight of an Arrow model

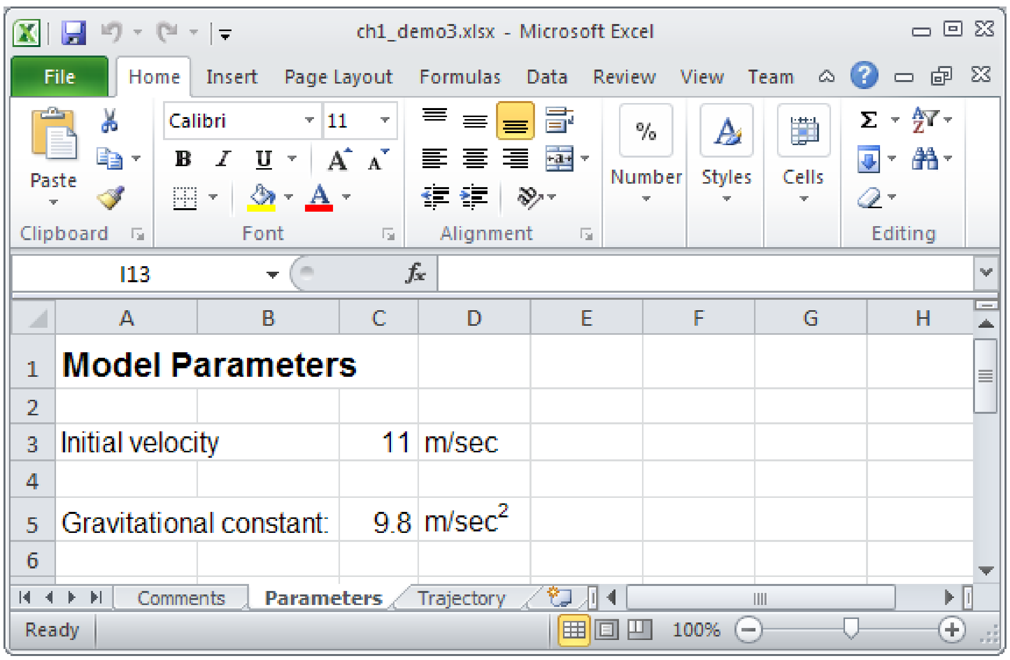

The Trajectory worksheet (Figure 0.5) contains a column of values for the initial angle and a column of values for the calculated distance the arrow travels, as well as a graph of the relationship.

Nested Functions

The formula to calculate the distance the arrow will travel is:

$$\text{distance} = \text{velocity}^2 * \text{sine}(\ 2*\text{angle}\ ) / g$$

Click on one of the values in the Distance column in the Trajectory worksheet and observe the formula as it is expressed in Excel. In this formula two functions are used – but one of them is nested within the parentheses of the other. The formula is:

```
=V*V*SIN(2*angle*PI()/180)/g
```

Launch Angle is in degrees (you can see this in the worksheet) but the SIN function expects the angle to be in radians. Therefore, the angle is converted by multiplying by π/180. And π is represented by the second function, PI(). Notice that the function PI() does not have any arguments, as indicated by the empty parentheses – it simply produces a very good approximation of value of π as its result.

Inter-Worksheet Formulas

The formula that calculates the Distance values in the Trajectory worksheet uses the named values V, angle, and g. However, V and g refer to values that are in a different worksheet, the Parameters worksheet, and hence the formula is actually what we will call an inter-worksheet formula.

Use Formulas/Name Manager to open the Name Manager window to see the list of names. You'll see that the cell reference is prefixed by the worksheet name – Parameters! in this case. This method of specifying a cell or range as coming from a particular worksheet is part of the syntax of Excel formulas. Take careful note of this syntax because you'll meet it frequently and sometimes you might want to be able to type it out yourself. Close the Name Manager window again.

The Name Manager has been removed in more resent versions of Excel. This is very unfortunate as it was a very handy tool for viewing the list of named ranges.

A Graph

Graphs of many different types can be created by a spreadsheet program – and either placed next to data on a worksheet (as in this example) or placed on a separate worksheet of their own. One interesting feature is that the graph is interactive in the sense that it is redrawn immediately if you change any of the data from which it is derived.

Observe the maximum distance that the arrow can travel (it is about 7.5 metres for an initial angle of 45 degrees) and then switch to the Parameters worksheet where you can change the value of the Initial Velocity. Type a new value and return to the Trajectory worksheet. You should see that the graph has changed to reflect the new data.

Figure 0.5 – the Trajectory worksheet in the Flight of an Arrow model.

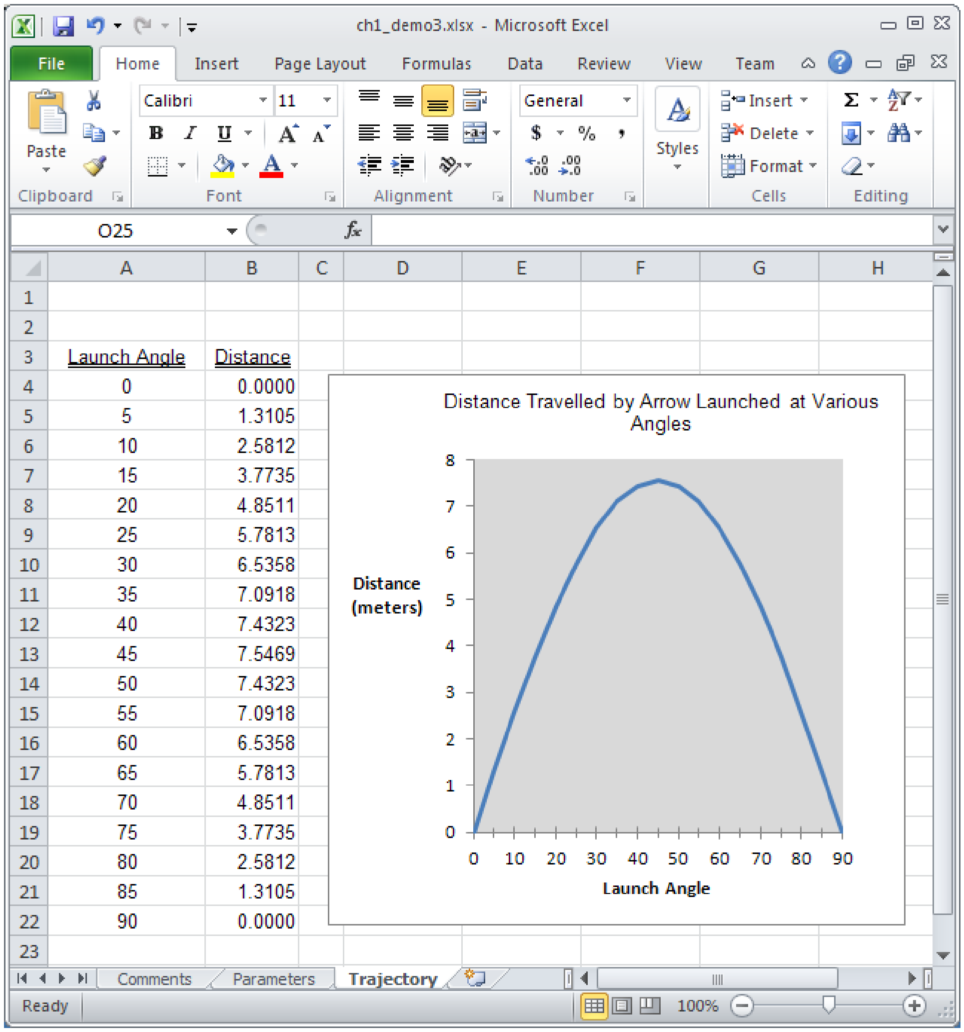

Demonstration Model 4

Open Demonstration Model 4. This model contains 4 worksheets – the Comments worksheet; a worksheet containing a graph and values for the initial velocity and angle of a canon ball fired over a hill; a Parameters worksheet containing values which define the shape of the hill (and the gravitational constant); and a Hill_and_Projectile worksheet containing the data that is to be graphed. The graph is shown in Figure 0.6.

<u>Figure 0.6 – the Canon Ball Trajectory</u>

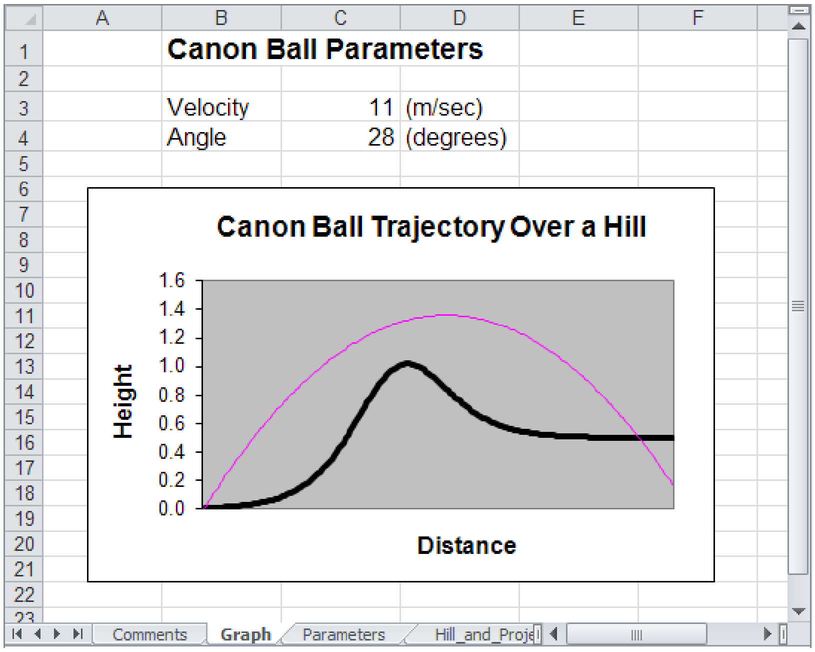

The main new feature of this model is the ease with which the model parameters can be adjusted in order to achieve some objective. The objective in this case might simply be to have the canon ball hit the far right corner of the land (i.e. where the black line reaches the edge of the graph). One approach is simply to change the velocity and/or angle parameters hoping by trial and error to hit the corner. Try this now.

New Functions and Operators

Besides this goal seeking demonstration the model also provides examples of the use of yet more functions. Examine the formula for the Z column and you'll see the use of the EXP function (which calculates the value of e^x, where x is the argument of the function). The formula that calculates the Hill Height column demonstrates the use of the exponentiation operator, ^, and the Projectile Height formula shows the use of further trigonometric functions, TAN and COS, in a fairly complex expression.

Demonstration Model 5

Open Demonstration Model 5. This model demonstrates the capability of a spreadsheet to analyze (in this case statistically) a large set of data, and to present the results visually. The model simulates the rolling of two dice and presents simple statistics (frequency counts) of the sum of the numbers on the two dice.

There are three worksheets in the model – a Comments worksheet, the dieValues worksheet which contains the values on the two dice and their sum, and a Countif worksheet which demonstrates one possible way of calculating the frequency counts.

The RAND Function

Select the dieValues worksheet and examine the formula that calculates a die value. It uses nested functions again. The RAND() function produces a random number between 0 and (less than) 1, which must then be multiplied by 6 and converted to an integer in order to obtain a value between 0 and 5 inclusive. The INT function does the conversion to an integer. Add one to produce the values on the dice.

The COUNTIF Function

To count the frequency of each possible roll of the dice, we list the possible values first (i.e. the values 2 to 12) and use them as a criterion by which to decide whether to count a particular value. The COUNTIF function uses the possible value of the sum of the dice in order to decide whether to include a value from the long list of dice "throws" in the frequency count for that possible sum value.

Select the first value in the Frequency column and examine the formula. It should be

```
=COUNTIF(dieValues!$C$2:$C$701,Countif!A3)
```

Looking at the last argument first, namely Countif!A3, observe that this refers to the first value in the Sum column, i.e. 2. The first argument refers to the large column called Sum in the dieValues worksheet. Thus this function arrives at this first value in the Frequency column by counting those values from the dieValues!Sum column which equal 2.