INTEGRAL UNIVERSITY, LUCKNOW

DIRECTORATE OF DISTANCE EDUCATION

BBA-305 Paper Code: BM/B

BUSINESS MATHEMATICS

CONTENTS

UNIT-1 BASIC MATHEMATICS

CONTENTS

INTRODUCTION

In this first chapter, a number of concepts are introduced that are fundamental to many areas of business mathematics. We begin by defining some key words and phrases.

Mathematical operations and brackets

The basic mathematical operations are addition, subtraction, multiplication and division; and there is a very important convention about how we write down exactly what operations are to be carried out and in what order. Brackets are used to clarify the order of operations and are essential when the normal priority of operations is to be broken. The order is:

- work out the values inside brackets first;
- Powers and roots;
- Multiplication and division are next in priority;
- Finally, addition and subtraction.

Basic Mathematics

For example, suppose you want to add 2 and 3 and then multiply by 4. You might write this as 2+3X4

However, the above rule means that the multiplication will take priority over the addition. What you have written will be interpreted as an instruction to multiply 3 by 4 and then add 2. It would be useful at this point to briefly digress and check your calculator. Type 2+3X4=' into it. If the answer is 14, you have a scientific calculator that obeys mathematical priorities. If the answer is 20, your calculator is non-scientific and perhaps will not be suitable for your CIMA studies.

Returning to the main problem, if you want to add 2 to 3 and then multiply by 4, you must use brackets to give priority to the addition – to ensure that it takes place first. You should write (2 + 3)X 4. The contents of the bracket total 5, and this is then multiplied by 4 to give 20.

Different types of numbers

A whole number such as - 5, 0 or 5 is called an integer, whereas numbers that contain parts of a whole number are either fractions – such as 3/ 4 – or decimals – such as 0.75.

Any type of number can be positive or negative. If you add a positive number to something, the effect is to increase it whereas, adding a negative number has the effect of reducing the value. If you add =B to any number A, the effect is to subtract B from A. The rules for arithmetic with negative numbers are as follows:

- adding a negative is the same as subtracting, that is $A + (-B) = A - B$;
- subtracting a negative is the same as adding, that is $A - (-B) = A + B$;
- if you multiply or divide a positive and a negative, the result is negative, that is $(+) \times (-)$ and $(-) \times (+)$ and $(+) \div (-)$ and $(-) \div (+)$ are all negative;
- if you multiply or divide two negatives, the result is positive, that is $(-) \times (-)$ and $(-) \div (-)$ are both positive.

Rounding

Quite often, numbers have so many digits that they become impractical to work with and hard to grasp. This problem can be dealt with by converting some of the digits to zero in a variety of ways.

Rounding to the nearest whole number

For example, 78.187=78 to the nearest whole number. The only other nearby whole number is 79 and 78.187 is nearer to 78 than to 79. Any number from 78.0 to 78 . 49 will round down to 78 and any number from 78.5 to 78 . 99 will round up to 79. The basic rules of rounding are that:

1. digits are discarded (i.e. turned into zero) from right to left;
2. reading from left to right, if the first digit to be discarded is in the range 0–4, then the previous retained digit is unchanged; if the first digit is in the range 5–9 then the previous digit goes up by one.

Depending on their size, numbers can be rounded to the nearest whole number, or 10 or 100 or 1,000,000, and so on. For example, 5,738 = 5,740 to the nearest 10; 5,700 to the nearest 100; and 6,000 to the nearest 1,000.

Significant figures

For example, 86,531 has five digits but we might want a number with only three. The '31' will be discarded. Reading from the left the first of these is 3, which is in the 0–4 range, so the previous retained digit (i.e. the '5') is unchanged. So 86,531 = 86,500 to three significant figures (s.f.).

Suppose we want 86,531 to have only two significant digits. The '531' will be discarded and the first of these, '5' , is in the 5–9 range, so the previous digit ('6') is increased by 1. So 86,531= 87,000 to two s.f. Zeros sometimes count as significant figures; sometimes they do not. Reading a number from the right, any zeros encountered before you meet a non-zero number do not count as significant figures. However, zeros sandwiched between non-zeros are significant. Hence, 87,000 has two s.f., while 80,700 has three.

Decimal places

The other widely used rounding technique is to discard digits so that the remaining number only has a specified number of decimal places (d.p.). For example, round 25.7842 to two d.p. The digits to be discarded are '42' , the first ('4') is in the 0–4 range and so the next digit ('8') remains unchanged. So 25.7842 =25.78 to two d.p. Strings of '9' can be confusing. For example, if we want to round 10.99 to one d.p., the first digit to be discarded is '9' and so the next digit, also '9' , goes up to '10' . In consequence, the rounded number is written as 11.0 to one d.p.

Rounding up or rounding down

A number to be rounded up will be changed into the next higher whole number so, for example, 16.12 rounds up to 17. A number to be rounded down will simply have its decimal element discarded (or truncated). Numbers can also be rounded up or down to, say, the next 100. Rounding up, 7,645 becomes 7,700 since 645 is increased to the next hundred which is 700. Rounding down, 7,645 becomes 7,600.

Powers and roots

Definitions

1. The n th power of a number, a, is the number multiplied by itself n times in total, and is denoted by an or a^n. For example, 2^5 or $2^\wedge 5 = 2 \times 2 \times 2 \times 2 \times 2 = 32$

2. Any number to the power of zero is defined to be 1. For example, $7^0 = 1$ a =n is the reciprocal of an , that is a = n = 1 = an = 1/ an

$$3^{-2} = \frac{1}{3^2} = \frac{1}{9}$$

4. The n th root of a number, a , is denoted by a1/n and it is the number that, when multiplied by itself n times in total, results in a. For example,

$$8^{1/3} = \sqrt[3]{8} = 2$$

The square root, $a^{1/2}$, is generally written as $\sqrt{a}$ without the number 2.

5. an/m can be interpreted either as the m th root of an or as the m th root of a multiplied by itself n times. For example, $9^{5/2} = (\sqrt{9})^5 = 3^5 = 243$

6. The rules for arithmetic with powers are as follows:

(i) Multiplication: am = an = am=n . For example,

$2^3 \times 2^4 = (2 \times 2 \times 2) \times (2 \times 2 \times 2 \times 2) = 2^3 \times 2^4 = 2^{3+4} = 2^7 = 128$

Division: $a^m \div a^n = a^{m-n}$. For example:

(ii) $3^5 \div 3^2 = (3 \times 3 \times 3 \times 3 \times 3)/(3 \times 3) = 3^{5-2} = 3^3 = 27$

Powers of powers: $(a^m)^n = a^{m \times n}$. For example,

(iii) $(4^2)^3 = 4^{2 \times 3} = 4^6 = 4,096$

Mathematical operations in Excel

When performing calculations in Excel the same mathematical rules that have been discussed in this chapter apply. The following examples use the data from Examples 1.2.1 and 1.2.2 and show how formulae in Excel would be created to arrive at the same results.

Notice from Figure 1.1 that most of the values have been addressed by the cell reference – but it is also possible to incorporate numbers into the formulae.

Rounding numbers in Excel

To accurately round numbers in Excel a built-in function called = round () is used. This can be used to set any degree of accuracy required and once the function is incorporated into a formula any future references to the cell containing the round function will use a value rounded to the specified number of decimal places. The following example (Figure 1.2) illustrates this and uses the data from Example 1.4.1. Notice that the third example requested the result be rounded to three significant figures, the formula is a little more complex and has been done here in two steps.

In the first step in cell f6 the arithmetic has been performed and the result rounded to three decimal places. Then in g6 the len and the int functions have been applied to further round the result to three significant figures.

It is sometimes preferable to take the integer value of a number as opposed to rounding it to the nearest whole number. The difference is that the integer value is a number without any decimal places. Therefore the integer value of 9.99 is 9 and not 10 as it would be if the number had been rounded to the nearest whole number. Figure 1.3 shows the table used in the rounding exercise but with the Excel int function in place of the round function

Figure 1.4 shows the results of the rounding and the integer formulae used in figures 1.2 and 1.3.

Looking at Figure 1.4, the different result produced through the use of the int function as opposed to the round function can be seen. In each case the result has been rounded down to the integer value.

Variables and functions

A variable is something which can take different values. Variables are often denoted by letters. Thus, the set of positive whole numbers can be considered as a variable. If we denote it by x, then this variable can have many values. x

$x = 1$
or $x = 2$
or $x = 3$, and so on.

Another example is the set of the major points of a compass. If this variable is denoted by c, then it can have more than one value, but only a limited number.

c = north
c = south
c = north-west, and so on.

These examples show that variables can take on non-numerical 'values' as well as numerical ones. In this text we shall concentrate on numerical variables, that is, those whose values are numbers, like the first case above.

A mathematical function is a rule or method of determining the value of one numerical variable from the values of other numerical variables. We shall concentrate on the case where one variable is determined by or depends on just one other variable. The first variable is called the dependent variable, and is usually denoted by y, while the second is called the independent variable, denoted by x. The relationship between them is a function of one variable, often referred to as a function, for brevity. Note that whilst functions are similar to formulae (see Section 1.8) there are specific conditions relating to the definition of a function, but these are outside the scope of this book.

A very useful way of stating a function is in terms of an equation, which is an expression containing an 'equals' sign. The equation of a function will thus take the typical form: y= a mathematical expression containing x

If we know the value of the independent variable x, then the expression will completely determine the corresponding value of the dependent variable, y.

Formulae

A formula is a statement that is given in terms of mathematical symbols: it is a mathematical expression that enables you to calculate the value of one variable from the value(s) of one or more others. Many formulae arise in financial and business calculations, and we shall encounter several during the course of this text. In this chapter, we shall concentrate on some of the more complicated calculations that arise from the application of formulae.

Calculate the value of A from the formula

$$A = \frac{B(C+1)(3-D)}{(2E-3F)}$$

when $B = 2$, $C = 3$, $D = -1.6$, $E = -1$ and $F = -2.5$

Solution

$$(C+1) = (3+1) = 4$$
$$(3-D) = (3-(-1.6)) = 3+1.6 = 4.6$$
$$(2E-3F) = 2\times(-1) - (3\times(-2.5)) = -2+7.5 = 5.5$$
$$\text{Hence } A = \frac{2\times 4\times 4.6}{5.5} = 6.69 \text{ to two d.p.}$$

Exponential numbers

In Excel a number is raised to a power by using the symbol referred to as a carat (∧). Some practitioners refer to this symbol as being the exponential operator. Thus, for example, to cube 4 the formula required would be = 4 ∧ 3. To find the square of 4, the formula required is = 4 ∧ 2. The carat can also used to find the square root. In this case the formula would be = 4 ∧ (1/2), or to find the cube root the formula would be = 4∧ (1/3). The method is used to find the 4th root, 5th root and so on. Some examples are demonstrated below.

	A	B	C	D	E
1					
2		3 squared	9	9	
3		3 cubed	27	27	
4		3 to the power of 4	81	81	
5		4 to the power of 4	256	256	
6					
7		Square root of 4	2		
8		Cubed root of 64	4		
9					

Solving equations

Linear equations with only one variable

An equation is linear if it has no term with powers greater than 1, that is, no squared or cubed terms, etc. The method is to use the same techniques as in changing the subject of a formula, so that the equation ends up in the form variable = something.

(a) Solve $6 - 3X = 0$

$6 = 3X$
$X = 6 \div 3 = 2$

(b) Solve $200 = 5(X - 2) + 80$

$200 = 5X - 1080 = 5X + 70$
$200 - 70 = 5X = 130$
$X = 130 \div 5 = 26$

(c) Solve $-\frac{50}{X} = \frac{24}{X-3}$

Multiply up by the two denominators:

$50(X - 3) = 24X$
$50X - 150 = 24X$
$50X - 24X = 150 + 26X$
$X = 150 \div 26 = 5.77$ to two d.p.

Quadratic equations with only one variable

A quadratic equation has the form $aX_2 = bX = c = 0$ where a, b and c are constants. The equation can be solved using a formula but if either the bX or c terms or both are missing the formula is not necessary. Examples will be used to illustrate the methods.

Solve the following simple quadratic equations (note that the variable used is Y, but as there is only one variable used, this is fine.).

a) $4Y^2 = 100$
b) $Y^2 - 9 - 0$
c) $Y^2 + 2Y = 0$
d) $(Y - 5)^2 = 0$

Solution

a) $Y^2 = 100 \div 4 = 25$

$Y - +\sqrt{25}$ and $-\sqrt{25} - \pm 5$

b) $Y^2 = 9$

$Y = +\sqrt{9} = +3$

c) $Y(Y + 2) = 0$

Either $Y = 0$; or $Y + 2 = 0$, so $Y = -2$

d) The only solution is that $Y - 5 = 0$, so $Y = 5$

You may have noticed that most quadratic equations have two roots, that is, two values for which the two sides of the equation are equal, but occasionally, as in (d) above, they appear to have only one. It is, in fact, a repeated (or double) root. For example, $Y^2 = -9$ has no real roots. We shall consider this again when we look at quadratic graphs in the next chapter.

For quadratic equations all of whose coefficients are non-zero, the easiest method of solution is the formula. If the equation is $aX^2 + bX + c - 0$, then the roots are given by.

$$X = \frac{-b \pm \sqrt{(b^2 - 4ac)}}{2a}$$

This formula is given in your exam so you don’t need to learn it.

Solve the equation $X^2 - 50X + 600 = 0$

Solution

$a = 1; b = -50; c = 600$

$$X = \frac{-(-50) \pm \sqrt{(-50)^2 - 4 \times 1 \times 600}}{2 \times 1}$$

$$X = \frac{50 \pm 10}{2} = \frac{60}{2} \text{ and } \frac{40}{2}$$

$$X = 30 \text{ and } 20$$

Notice that the equation has real roots only if $b^2 - 4ac$ is positive, since negative numbers do not have square roots.

Simultaneous linear equations

These are equations of the type:

$$3X + 4Y = 18 \qquad \text{(i)}$$

$$5X + 2Y = 16 \qquad \text{(ii)}$$

which must both be satisfied by the roots X and Y. Provided you multiply both sides of an equation by the same amount, it continues to be true. In the solution of these equations, one or both of the equations are multiplied by numbers chosen so that either the X or the Y terms in the two equations become numerically identical. We have labelled the equations (i) and (ii) for clarity. Suppose we were to multiply (i) by 5 and (ii) by 3. Both equations would contain a 15 X -term that we could eliminate by subtraction, it being the case that you can add or subtract two equations and the result remains true. In this case, however, the simplest method is to multiply equation (ii) by 2, so that both equations will contain 4 Y and we can subtract to eliminate Y. The full solution is shown below.

$$3X + 4Y = 18 \qquad \text{(i)}$$

$$5X + 2Y = 16 \qquad \text{(ii)}$$

Multiply (ii) by 2:

$$10X + 4Y = 32 \qquad \text{(iii)}$$

Subtract (iii) − (i):

$$7X + 0 = 14$$

$$X = 14 \div 7 = 2$$

Substitute $X = 2$ into (i)

$$6 + 4Y = 18$$

$$4Y = 18 - 6 = 12$$

$$Y = 12 \div 4 = 3$$

Check the results in (ii):

$$5 \times 2 + 2 \times 3 = 16$$

The solution is $X = 2$, $Y = 3$.

Had we chosen to substitute $X = 2$ into equation (ii) it would not have affected the result but we would then have checked in the other equation (i).

Had we chosen to substitute X = 2 into equation (ii) it would not have affected the result but we would then have checked in the other equation (i).

Notes

Manipulating inequalities

Inequalities are treated in almost exactly the same way as equations. In fact an inequality says much the same thing as an equation, except that one side will be less than or greater than the other, or less than and greater than the other.

Inequalities can be manipulated in the same way as equations, except that when multiplying or dividing by a negative number it is necessary to reverse the inequality sign.

Example

Solve for x $3x + 10 > 40$

Solution $3x > 40 - 10$

$3x > 30$

$x > 10$

Example

Solve for x $5x + 20 < 60$

$5x < 60 - 20$

$5x < 40$

$x < 8$

Percentages and ratios

Percentages and ratios (or proportions) occur in many financial calculations. Basically, a percentage (denoted '%' or 'per cent') is expressed out of 100, whereas a ratio is one number divided by another. A simple example will illustrate.

Example

(a) Express 4.6 as:

(i) a ratio of 23.0;

(ii) a percentage of 23.0.

(b) Evaluate 30 per cent of 450.

(c) The ratio of the earnings from a certain share to its price is 18.5. If the price is £1.50, what are the earnings?

(d) If a variable, A , increases by 8 per cent, what does it become?

(e) If a variable, B , changes to 0.945 B , what percentage change has occurred?

$$\frac{4.6}{23.0} \text{ or } 0.2$$

(ii) In a basic example like this, the percentage is 100 times the ratio:

$$0.2 \times 100 = 20\%$$

(b) Thirty per cent is 30 out of 100. Thus, out of 450

$$\frac{30}{100} \times 450 = 135$$

(c) We are told that the ratio

$$\frac{\text{Earnings}}{\text{Price}} = 18.5$$

If the price is £1.50

$$\frac{\text{Earnings}}{£1.50} = 18.5$$

$$\text{Earnings} = £1.50 \times 18.5 = £27.75$$

(c) An increase of 8 per cent of A is

$$\frac{8}{100} \times A \text{ or } 0.08A$$

The variable therefore becomes

$$A + 0.08A = 1.08A$$

(e) If a variable has decreased by

$$B - 0.945B = 0.055B$$

As a percentage, this is

$$\frac{0.055B}{B} \times 100 = 5.5\%$$

The next example will demonstrate the use of percentages in financial calculations.

Example

(a) During a certain year, a company declares a profit of £ 15.8 m, whereas, in the previous year, the profit had been £ 14.1 m. What percentage increase in profit does this represent?

(b) A consultant has forecast that the above company's profit figure will fall by 5 per cent next year. What profit figure is the consultant forecasting for the next year?

(c) If this year's profit is £ 6.2 m, and if the increase from last year is known to have been 7.5 per cent, what was last year's profit?

Solution

(a) The increase in profit is £1.7m, which as a percentage of the previous year's profit is:

$$\frac{1.7\text{m}}{14.1\text{m}} \times 100\% = 12.1\% \text{ to one d.p.}$$

(b) The forecast decrease in profit is 5 per cent of £15.8m

$$\frac{5}{100} \times 15.8 = £0.79\text{m}$$

Hence, the forecast profit for the following year is £15.01m.

(c) This year's profit is 107.5 per cent of last year's

$$\text{Last year's profit} \times \frac{107.5}{100} = \text{This year's profit}$$

$$\text{last year's profit} = £6.2\text{m} \div 1.075 = £5.77\text{m to three s.f.}$$

Accuracy and approximation

All business data are subject to errors or variations. Simple human error, the rounding of a figure to the nearest hundred or thousand (or whatever), and the inevitable inaccuracies that arise when forecasting the future value of some factor, are examples of why business data may not be precise.

In certain circumstances, errors can accumulate, especially when two or more variables, each subject to error, are combined. The simplest such forms of combination are addition and subtraction.

Errors from rounding

Suppose an actual value is 826 and you round it to 830 (two s.f.). Your rounded value contains an error of 4. Someone else using the rounded figure does not know the true original value but must be aware that any rounded figure is likely to be erroneous.

The rounded value 830 could represent a true value as low as 825, or one as high as 835 (or, strictly speaking, 834.9999). There is a possible error of = 5. In general, rounded values have a possible error given by substituting = 5 in the position of the first discarded digit. For example, with a value of 830, the first discarded digit is in the position of the '0' , which is the units position. This gives a possible error of = 5 units. If the rounded figure were 82.391 (to three d.p.), the first discarded digit is immediately to the right of the '1' and the possible error is = 0.0005.

Example.

State the maximum possible errors in the following rounded figures:

(i) 67,000
(ii) 5.63
(iii) 10.095

Solution

(i) The first discarded digit is in the '0' position immediately to the right of the '7', so the maximum possible error is ±500.
(ii) The first discarded digit is immediately to the right of the '3', so the maximum possible error is ±0.005.
(iii) The first discarded digit is to the right of the '5', so the maximum possible error is ±0.0005.

Using Excel to produce graphs of Linear and Quadratic Equations

Excel can be used to produce graphs of linear and quadratic equations. The first step is to produce a single linear equation, from which a graph can be drawn.

Producing a single linear equation in Excel

The form of the equation that will be used is

$y = mx + c$

This equation will be drawn for a given value of c (in this example we will use 20) and a range of 10 values of x (from 1 to 10), calculating corresponding values of y . Thus in this example the formula will be represented as y = 3 x + 20.

Figure 1.7 shows the data for x and the results of entering the formula in the adjacent column.

Single linear equation		
	values for x	$y = 3x + 20$
	1	23
	2	26
	3	29
	4	32
	5	35
	6	38
	7	41
	8	44
	9	47
	10	50

To show these results graphically in Excel, select the two columns and click on the Chart icon on the Standard Toolbar. This will produce a choice of graph types. Select xy and then choose the joined up line option. Click Finish to complete the chart. Figure 1.8 shows the resulting graph.

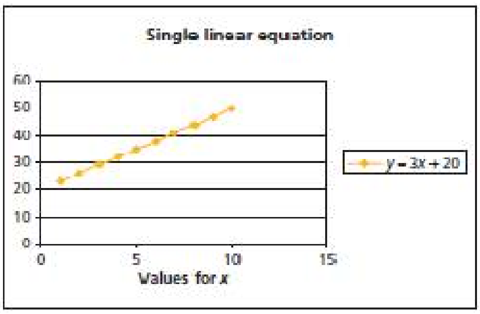

Graph showing single linear equation

Drawing multiple equations on a single graph

It is possible to produce multiple equations and plot the results onto a single graph, which can be useful for comparison purposes. Figure 1.9 uses the same set of data for x and the results of two different equations are shown in the adjacent two columns.

values for x	$y = 3x + 20$	$y = 6x + 1$
1	23	7
2	26	13
3	29	19
4	32	25
5	35	31
6	38	37
7	41	43
8	44	49
9	47	55
10	50	61

The graph is produced in the same way as the first example, by selecting the three columns and clicking on the Chart icon. The results of plotting these two lines onto an xy

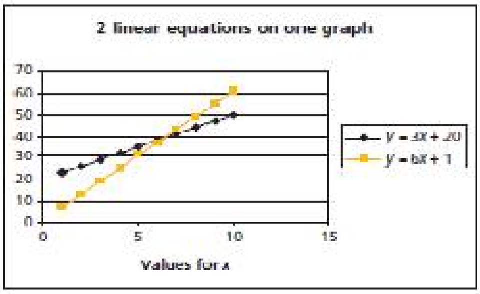

Single quadratic equation

The form of the equation that will be used is

$y = ax^2 + bx + c$

This equation will be drawn for a given value of a, b and c, where in this example we will use a =1, b = 5 and c =10 and a range of 10 values of x (from- 25 to 20), calculating corresponding values of y.

Thus in this example the formula will be represented as y= x^2+5 x+ 10. Figure 1.11 shows the data and the formula calculated in the adjacent column.

values for x	x^2 + 5x + 10
–25	510
–20	310
–15	160
–10	60
–5	10
0	10
5	60
10	160
15	310
20	510

Using the same method as before a graph can be drawn to show these results and this is shown in Figure below.

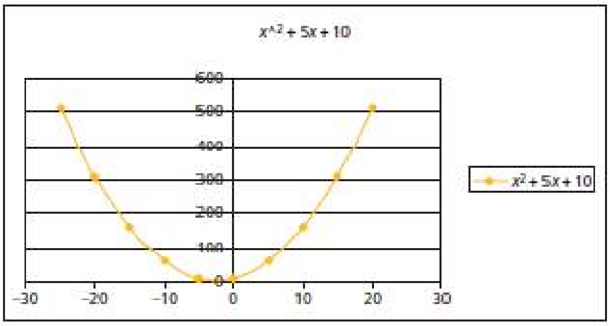

Two equations on one graph

It is possible to produce multiple quadratic equations and plot the results onto a single graph, which can be useful for comparison purposes. Figure below uses the same set of data for x and the results of two different equations are shown in the adjacent two columns.

values for x	$x^{\wedge}2+5x+10$	$-2x^{\wedge}2-x+100$
-25	510	-1125
-20	310	-680
-15	160	-335
-10	60	-90
-5	10	55
0	10	100
5	60	45
10	160	-110
15	310	-365
20	510	-720

Graph showing a single quadratic equation

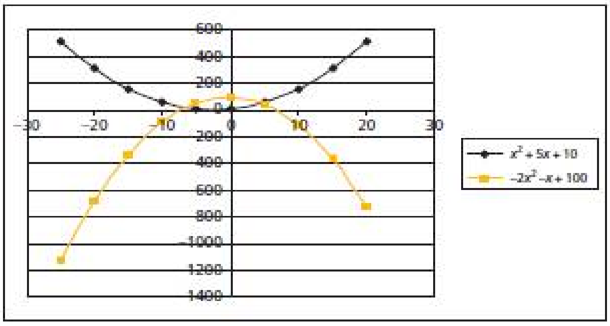

The graph is produced in the same way as the previous example, by selecting the three columns and clicking on the Chart icon. The results of plotting these two lines onto an xy line graph can be seen in Figure 1.14 .

Using Excel to produce the graph of a hyperbola

The formula of a hyperbola takes the form of y = a(1/x) = n. In the example below values of x from 1 to 27 are used. The constant a is 10 and a value of 1 was used for n. The results are shown in Figure below .

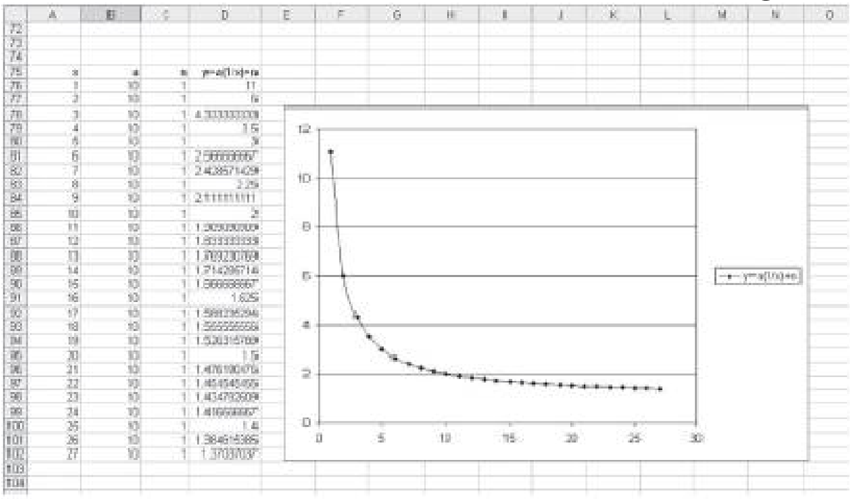

Basic mathematics covers a wide range of topics and underlies virtually all the elements of business mathematics. The key contents of the chapter are:

- The rules for the order of mathematical operations and the use of brackets – a source of many errors in calculations;
- dealing with negative numbers;

- rounding to the nearest number or to various numbers of decimal places or significant figures;
- dealing with powers and roots;
- manipulating formulae;
- dealing with exponential numbers;
- solving equations;
- manipulating inequalities;
- dealing with percentages and ratios;
- rounding errors;
- creating graphs in a excel to draw linear and quadratic equations and graph of a hyperbola.

REVIEW QUESTIONS

1. The number 268.984 is to be rounded. In each case write the correct answer, to the accuracy specified.

(A) to two decimal places.
(B) to one decimal place.
(C) to the nearest whole number.
(D) to the nearest 100.
(E) to three significant figures.
(F) to four significant figures.

2. Evaluate the following without rounding.

(A) 7 + 2 X 5
(B) (5 + 2) X 8
(C) 28 – 48/4
(D) (7 + 3)/5
(E) 8+ 4 X5 - 2
(F) (8 -4)X (3+ 7).

3 A small company produces specialised posters. The total cost is made up of three elements = materials, labour and administration = as follows:

- Materials £ 0.50 per poster
- Labour £ 15 per hour
- Administration £ 10 per 100, plus £ 50.

4. The set-up time for printing takes two hours, and the posters are run off at the rate of 300 per hour.

5. If the number of posters produced is denoted by N, write down the formulae for the following in £ :

(A) the total cost of materials,
(B) the time taken to produce the posters,
(C) the total labour costs of production,
(D) the administration costs.

6. Calculate the solutions to Question 1.2.1 using Excel. You will first have to create a spreadsheet to represent the data for the printing company.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-2 OBTAINING DATA

CONTENTS

INTRODUCTION

Many of the problems that accountants face require the acquisition, communication or analysis of data as part of their solution. We look at each of these aspects in turn. First of all, there is the question of how to obtain data, as the individual facts and figures are known. Data can be classified in two ways: primary and secondary. Primary data is that collected specifically for the problem in hand, while secondary data is collected (by others) for some other purpose. Thus an accountant, working in the budgeting department of a manufacturing company, might get information on raw material costs by contacting the suppliers himself or herself, and so obtain primary data. Alternatively, he or she could use secondary data in the form of a list of quotations compiled for its own purposes by the company's buying department. Primary data is the more reliable, since you have obtained it yourself (or have had it collected) and because it relates precisely to the particular problem you are facing. Its actual collection, however, does take time; obtaining it, therefore, tends to be costly, and there may be a considerable delay before the information is ready to use.

On the other hand, secondary data, if available, is relatively inexpensive and quick to obtain: often simply a reference to some relevant publication. The disadvantages here arise from the possibility that there may be no suitable sources for the information. Even if there are, the data may not match your requirements too well. In addition, although official or government statistics may be considered reliable, other secondary sources may not.

Obtaining Data

Immediately after collection, in what is often termed its raw form, data is not very informative. We cannot learn about the situation from it or draw conclusions from it. After it has been sorted and analysed, data becomes information that, it is to be hoped, is understandable and useful.

The difference between information and data

Sometimes the issue of the quality of data is raised and often there is not a clear understanding of this issue. Quality data has several characteristics including being:

- error free;
- available at the right time;
- available at the right place;
- available to the appropriate individuals.

The arrival of the Internet has made it much easier for organisations and individuals to access data at the right time and the right place. However, at the same time the Internet have opened up questions about data being error free and about who can have access to it. As well the issue of data quality there is the question of how data, information and knowledge relate to one another. Russell Ackoff was one of the first people to speak of there being a hierarchy which he referred to as the Data Information Knowledge Wisdom (DIKW) Hierarchy. According to this model, data (which is by the way sometimes said to be a plural word as it is the actual plural for the word datum) are simple facts or figures or maybe even a photograph or an illustration. In this form data is unstructured and uninterrupted. Information comes from processing or structuring data in a meaningful way. Another way of looking at this is that information is interpreted data. An interesting story is told by Joan Magretta in her book What Management is ? about Steve Jobs which clearly illustrates the difference between data and information. Despite its small share of the total market for personal computers, Apple has long been a leader in sales to schools and universities. When CEO Steve Jobs learned that Apple’s share of computer sales to schools was 12.5 per cent in 1999, he was dismayed, but unless you’re an industry analyst who knows the numbers cold, you won’t appreciate just how dismayed he was. That’s because, in 1998, Apple was the segment leader with a market share of 14.6 per cent. And, while Apple slipped to the number two spot in 1999, Dell grew and took the lead with 15.1 per cent. Alone each number is meaningless. Together they spell trouble, if you’re Steve Jobs, you see a trend that you’d better figure out how to reverse. This isn’t number crunching, it’s sense making. In this example the 12.5 per cent was data and when it was seen in conjunction with the 15.1 per cent it became information.

Knowledge is again different to data and information. Knowledge is much more personal and the presence or absence of knowledge can normally only be seen through the actions of individuals. When knowledge is written down it effectively becomes information.

Finally with respect to wisdom it is difficult to define this concept. Wisdom has something to do with understanding or insight. It is to do with achieving a good long-term outcome in relation to the circumstances you are in.

Primary data: sampling

We begin by considering how primary data can be obtained. All the data relating to a problem is known as the population . For reasons of finance and practicability, it is rarely possible to obtain all the relevant data, so you normally have to use only part of the population, that is, a sample . It is clear that, if the sample data is to be of any use, it must be representative of the population as a whole. If a sample is representative of its population, it is said to be unbiased ; otherwise it is biased . Another fundamental point to note is that, because a sample is only a subset – that is, some information has been omitted – any results arising from it cannot be exact representations of the whole population. This deficiency is said to constitute sampling error .

By careful choice of sampling method, it is possible to ensure that a sample is representative of its population, thereby avoiding bias. Since the very act of sampling omits some of the data, sampling error is inevitable. In general, if you increase its size, the sample will represent a larger proportion, and so will be 'nearer to' (or more representative of) the population. Increasing sample size thus tends to reduce sampling error. An example will illustrate these new concepts and terms.

Example

You work as an assistant to the chief accountant of a company that owns a large number of retail stores across the country. The chief accountant asks you to provide him with some up-to-date information on the weekly turnover figures of the stores. Discuss how you would set about this.

Solution

Secondary financial data on the stores will no doubt be available, but it may not consist of weekly turnover figures and probably will not be up to date. Thus, provided that enough time and resources are available, you should consider obtaining primary data. We shall leave discussion of possible methods of actually collecting the information until later in this chapter and concentrate here on the meaning of the various concepts defined above, as applied to this example.

The population here consists of all the recent weekly turnover figures for all the stores owned by the company. Clearly it would be practically impossible to collect all this data, and so a sample will have to be taken. This will consist of a selection of a certain number of the weekly turnover figures, for example 100 or 1,000 of them: exact number will depend on the time and resources available, and the level of accuracy required by the chief accountant. If the selection of 100 or 1,000 (or whatever) weekly turnover figures is representative of all the possible data, then we shall have an unbiased sample. However, because a sample consists of only part of the population (possibly just a small proportion), it will give only an approximation to the overall picture: sampling error will be present. In general, increasing the sample size from, say, 100 to 1,000 will reduce this error.

Probability sampling methods

We now look at various ways of selecting samples, beginning with probability sampling methods (also called random sampling methods): those in which there is a known chance of each member of the population appearing in the sample. Such methods eliminate the possibility of bias arising through (subjective) human selection: it can now only arise by chance. In addition, it is possible to undertake calculations concerning the effects of sample error when probability sampling methods are used.

Simple random sampling

Of these methods, the most basic is simple random sampling , in which each element of the population has an equal chance of being selected. Such a sample could be chosen by drawing names out of a hat, or by employing more sophisticated models using random numbers (which will be illustrated later). As we shall see in later parts of the text, random sampling is one of the most important sampling methods in statistics, since the accuracy of estimates made from the sample can be calculated.

Stratified random sampling

In this case, the population is divided into coherent groups or strata, and the sample is produced by sampling at random within each stratum. This process takes more time than simple random sampling, but should result in more representative samples and hence should reduce the sample error.

Example

How would a simple random sample and a stratified random sample be drawn in the situation described in Example?

Solution

To begin with, we shall need a list of the population. If, for example, the company owns 100 stores, and the investigation is confined to the last year's (50 trading weeks') trading, then the population will consist of 5,000 pieces of data. These might be arranged as follows:

Turnover of store 1 week 1 0001
Turnover of store 1 week 2 0002
. . . .
Turnover of store 1 week 50 0050
Turnover of store 2 week 1 0051
. . . .
Turnover of store 100 week 50 5000

The members of the population can be numbered from 1 to 5,000, using a four-digit notation, as shown in the right-hand column of the above table. The table is not set out in full for reasons of space: however, it is easy to calculate that the turnover of store 7 in, say, week 47 is represented by the number 347 (i.e. 6 x 50 + 47). The problem of selecting, say, 200 of these weekly turnover figures, each with an equal chance of being picked, can now be translated into the task of obtaining a similar sample from the digits 0001 to 5,000.

There are a number of ways of obtaining the necessary random numbers, such as computer generation or by using random number tables. The latter method consists of reading off four-digit numbers from the tables. (Because there are 5,000 elements in the population, we ensured that each one was numbered with a four-digit number: hence 0001, etc.) For example, the random numbers 0347, 4373 and 8636 would be interpreted as follows:

0347: week 47 of store 7 is the first member of the sample
4373: week 23 of store 88 is the second
8636: this is bigger than 5,000, and so is ignored.

We proceed in this way until we have selected the desired sample size (200 items). Note that because each digit in the table has an equal chance of being 0, 1, 2, …, 9, we have ensured that each element of the population has an equal chance of being sampled, that is, we have a simple random sample.

Following the above procedure, it is possible, by pure chance, that the 200 sampled weeks may come from just a few stores' figures, with many stores not being represented. If the 100 stores are similar in their trading patterns, this may not matter, but, if we wish to ensure that there is a good spread from the stores within the sample, we can stratify before sampling. By way of illustration, this can be done by taking each store's figures as a stratum and taking a sample of 2 weeks within each. Thus, the data on store 1 would form the first stratum:

Turnover, week 1 01

Turnover, week 2 02

. . . . . .

Turnover, week 50 50

Note that we need only two-digit random numbers now. The random numbers 33, 67, 00, 98 and 09 would be interpreted as follows:

33 the turnover of week 33 is selected

67 bigger than 50, so this is ignored

00 too small (01 is the lowest in our numbering)

98 too big, so ignored

09 the turnover of week 9 is selected.

These two figures would form the contribution to the sample from the first stratum. Repeating this procedure for the 2nd to the 100th stratum (store's figure) would now produce a randomised sample in which every store was represented.

Before we leave this example, it should be noted that there are many other ways of stratifying. For example, it may be that stores number 1–10 are far bigger than the others and so should be better represented in the sample. In such a case, you might work with just two strata (larger stores' turnover figures and small stores' turnover figures) and then sample (at random) 100 values from each. In practice, situations like this demand the use of personal judgement when determining how to stratify and what proportion of the sample to include from each stratum. In other situations, it is possible to be a little more precise, as the following example illustrates.

Other sampling methods

If we reconsider the examples in the preceding section, we shall see that there can be great practical problems in using random methods of sampling. First of all, in order to use random number tables (or any other method of ensuring randomness), we need a list of the population. This list is often called a sampling frame. As we shall see, there are many instances where a sampling frame is unavailable. Even when one is, stratified random sampling may be impossible. For instance, in the previous example, to stratify the population it would be necessary to divide the sampling frame (possibly the town's electoral register) into the four categories shown. This would not be possible, since an individual's age could not easily be determined from the register. Note that it would not be sensible to try to find the age of every individual in the town, as one would then be contacting every member of the population – the whole idea of sampling is to avoid such an onerous task! A second practical problem lies in the cost of random sampling. Imagine taking even a simple random sample of 1,000 people from a town's electoral register. This would involve counting down the register to find person number 139,103 (say), which our random number tables gave us, then contacting her or him to conduct an interview. Multiply this by 1,000 and you will appreciate the immense time and expense involved. A number of alternative sampling methods have been devised to get round some of the problems listed above.

- Cluster sampling consists of taking one definable subsection of the population as the sample. Thus, you could take the inhabitants of one street as a cluster sample of the population of a town. However, such a sample is unlikely to be representative, and so a variation might be to take the inhabitants of five streets (chosen at random from an alphabetical list) as the sample. The latter example could still be unrepresentative but would be a great deal easier and cheaper to survey than a sample randomly spread all over the town. Cluster sampling is a random method provided that the clusters are randomly selected, but it does not give data as reliable as that given by either simple random or stratified random sampling. It is widely used for reasons of speed, cheapness and convenience.
- Systematic sampling involves taking every nth member of the population as the sample. This can be applied with or without a sampling frame, and so may or may not be a probability method. As an example, a sample could be drawn from an electoral register by selecting every 1,000th on the list. In a quality-control situation, we could take every 100th batch coming off a production line for testing. Note that, in this latter case, we have only part of the sampling frame, those batches produced near the time of sampling. A complete sampling frame, consisting of all past and future batches, would be impossible to obtain. Systematic sampling from a sampling frame provides a good approximation to simple random sampling without the bother of using random numbers. Problems of bias arise only if there is a cycle in the data that coincides with the sampling cycle. For instance, if, in the quality-control situation, 100 batches were produced every two hours, then you might find yourself always sampling the last batch prior to staff taking a break: this might not be representative of the general quality of production. In general, however, systematic sampling provides random samples almost as reliable as those of simple random sampling and in a much more convenient manner.
- Quota sampling is essentially non-random stratified sampling. The members of the various strata are called quotas , which are not chosen at random but are selected by interviewers. The market research agency in the previous example could draw its sample by issuing quotas of:

347 females, 35 and over
163 females, under 35
333 males, 35 and over
157 males, under 35

The actual members of the sample would be selected by the interviewers, as they moved around the town. When the respondent's age was determined, he or she could be included as part of one of the above quotas. When a quota was complete, no more people from that category would be included. This subjective element is an obvious source of bias which can be reduced in practice by training interviewers to choose a 'spread' within each quota. For example, they would be encouraged not to choose all 40-year-olds in the '35 and over' quotas, but to try to achieve a variety of ages within the range. This method is cheap, quick and easy and, in particular, it does not require a sampling frame. It has all the advantages of stratified sampling at a much lower cost, but it is not random: information obtained from it may well be biased and there is no means of measuring its reliability.

Multistage sampling

We conclude the discussion of sampling techniques by looking at a method commonly used when a survey has to cover a wide geographical area without incurring great expense. Multistage sampling , as its name implies, involves splitting the process into a number of (typically three) separate steps. An example will illustrate.

Example

If you were organising a nationwide opinion poll, how would you set about organising the sample, using a multistage technique?

Solution

The first stage is to divide the country into easily definable regions: in the case of a political survey such as we have here, the 651 parliamentary constituencies are ideal for this. It is now a straightforward matter to select, at random, twenty (say) of the constituencies.

In the second stage, the regions are split into smaller, more manageable, districts. There is an ideal subdivision in this example, namely the political wards within each constituency. A random sample of (say) three wards might now be selected within each of the twenty constituencies obtained in the first stage. Note that we have now a sample of sixty wards, randomly selected, and that these could be obtained by one person in a matter of minutes, provided that a complete list of constituencies and wards and a set of random number tables are available.

The time-consuming stage is the third and final one, that of contacting and interviewing a sample of (say) thirty voters in each of the sixty wards. This could be done at random from the electoral registers of the wards, thereby ensuring that the whole process is random. A faster and less expensive alternative, albeit one that risks the introduction of bias, would be a non-random method such as quota sampling. The quotas might be by gender and age or, working with a larger budget, might be more sophisticated.

This sample involves two stages of cluster sampling and a third that may be either systematic or quota sampling.

One criticism of multistage sampling is that the regions in the first stage and the districts in the second stage will not each contain the same number of people. In the above example, one constituency may have 40,000 electors, whereas another may have 60,000. As the two are equally likely to be chosen in stage 1, an elector in the former constituency will be 1.5 times as likely to appear in the final sample than an elector in the latter constituency (one in 40,000 as opposed to one in 60,000).

A variation that redresses this imbalance is sampling with probability proportional to size. Hence, we weight the chances of choosing a region (stage 1) and then a district (stage 2) according to the number of people in the region or district; larger regions or districts having proportionately more chance of selection than smaller ones. In the above example, we would ensure that a constituency of 40,000 electors has two-thirds the chance of being selected as does a constituency of 60,000 electors, by allocating two-thirds as many random numbers to the former as compared with the latter. This proportion will then exactly compensate for the imbalance referred to in the preceding paragraph.

You will virtually certainly have a question in your assessment which requires you to differentiate between the different types of samples.

Secondary data: sources

There are numerous potential sources of secondary data that could be used to provide the required information in a given situation. Searching out sources is usually not too problematical: the real difficulty lies in judging whether the data adequately matches the requirements or whether primary data should be sought.

Sources of secondary data can be categorised into three types. First of all, there are data collected and compiled internally by the organisation, such as its financial reports and accounts, personnel records, and so on. Second, there are business data produced by sources external to the organisation. Under this heading come the results of surveys by the CBI, the financial press and similar sources. Finally, we have the many government produced statistics on a whole range of commercial and demographic topics, any one of which might be applicable in solving a business problem.

These publications are too numerous to list here, but the Office for National Statistics Guide to Official Statistics is published annually and gives a comprehensive catalogue of sources of official statistics.

Questionnaires

In many situations, data are collected by asking a sample of people a series of questions. The printed form used by interviewers to pose the questions or for respondents to complete themselves is termed a questionnaire . Although it is unlikely that management accountants will have to design questionnaires, they may have to commission their design, or may encounter them in some other way. It is therefore useful to be able to judge whether a questionnaire is well designed; as poor design could lead to the collection of unreliable data.

The basic ideas are to make it as easy as possible to answer the questions accurately and to encourage respondents to complete the questionnaire, if interviewers are not being used. Overall, therefore, the questionnaire needs to be as brief as possible, consistent with the data that needs to be collected: long documents can be off-putting. In addition, it should be logically structured and well laid out, otherwise errors can be introduced through confusion or, again, respondents may be discouraged from completing the forms. The individual questions, in particular, need attention. There are a number of do’s and don’ts to be considered when drafting questions, all again concerned with obtaining accurate and reliable responses from the sample. These are listed below.

- Do not ask ambiguous questions.
- Do not use leading questions.
- Do not pose questions that require technical knowledge or that use a complicated vocabulary.
- Give a brief, simple list of possible answers, whenever you can.
- Put personal or difficult questions at the end of the questionnaire.
- Do not ask questions that rely too much on memory.
- Try to avoid open-ended questions.

Contact with respondents

The following methods of making contact with respondents are all in wide use and have different advantages and disadvantages.

- Interviewers undoubtedly get the best results. The response rate is high, the questionnaire is filled in immediately and in an accurate manner, and any misunderstandings on the part of the respondent can be rectified. However, the method is not without disadvantages. It is time-consuming and costly. The interviewer needs to be well trained in order to ensure that he or she does not introduce interviewer bias by posing questions in such a way that respondents are inclined to give certain answers. Respondents may not be prepared to discuss certain topics in an interview, but might be prepared to complete a questionnaire privately. Finally, in interviews the respondent cannot take time to 'mull over' the questions. Indeed, they may answer quite thoughtlessly in order to speed up the interview. Notwithstanding the disadvantages, face-to-face interviews are generally regarded as the best method of contact with respondents because low response rates and poorly completed questionnaires are so devastating in surveys.
- Enumerators deliver the questionnaires and explain to the respondents what the survey is about and so on. A few days later they collect the completed questionnaires and deal with any problems that have arisen. The method secures quite good response rates, although not as good as face-to-face interviews; it gives the respondent privacy, the time to think about questions and some degree of assistance and it reduces the likelihood of interviewer bias. Its disadvantages are that it is costly and time-consuming and that the forms may not be filled in very well.
- Telephone interviews are the other method in which there is some personal contact. It is easier to refuse an interview if you are not face-to-face, so telephone interviews suffer from a still lower response rate. The method has most of the disadvantages of face to face interviews, namely interviewer bias, lack of privacy and lack of time to reflect on questions, plus the additional problem that a proportion of the population, especially among the elderly and the poor, do not have telephones. On the other hand, telephone interviews avoid the costs and problems of having to travel and are not as time consuming or as costly as interviewing face to face.
- Postal surveys include any methods of delivering questionnaires without making personal contact with respondents, such as leaving questionnaires on desks or in pigeonholes as well as sending them by post. They suffer from very low response rates and are frequently poorly

completed. They are, however, undoubtedly the cheapest and easiest method, they are free from interviewer bias and they give the respondent privacy and the option to take time over answering. Postal surveys are the only method that can absolutely guarantee confidentiality, since in all other cases respondents cannot be certain that their names will not be associated with the completed questionnaire. However, inducements such as small gifts or entry into a lottery are often given in order to overcome the low response problem and hence confidentiality is often lost.

We saw earlier that questionnaires can typically be used in interviewer-based or postal surveys. Each type of survey has its own implications when it comes to the design of the questionnaire to be used. For example, interviewers can explain an unclear point on a form, or can (with good training) 'probe' for deeper answers. With a postal survey, therefore, even greater care is needed when drafting questions, and some types of 'probe' or follow-up question may not be possible at all.

Importing data to Excel

Having identified data to be used for an application that can be managed in Excel it is sometimes possible to import that data from its original source as opposed to having to re-enter the information directly into the spreadsheet through the keyboard. Excel has a number of importing options, some of which are described here.

Importing data from Word

The data shown in Figure 2.2 has been entered into Word using the Tab key to space it out. This information can be copied and pasted directly into Excel and the tabs indicate to Excel where the cell change occurs. Figure 2.3 shows the data after copying into Excel.

Month	Mon	Tues	Wed	Thurs	Fri
Jan	570	539	580	563	497
Feb	520	1480	510	1500	490
Mar	1562	588	502	516	540
Apr	568	516	550	1562	556
May	1555	562	548	548	1554
June	562	1553	560	498	554
July	562	553	1575	539	531
Aug	1586	567	509	529	587
Sept	596	577	574	555	580
Oct	569	1550	557	558	563
Nov	562	519	569	1530	560
Dec	567	553	524	501	1550

	A	B	C	D	E	F	G	H
1								
2								
3		Month	Mon	Tues	Wed	Thurs	Fri	
4		Jan	570	539	580	563	497	
5		Feb	520	1480	510	1500	490	
6		Mar	1562	588	502	516	540	
7		Apr	568	516	550	1562	556	
8		May	1555	562	548	548	1554	
9		June	562	1553	560	498	554	
10		July	562	553	1575	539	531	
11		Aug	1586	567	509	529	587	
12		Sept	596	577	574	555	580	
13		Oct	569	1550	557	558	563	
14		Nov	562	519	569	1530	560	
15		Dec	567	553	524	501	1550	
16								

Using the Excel Text to Columns feature

Sometimes it is not possible to copy and paste information from another source into Excel and have the data automatically drop into cells correctly. This might be when data is being selected from a database system or an accounting system. In such cases it is usually possible to instruct the external source to export data with a pre-defined separator, sometimes referred to as a 'delimiter' . Figure 2.4 shows some data that has been separated by commas.

```
Jones,34,car,45,finance
Brown,42,bicycle,8,sales
McDuffy,32,car,20,sales
Greggory,23,walk,2,admin
Hafeez,29,car,15,personnel
Bundi,54,bicycle,10,finance
```

When this data is copied into Excel it appears as shown in Figure 2.5 . The Excel Data Text to Columns command can now be used to separate this data into different cells. First select the data in the range B4:B9 and then select Data Text to Columns. A dialogue box is displayed on the screen and in this example the delimiter should be set to comma. Click next and next again and then finish.

In obtaining data, various decisions have to be made. You should understand what the following mean and be able to discuss their relevance and advantages and disadvantages.

- What resources of staff, time and money are available?
- Will you use primary or secondary data?
- Will you sample or survey everyone?
- If sampling, is there a sampling frame?
- Can you stratify?
- Depending on the above, what sampling method will you use? You should know the relative merits of simple random, stratified random, systematic, cluster, multistage and quota sampling.
- How will you approach the respondent? You should know the advantages and disadvantages of using interviewers, enumerators, telephone interviews and postal questionnaires.
- How will you design the questionnaire? You should know the pitfalls to avoid and the features of good design to suit your chosen method of contact with respondents.
- Having obtained the data, subsequent chapters will show you how to turn it into information.

REVIEW QUESTIONS

Q1-Describe the Difference between Information and Data.

Q2-What are the sources of Primary Data? Describe Sampling.

Q3-What are Probability Sampling Methods? Describe.

Q4-What are the Sources of Secondary Data?

Q5-Define Questionnaires and its use in data collection.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-3 PRESENTATION OF DATA

CONTENTS

INTRODUCTION

Data, when first collected, are often not in a form that conveys much information. Such raw data, as they are called, may just consist of a list or table of individual data values: if the list or table is of any appreciable size then it may need some refinement before anyone can draw conclusions from it. In this chapter we look at ways in which raw data can be collated into more meaningful formats, and then go on to see some pictorial representations of data that provides convenient ways of communicating them to others. We begin by looking at linear and other important graphs.

Linear graphs

For each function of one variable there is an associated graph . This is a pictorial representation in which every value of x , with its associated value of y , is shown. In order to do this, pairs of values of x and y are plotted on graph paper as illustrated in the example below.

Example

Calculate the values of y in the function

$y = 3 + 2x$

corresponding to the values: x= 2; x = 1; x= 3; x = 4; x= - 1.

Plot the five corresponding pairs of values on a graph and hence draw the

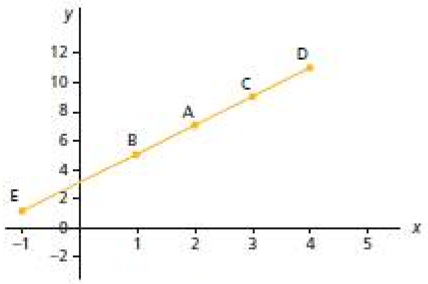

Figure 3.1 Plot of points and line

graph of the function.

Plot of points and line

Solution

Substituting x = 1 gives

$y = 3 + 2 \times 1 = 3 + 2 = 5;$

In the same way, it can be seen that x= 3 gives y= 9; and x= 4 gives y = 11.

In the final case, x= - 1 gives

$y = 3 + 2 \times -1 = 3 - 2 = 1$

The five pairs of values are shown as the points A, B, C, D and E, respectively, in Figure 3.1 . Two things should be noted about this figure:

1. Values of x are always measured in a horizontal direction, along the x-axis: positive values to the right, negative to the left. Values of y are measured in a vertical direction, along the y-axis: positive values upwards, negative downwards. Thus, the point A is plotted by moving from the point where the axes cross (the origin, where both variables have a value of zero).

- 2 to the right (plus 2);
- 7 upwards (plus 7).

These values are known as the x- and y-coordinates of A, respectively. Point E has a negative x-coordinate of = 1 and a positive y -coordinate of 1, and so is plotted as 1 to the left and 1 upwards.

- Note the scales on the two axes. These distances are marked off on each axis as an aid to plotting. They need not be the same as each other (indeed, they are not in this case) but they must be consistent. In other words, if you decide that one square on your graph paper equals one unit (or whatever) along the x-axis, you must keep this scale throughout the x-axis. The scales should be chosen so that the range of x and y coordinates can be displayed appropriately.

We can now join the points up, in order to form the graph of the function. In this case, the points lie exactly on a straight line, and so the joining up can be done best with a ruler. This is an example of a linear graph.

Solving simultaneous linear equations using graphs

We have seen that breakeven occurs when the cost and revenue graphs cross. This is the graphical interpretation of the solution of simultaneous linear equations, and a graphical method could be used instead of an algebraic method (provided that the scale was big enough to give the required accuracy).

Example

Solve the simultaneous equations

$$2x + 3y = 8 \qquad \text{(i)}$$

$$5x - 2y = 1 \qquad \text{(ii)}$$

by graphing the lines using the values x = 0 and x= 5.

If you feel that you need more practice solving such equations algebraically, you could do that as well. The algebraic solution of simultaneous linear equations is often slightly easier and lends itself to greater accuracy (as you are not reading from a graph).

Solution

n (i), when $x = 0$, $3y = 8$ so $y = 8 \div 3 = 2.67$. When $x = 5$, $3y = 8 - 10 = -2$, so $y = -2/3$.
In (ii), when $x = 0$, $-2y = 1$, so $y = -1/2$. When $x = 5$, $-2y = 1 - 25 = -24$, so $y = -24 \div -2 = 12$.
These values are plotted in Figure 3.6.

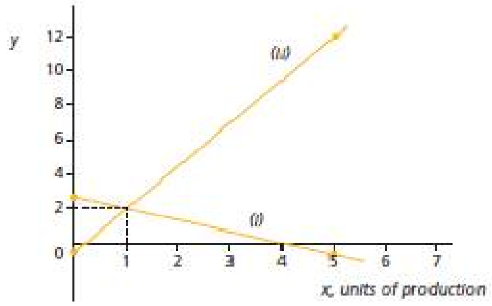

Quadratic graphs

We saw esrliar that the general quadratic function is:

$$Y = aX^2 + bX + c$$

It has one of two basic graph shapes, as shown in Figure below

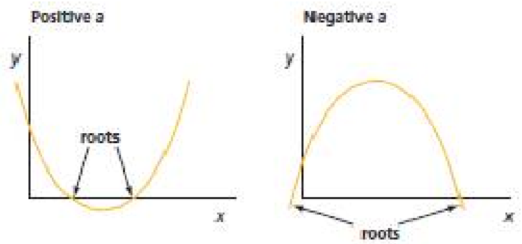

It is a symmetrical ‘U’ -shape or ‘hump’ -shape, depending on the sign of a . Another way of saying the same thing is that the y -values drop to a minimum and then rise again if a is positive, whereas they rise to a maximum and then fall if a is negative. This is of some importance in your later studies, when you may need to investigate the maximum or minimum values of functions for profits or costs, etc. The roots or solutions of the equation

$$aX^2 + bX + c = 0$$

are given by the intercepts on the x -axis, and by symmetry the maximum or minimum is always halfway between them. Some quadratic equations do not have real roots, and in these cases the graph simply does not cut the x -axis at all = as shown in Figure above.

Quadratic equation with no roots

With computer-based assessment you cannot at present be required to actually draw graphs, so questions are likely to ask for the labels of axes, coordinates of particular points and the information given by particular points.

Tallying frequency distributions

An example will illustrate a simple way of converting raw data into a concise format.

Example

In order to monitor the efficiency of his department, the head of the finance section of a large company spot-checks the number of invoices left unprocessed at the end of each day. At the end of the first period of this check (26 working days), he has collected the following data:

1	5	3	3	2
3	0	4	1	4
3	3	2	1	2
1	1	0	3	6
5	0	3	4	2
3				

Collate this raw data into a more meaningful form.

Solution

By scanning the table we can see that all the values lie between 0 and 6 inclusive. It might be useful to find out how often each value in this range occurs in the table. This could be achieved simply by counting, but there are no safeguards against human error in doing this. Instead we use a tallying procedure, which is more accurate than counting, especially with large tables of figures. After going along the first row, the tally will look like:

Number of invoices left unprocessed	*Tally*
0	
1	I
2	I
3	II
4	
5	I
6	

counting, when each fifth notch is reached, it is marked thus: IIII Number of invoices left

Number of invoices left unprocessed	*Tally*	*Total*
0	III	3
1	~~IIII~~	5
2	IIII	4
3	~~IIII~~ III	8
4	III	3
5	II	2
6	I	1
		26

The 'totals' in the above table are called frequencies and the table is called the frequency distribution of the sample. Thus the frequency of 0 invoices is 3 and so on.

Discrete and continuous variables

There is an essential difference between the variables considered in Examples 3.5.1 and 3.5.3. The former is discrete, whereas the latter is continuous. That is to say, the number of invoices can only consist of certain values:

0 or 1 or 2 or …

but never 1.6, 2.3 and so on.

On the other hand, the time taken to undertake a certain operation can theoretically take a value to any level of precision:

20.2 minutes

20.19 minutes

20.186 minutes

20.1864 minutes and so on.

In Example 3.5.3, the management services staff chose to measure to one decimal place: theoretically, they could have chosen to measure to two, three or any number of places. A number of invoices cannot be measured any more accurately than in whole numbers. This distinction has a number of consequences. Here, it can affect the way we tally. Continuous variables, such as the times to undertake a certain operation, can rarely be tallied as individual values, since few of them will coincide to give meaningfully large frequencies. Classifying is therefore almost always necessary with continuous variables. As Example 3.5.1 demonstrated discrete variables can sometimes be tallied with single values. However, with a wider range from (0 to 100, for example), the problem of having frequencies being mostly 0, interspersed with a few 1 s, could still arise: it is therefore sometimes necessary to use classes for discrete data too.

There are numerous ways of classifying when it is necessary. For instance, in Example 3.5.3, we could have used

15	to	17	
17.1	to	19	
19.1	to	21	and so on

15	to	16.95	
16.95	to	18.95	
18.95	to	20.95	and so on.

Both of these could be problematical if the measurements were later taken to the nearest twentieth (0.05) of a minute, as we should have difficulty placing 17.05 minutes in the former classification and 18.95 minutes in the latter. For this reason, we recommend that continuous variables are always classified:

15 to under 17

17 to under 19 and so on.

PRESENTATION OF DATA

Example

At the factory mentioned in Example 3.5.4 the daily outputs of a different product (Q) are measured to the nearest kg and are recorded as:

Daily output, kg

			383	351
362	377	392	369	351
368	382	398	389	360
359	373	381	390	354
369	375	372	376	361

Tally these data into a frequency distribution using the intervals 350–under 360; 360–under 370 and so on.

Solution

Output of Q kg	No. of days (frequency)
350–under 360	4
360–under 370	6
370–under 380	5
380–under 390	4
390–under 400	3

Cumulative frequency distribution

It is sometimes helpful to develop the idea of frequency further and to look at cumulative frequencies. These are the number of data values up to – or up to and including – a certain point. They can easily be compiled as running totals from the corresponding frequency distribution, as the following will illustrate.

Example

Form the cumulative frequency distributions from the data given below Hence estimate:

(a) how often there are more than four invoices left unprocessed at the end of the day;

(b) how often the time taken beats the target of 23 minutes.

Solution

The frequency distribution of the number of unprocessed invoices can be used to obtain:

Number of invoices left unprocessed (less than or equal)	Cumulative frequency	
0	3	(simply the frequency of '0')
1	8	(i.e. 3 + 5)
2	12	(i.e. 8 + 4)
3	20	
4	23	
5	25	
6	26	

In the same way, for the distribution of times taken to undertake the operation:

Time (minutes) (less than)	Cumulative frequency	
15	0	(no values below 15 minutes)
17	3	(frequency of the first class)
19	8	(i.e. 3 + 5)
21	18	(8 + 10)
23	25	
25	29	
27	30	

In the latter example, we have to take the upper limit of each class to ensure that all the values in the class are definitely less than it. We must use 'less than' as opposed to 'less than or equal' here because it corresponds to the way the frequency table has been compiled.

It is now a simple matter to estimate:

(a) 26 - 23 =3 occasions out of 26: that is, 11.5 per cent;

(b) 25 occasions out of 30: that is, 83.3 per cent.

How reliable these estimates are depends on how typical or representative is the period or month in which the samples are taken.

We shall see further applications of cumulative frequency in the following section.

Histograms and ogives

Many people find it easier to understand numerical information if it is presented in a pictorial form, rather than as a table of figures. In this section, therefore, we look at diagrammatic representations of frequency and cumulative frequency distributions. A histogram is a graph of a frequency distribution. The x -axis is the variable being measured and the y -axis is the corresponding frequency. It differs from the graphs drawn earlier since, in the examples so far, the frequency is represented by the height of a block. The base of the block corresponds to the class being represented, so that it is usual to draw histograms only for continuous variables. .

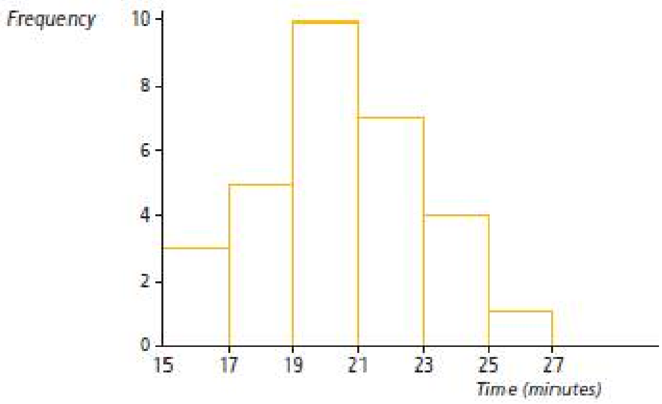

An ogive is a graph of a cumulative frequency distribution. The x -axis is the variable being measured and the y -axis is the corresponding cumulative frequency, the x – and y -values being plotted in exactly the same way as we discussed earlier. With a discrete variable, intermediate x -values have no meaning in reality (recall 1.6 invoices) and so the ogive would consist of a series of discrete points. It is usual therefore not to draw it. With a continuousvariable, the intermediate values do have a meaning, and so it makes sense to join the plotted points.

This can be done with a series of straight lines, which is tantamount to assuming that the values are evenly spread throughout their classes. The ogive of the continuous cumulative frequency distribution encountered earlier is shown in Figure below.

Before leaving these graphs, we look at one simple example to show how cumulative frequency distributions and ogives can be used in practice.

Example

Plot the histograms and ogives to represent the data in Example.

Solution

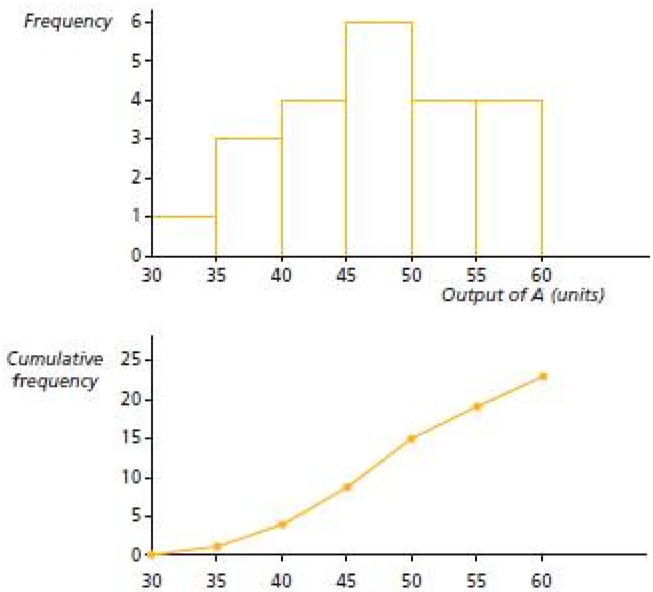

Histogram and ogive

Example The management of the company discussed in Example 3.5.3 wishes to reduce the target time for the operation to 22 minutes. Assuming the distribution of times remains unaltered, how often will this target be met?

Solution

First of all, it is not possible to answer this as a straight reading from the cumulative frequency distribution, as 22minutes does not correspond to a value in the table. If we look at the ogive in Figure 3.10 , however, we can estimate how many of the 30 occasions took less than 22 minutes, by reading off the graph, as shown. Thus, we estimate that the target will be met on 21.5 out of every 30 occasions: that is, 72 per cent of the time.

Example

Use the ogive from Example 3.8.1 to estimate:

(a) the percentage of days when output of A is less than 47 units;

(b) the percentage of days when output of A is more than 52 units.

Solution

(i) From the graph, the cumulative frequency corresponding to 47 units is 10.5, so the required percentage is 100 x10.5/22 = 48 percent (approximately).

(ii) From the graph, the cumulative frequency corresponding to 52 units is 15.5, so the percentage less than 52is 100 = 15.5/22 = 70 per cent and the percentage greater than 52 is 30 per cent (approximately).

We now give two examples to demonstrate problems that can arise when drawing histograms and ogives. We shall also at this point introduce a new diagram called a frequency polygon. For ungrouped, discrete data this consists of a simple graph of frequencies on values. For grouped data it is a graph of frequency density on interval mid-points and it may be obtained by joining the mid-points of the tops of the bars of the histogram. The first example deals with the important topic of unequal class intervals.

The resulting correct version of the histogram is shown in Figure 3.12(b) . Formally, it is the area of the block that is proportional to the frequency. It will be noted that the areas of the blocks are now in the correct proportion and that the vertical axis of the graph can no longer be labelled 'frequency', but is now 'frequency density'.

Before leaving this example, it is worth pointing out that the ogive of this distribution (Figure 3.13) would present no extra problems. As this consists only of plotting the upper limit of each class against cumulative frequency, the unequal class intervals do not affect matters.

Example

Plot the histogram for the following distribution:

Time taken to complete repeated task (minutes) Frequency

Time taken to complete repeated task (minutes)	Frequency
10–under 20	63
20–under 30	52
30–under 40	46
40–under 60	60
60–under 80	48
80–under 120	40

Solution

Time taken (minutes)	Class width	Frequency	Frequency density
10–20	10	63	63
20–30	10	52	52
30–40	10	46	46
40–60	20	60	60/2 = 30
60–80	20	48	48/2 = 24
80–120	40	40	40/4 = 10

Solution

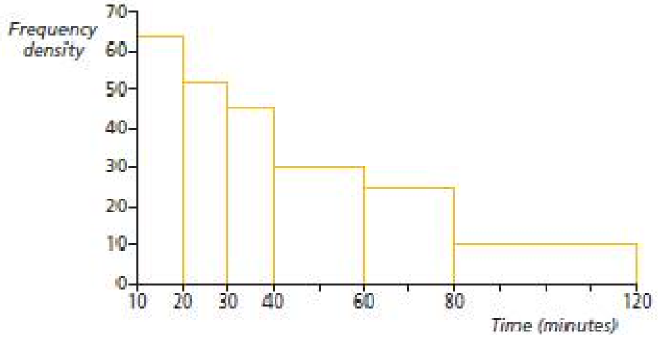

We have taken the standard class width to be ten minutes. For the two classes whose widths are twice the standard, we have divided frequency by two to get the frequency density. We have divided by four for the final class, whose width is four times the standard. Figure 3.14 shows the histogram.

Annual salary (£)	Number of graduate entrants
under 10,000	64
10,000–under 12,000	131
12,000–under 14,000	97
14,000–under 15,000	40
15,000 and over	11

Example

The compiler of the careers guide also receives, from a different source, information on the graduate salaries in another profession:

Annual salary (£)	Number of graduate entrants
under 10,000	64
10,000–under 12,000	131
12,000–under 14,000	97
14,000–under 15,000	40
15,000 and over	11

What problems would the compiler have when drawing the histogram and the ogive of this distribution?

Solution

We have seen how to deal with the unequal class intervals, but here we have the extra problem of open-ended classes. When drawing the histogram, we can either omit the first and last class or estimate 'closing' values for these two classes. The former would leave the histogram looking rather sparse, and, indeed, it is often necessary to close the classes so as to make certain calculations. It might therefore be advisable to estimate values such as

9,000 under 10,000 or 15,000 under 16,000

to draw the histogram. In the case of the ogive (Figure 3.15b), only the upper limit of each class is needed, and so it would be necessary to close only the last class in order to draw the whole ogive. It would also be possible to draw part of it from just the first four classes, omitting the last.

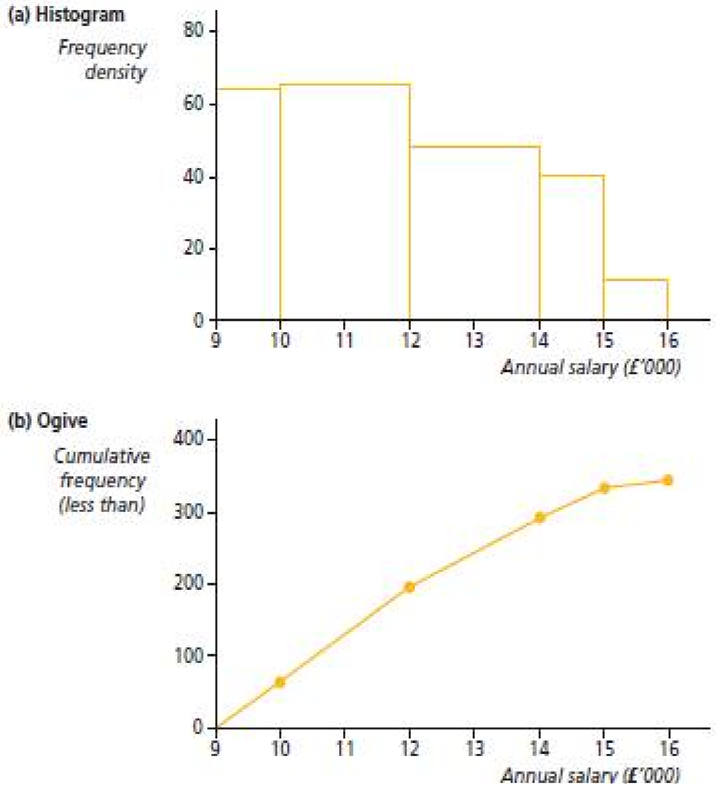

Pie charts

There are a number of other, more general, charts and graphs commonly used to represent business data. In this section we look at one of the most basic: pie charts. Pie charts are a very easily understood way of depicting the percentage or proportional breakdown of a total into various categories. They are so called because the total is represented by a circle, with each component shown as a sector with area proportional to percentage. Overall, the chart looks rather like a 'pie' with 'slices' in it. Sometimes two pie charts are used to compare two totals, along with the manner in which they are broken down. In such cases the areas of the pies, in other words the squares of their radii, are proportional to the total frequencies.

A company trades in five distinct geographical markets. In the last financial year, its turnover was:

	£m
UK	59.3
EU, outside UK	61.6
Europe, outside EU	10.3
North America	15.8
Australasia	9.9
Total	156.9

Display these turnover figures as a pie chart.

The first step is to calculate the percentage of the total turnover for each region:

	%
UK: (59.3/156.9) =	37.8
EU	39.3
Europe	6.6
North America	10.1
Australasia	6.3

Second, in order to make each 'slice' of the 'pie' proportional in area to these percentages, the whole circle (360°) has to be apportioned into five sections:

	Angle,°
UK: 37.8% of 360° =	136.1
EU	141.5
Europe	23.8
North America	36.4

Alternatively, the angles can be calculated directly as proportions of 360°, for example,

$360° \times (59.3/156.9) = 136.1°$
$360° \times (61.6/156.9) = 141.5°$, etc.

The resulting pie chart is shown in Figure 3.17

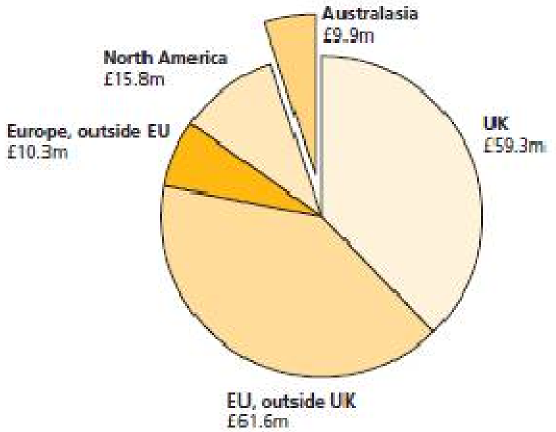

Bar charts

Bar charts represent actual data (as opposed to percentage breakdowns) in a way similar to earlier graphs. They appear similar to histograms, but with one essential difference: whereas distances against the vertical axis are measurements and represent numerical data, horizontal distances have no meaning. There is no horizontal axis or scale , there are only labels. Other than this proviso, the construction of bar charts is straightforward, as an example will illustrate.

Example

Represent the data of Example.

Solution

To draw this chart, it is simply a matter of drawing five vertical 'bars', with heights to represent the various turnover figures, and just labels in the horizontal direction .

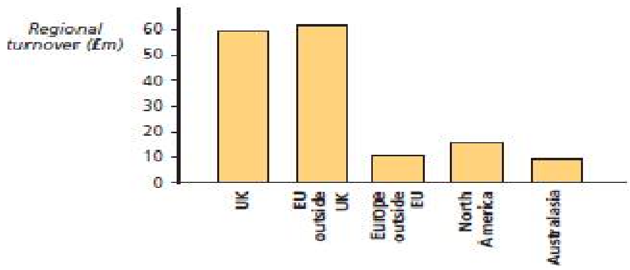

Figure Bar chart: breakdown of company turnover, last financial year

There are a number of variations on such a basic bar chart, used to display more data or more complex data.

An example will show just one.

Creating Bar charts using Excel

There are a wide range of different bar charts that can be created in Excel. Figure shows a basic bar chart that uses the same data as Example At present you cannot be asked to actually draw charts during a computer based assessment. Exam questions therefore take the form of labelling charts, calculating particular values, selecting a type of chart appropriate to particular data and drawing conclusions from charts.

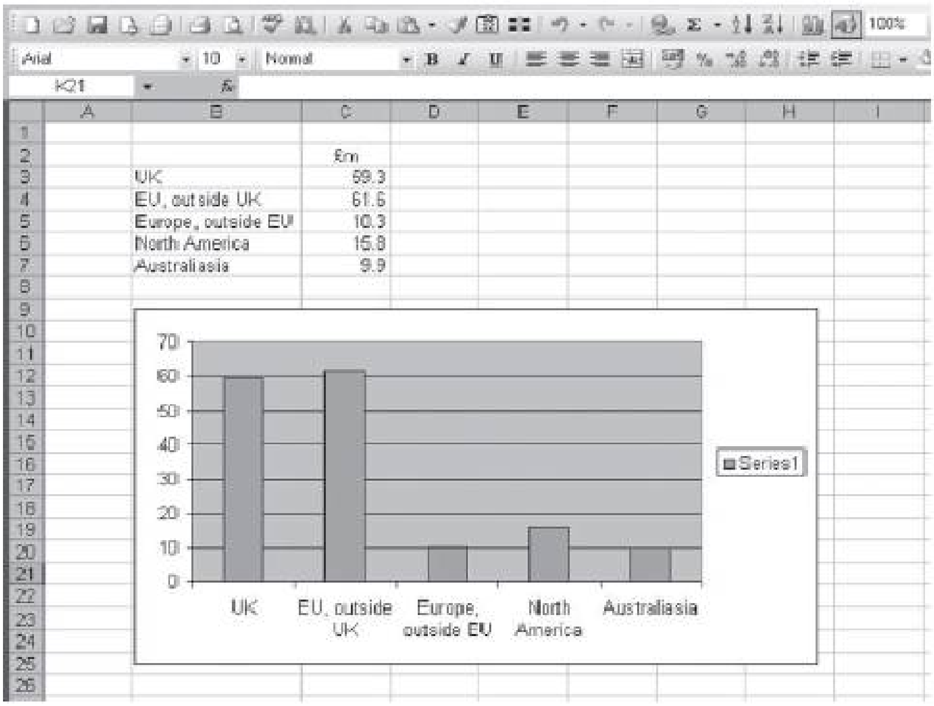

The chart in Figure is created by selecting the range b3:c7 clicking on the chart button on the toolbar and selecting the first sub-type option. This chart is representing a single set of data. Example introduced a second set of data which was represented using two different types of bar chart – a side-by-side and a stacked. Figures shows these charts produced in Excel.

In the case of the side-by-side bar chart the turnover for both companies for each country can easily be compared.

A slightly different version of the stacked bar is shown in Figure . This is referred to as a 100 per cent stacked bar and it compares the percentage each company contributes to the total across each country.

The stacked bar chart compares the contribution of each company to the total across the different countries.

Excel offers a variety of options when plotting data onto charts. Bar charts do not have to be represented by upright bars, but can be pyramids, cones or cylinders. Figure shows side-by-side bar in a pyramid shape. However, you should in each case, consider the most appropriate way is which to display date – often, the simpler methods of display are the most effective.

Furthermore bars can be positioned horizontally as opposed to vertically. This can be seen in Figure

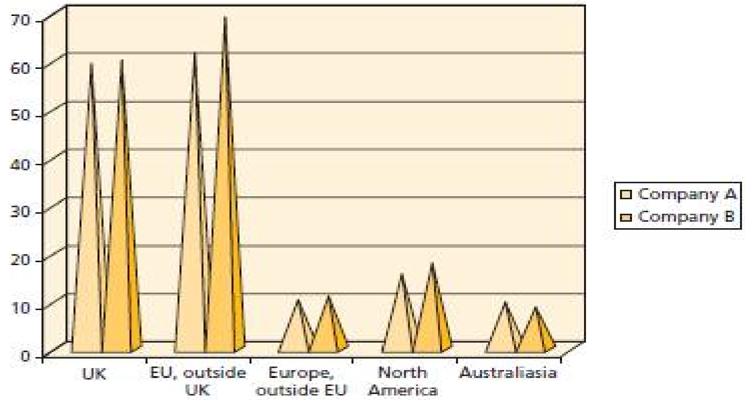

Figure a Bar chart using a pyramid shape

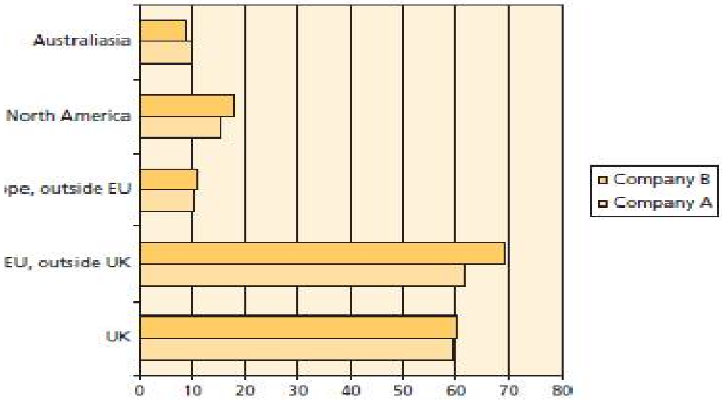

Figure Horizontal bar chart

Another way of looking at this data is to chart the turnover of the two companies showing the contribution from each country (as shown previously in Figure. This is achieved in Excel by selecting the data as before and selecting the stacked-bar chart option from the chart sub-options. However after clicking next, change the series setting to by rows instead of by columns. The results are shown in Figure.

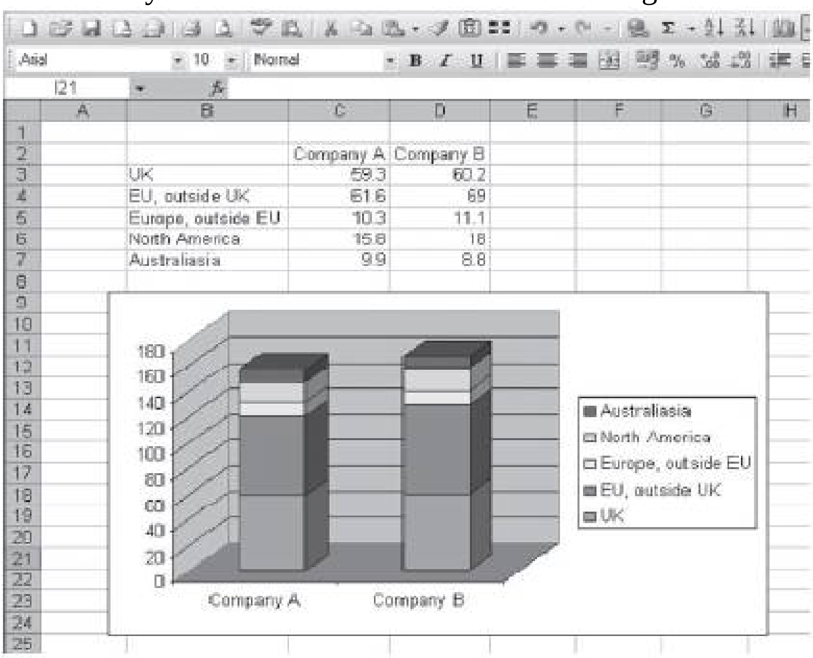

The examples shown so far in this section have been replicating the charts produced in an earlier section of this chapter.

Tabulation

Sometimes you will need to display actual figures rather than an illustration of them, and in such cases (especially if the data is to be presented to a numerate audience) a table may be preferable to a chart.

Example

In country A there were 22,618,462 dwellings. Of these 9,875,380 were owner-occupied, 6,765,097 were council rentals, 3,476,907 were private rentals and the remainder were held under a variety of tenures. In country B there were 1,846,093 in total and the numbers in the above categories were 569,043, 903,528 and 278,901, respectively. Display this information in a table.

Solution

The first step is to realise that the figures are 'far too accurate'. The table will be too cluttered if we retain this accuracy, so we will round each to the nearest thousand. We shall also need to calculate numbers in the 'others' category.

The next step is to decide upon the basic shape of the table. It shows four categories of housing in two countries, so the basic shape of the table is one column per country and one row per category.

Type of dwelling	*Country A*		*Country B*	
	Number ('000)	*%*	*Number ('000)*	*%*
Owner-occupied	9,875	44	569	31
Council rental	6,765	30	904	49
Private rental	3,477	15	279	15
Others	2,501	11	94	5
Total	22,618	100	1,846	100

Were you to draw the table in rough at this point (always a good idea) you would see that another row is needed for the totals. Furthermore, you would probably realise that country A has so many more dwellings than country B that they are not really comparable. This is dealt with by turning the frequencies into percentages, which will require two more columns. Figure shows the finished table. By using percentages we have been able to pinpoint the way in which patterns of owner-occupation and council rental are almost directly opposite in the two countries.

The basic rules of tabulation are:

- the aim is to present the data in an orderly fashion so that patterns can be seen;
- the table, with its title and headings, should be self-explanatory;
- It should be as simple as possible;
- combine small, unimportant categories if this simplifies the table;
- eliminate unnecessary detail by rounding;
- Find totals, subtotals and percentages where appropriate;
- think about who the table is aimed at – this may decide what other figures you calculate and how you position them;
- figures likely to be compared should be adjacent or at least close to one another;
- state units;

- State the source of the data if known.

Pareto analysis – The 80-20 rule

Pareto analysis was proposed by an Italian economist Vilfredo Pareto to describe how a relatively small part of a population may be so important. Pareto initially referred to wealth among individuals. He pointed out that a small number of individuals own a large portion of the wealth of any society. This idea is sometimes expressed more simply as the Pareto principle or the 80-20 rule.

In a business context the 80-20 rule states that a small number of clients will be responsible for a large proportion of the turnover, or a small number of inventory items will be responsible for a large amount of sales, or a small number of staff will present a disproportionate level of challenges to the management. When the term '80-20 rule' is used in this it does not always mean that it will actually be 20 per cent of the clients that produce 80 per cent of the profit as sometime it may be a smaller percentage like 10 which will produce the over whelming share of the profit.

In business the 80-20 rule essentially says that one should identify the really important elements of the business and focus the majority of one's time and effort on these elements.

It is possible to express the 80-20 rule in terms of a mathematical function and draw a probability density function. However, below is described a more practical approach to the 80-20 rule using simple commands in Excel. The managing director wants to know on which of these products the sales force concentrate should. There is an understanding that the best-selling products have the most upside potential. Therefore, it has been decided to apply 80-20 rule type thinking and highlight the best-performing products. There is a three-step process involved.

The first step is to sort the original data by the unit sales . This is achieved by selecting the range c 6: e 27 and selecting data sort . In the top box select unit sales , check the Descending box and make sure that Header Row is checked and click ok to perform the sort. Figure 3.30 shows the sorted data by unit sales .

The second step is to calculate the percentage of the total that each product line represents. To do this the following formula is entered into cell f 7.

= E7 / E28

This formula can be copied to the range f 8: f27

Now the third step is to calculate the cumulative sales in column g . Into cell g 7 enter:

= F7

And into cell g 8 enter

= G7+F8

This formula can be copied into the range g 9: g 27. Figure 3.31 shows the completed 80-20 rule table.

From Figure 3.31 it may be seen that Products F, U, L, T, I, D and A are the best performers and it is on these products that most effort should be expended.

Using spreadsheets to produce histograms, ogives and pie charts

Excel is a useful tool to create professional-looking charts including histograms, ogives and pie charts.

Creating a histogram in Excel

In order to show how to create a histogram in Excel, the data from Example will be used. Figure 3.32 shows this data in the spreadsheet.

As previously explained in this chapter, a histogram is a graph of frequency distribution, where the x axis is the variable being measured and the y axis is the corresponding frequency. In order to calculate the frequency distribution of the time taken to complete 30 repetitions of a task the Excel frequency function is used. The format of the function is = frequency(datarange,bin), where Bin refers to the required range of to be used for the x axis – in this example, time in minutes. Figure 3.33 shows the calculated frequency table in Excel.

Before entering the frequency formula into the spreadsheet first select the range j 3: j 9 and then enter:

= frequency (b 3: g 7, i3: i 9) and hold down the ctrl key and the shift key whilst pressing enter .

The formula is entered into all the cells in the range – this is referred to as an array function in Excel.

We now have the data that we want to plot onto a histogram, so select the range j 2: j 9 and click on the chart icon. Select Bar chart and take the first option of side-by-side bars. Click next to see the current chart. The chart so far is shown in Figure 3.34.

Figure First step in creating the histogram

The next step is to take the numbers in the range i 3: i 9 and use them as the labels for the x axis on the chart. To do this, click on the series tab and then click in the box next to the prompt 'Category x -axis labels' . Now select the range i 3: i 9 and then click finish. The chart is drawn as shown in Figure.

Histograms are generally shown with the side-by-side bars touching and using the chart formatting options in Excel this can be achieved here. Right click on one of the bars on the chart and select format data series . Then select the options tab. Set the gap width to zero and click OK. To finish off the chart it is helpful to add titles to the x and y axes. Right click on the white area surrounding the chart and select chart options, and then select the titles tab.

Readings

Company reports contain a mass of statistical information collected by the company, often in the form of graphs, bar charts and pie charts. A table or figures may not provide a very clear or rapid impression of company results, while graphs and diagrams can make an immediate impact and it may become obvious from them why particular decisions have been made. For example, the following graph provides a clear idea of the future of the company sales manager:

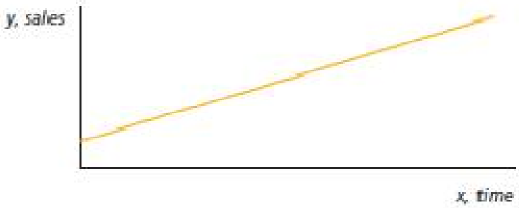

But pictures and diagrams are not always as illuminating as this one – particularly if the person presenting the information has a reason for preferring obscurity.

The gee-whiz graph

There is terror in numbers. Humpty Dumpty's confidence in telling Alice that he was master of the words he used would not be extended by many people to numbers. Perhaps we suffer from a trauma induced by early experiences with maths. Whatever the cause, it creates a real problem for the writer who yearns to be read, the advertising man who expects his copy to sell goods, the publisher who wants his books or magazines to be popular. When numbers in tabular form are taboo and words will not do the work well, as is often the case, there is one answer left: draw a picture. About the simplest kind of statistical picture, or graph, is the line variety. It is very useful for showing trends, something practically everybody is interested in showing or knowing about or spotting or deploring or forecasting. We'll let our graph show how national income increased 10 per cent in a year.

Begin with paper ruled into squares. Name the months along the bottom. Indicate billions of dollars up the side. Plot your points and draw your line, and your graph will look like this:

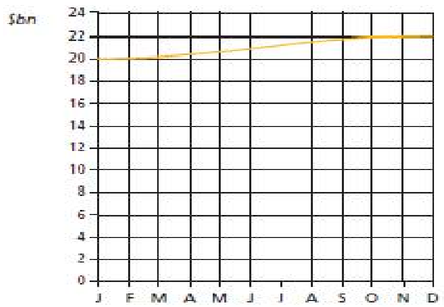

Now that's clear enough. It shows what happened during the year and it shows it month by month. He who runs may see and understand, because the whole graph is in proportion and there is a zero line at the bottom for comparison. Your 10 per cent looks like 10 per cent – an upward trend that is substantial but perhaps not overwhelming. That is very well if all you want to do is convey information. But suppose you wish to win an argument, shock a reader, move him into action, sell him something. For that, this chart lacks schmaltz. Chop off the bottom.

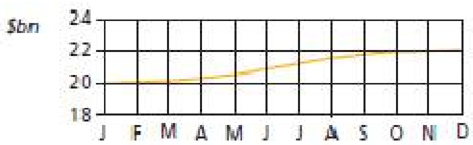

Now that's more like it. (You've saved paper too, something to point out if any carping fellow objects to your misleading graphics.) The figures are the same and so is the curve. It is the same graph. Nothing has been falsified – except the impression that it gives. But what the hasty reader sees now is a national-income line that has climbed half-way up the paper in twelve months, all because most of the chart isn't there anymore. Like the missing parts of speech in sentences that you met in grammar classes, it is 'understood'. Of course, the eye doesn't 'understand' what isn't there, and a small rise has become, visually, a big one.

Now that you have practised to deceive, why stop with truncating? You have a further trick available that's worth a dozen of that. It will make your modest rise of 10 per cent look livelier than one hundred per cent is entitled to look. Simply change the proportion between the ordinate and the abscissa. There's no rule against it, and it does give your graph a prettier shape. All you have to do is let each mark up the side stand for only onetenth as many dollars as before.

That is impressive, isn't it? Anyone looking at it can just feel prosperity throbbing in the arteries of the country. It is a subtler equivalent of editing 'National income rose 10 per cent' into '. . . climbed a whopping 10 per cent' . It is vastly more effective, however, because it contains no adjectives or adverbs to spoil the illusion of objectivity. There's nothing anyone can pin on you.

And you're in good, or at least respectable, company. A news magazine has used this method to show the stock market hitting a new high, the graph being so truncated as to make the climb look far more dizzying than it was. A Columbia Gas System advertisement once reproduced a chart 'from our new Annual Report' . If you read the little numbers and analysed them you found that during a ten-year period living costs went up about 60 per cent and the cost of gas dropped 4 per cent. This is a favourable picture, but it apparently was not favourable enough for Columbia Gas. They chopped off their chart at 90 per cent (with no gap or other indication to warn you) so that this was what your eye told you: living costs have more than tripled, and gas has gone down one-third!

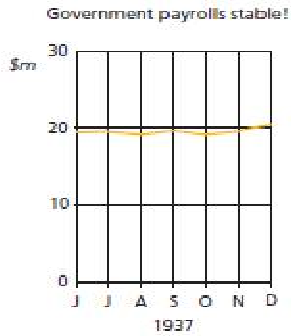

Steel companies have used similarly misleading graphic methods in attempts to line up public opinion against wage increases. Yet the method is far from new, and its impropriety was shown up long ago – not just in technical publications for statisticians either. An editorial writer in Dun's Review back in 1938 reproduced a chart from an advertisement advocating advertising in Washington, D.C., the argument being nicely expressed in the headline over the chart: **Government payrolls up** ! The line in the graph went along with the exclamation point even though the figures behind it did not. What they showed was an increase from near the bottom of the graph clear to the top, making an increase of under 4 per cent look like more than 400. The magazine gave its own graphic version of the same figures alongside – an honest red line that rose just 4 per cent, under this caption:

Government payrolls stable!

Postscript

Remember that if you are presenting information your objective will be to inform. Visual images can have an impact out of all proportion to the supporting detailed numbers. It is not sufficient to get the numbers right while your visual representations are slapdash – the end result may be to mislead or confuse your audience. As a user of statistical information, remember that the pictures you see may not give the same impression as a detailed numerical analysis would reveal. It is often essential to get behind the graphics and delve into the raw statistics.

REVIEW QUESTIONS

1. The managers of a sales department have recorded the number of successful sales made by their 50 telesales persons for one week, and the raw scores are reproduced below:

20	10	17	22	35	43	29	34	12	24
24	32	34	13	40	22	34	21	39	12
10	49	32	33	29	26	33	34	34	22
24	17	18	34	37	32	17	36	32	43
12	27	43	32	35	26	38	32	20	21

Sales persons who achieve fewer than twenty sales are required to undertake further training.

a Complete the table displaying the data as a grouped frequency distribution. Sales Frequency

Sales	*Frequency*
10–14	6
15–19	4
20–24	A
25–29	B
C	14
D	6
40–44	E
45–49	F

b Suppose the frequency distribution for the data was as follows:

Sales Frequency Cumulative frequency

Sales	Frequency	Cumulative frequency
10 and under 15	6	G
15 and under 20	12	H
20 and under 25	17	I
25 and under 30	7	J

Find the cumulative frequencies G, H, I and J.

c. Suppose the frequency distribution and cumulative frequencies were as follows:

Sales	Frequency	Cumulative frequency
10 and under 15	7	7
15 and under 20	16	23
20 and under 25	13	36
25 and under 30	4	40

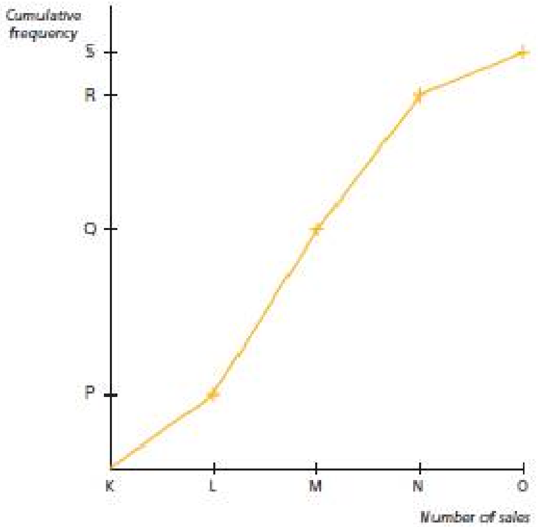

The chart is the cumulative frequency diagram illustrating this data. Find the values corresponding to the letters K–S.

2. The data below shows the number of people who passed through an airport terminal each day for a four-week period.

	A	B	C	D	E	F	G	H	I
1	No. of people passing through an airport terminal per day over a 4 week period								
2	All nos. in thousands								
3		Mon	Tues	Wed	Thur	Fri	Sat	Sun	
4	Week 1	122	177	103	161	202	115	147	
5	Week 2	178	133	109	125	106	139	164	
6	Week 3	124	191	152	145	140	169	153	
7	Week 4	159	154	113	144	167	128	149	

(A) Using this data produce a frequency distribution table to show the variation in the volume of traffi c through the terminal.

(B) Produce a histogram to graphically illustrate the frequency distribution.

(C) Calculate the cumulative frequencies.

(D) Produce an ogive chart to illustrate the cumulative frequency

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-4 DESCRIPTIVE STATISTICS

CONTENTS

INTRODUCTION

In Chapter 3 we saw how a set of raw data can be made more meaningful by forming it into a frequency distribution. Often it is advantageous to go further and to calculate values that represent or describe the whole data set; such values are called descriptive statistics. The most important are the various averages that aim to give a typical or representative value for the distribution. The other major group of descriptive statistics are the measures of spread, which tell us how variable the data are.

The arithmetic mean

Most people would understand an 'average' to be the value obtained by dividing the sum of the values in question by the number of values. This measure is the arithmetic mean, or, where there is no possibility of confusion, simply the mean. Further, if the data being considered is sample data, we refer to the sample mean to distinguish it from the mean of the population from which the sample is drawn. To understand the notation, consider the following example.

Example

A shopkeeper is about to put his shop up for sale. As part of the details of the business, he wishes to quote the average weekly takings. The takings in each of the last 6 weeks are:

£1,120 £990 £1,040 £1,030 £1,105 £1,015

Determine the mean weekly takings that the shopkeeper could quote.

Solution

If the weekly takings are denoted by the variable x, then the sample mean value of x is written as $\bar{x}$, pronounced 'x-bar'. Thus:

$$\bar{x} = \frac{\text{Sum of the values of } x}{\text{Number of values of } x}$$

$$\text{or } \bar{x} = \frac{\Sigma x}{n}$$

where Σ, a Greek capital letter 'sigma', is the mathematical symbol for 'add up', and n is the number of values of x. In this example:

$$\bar{x} = \frac{1{,}120 + 990 + 1{,}040 + 1{,}030 + 1{,}105 + 1{,}015}{6} = \frac{6{,}300}{6} = 1050.$$

The shopkeeper could therefore quote a sample mean weekly takings figure of £1,050.

As we can see, this formula is very easy to apply and, as indicated above, merely reflects the arithmetical procedures most people would recognise as the determination of an average. It will, however, need some modification before it can be used to determine the mean from a frequency distribution, a form in which many data sets appear.

Note how, in general, each x- value is multiplied by its corresponding frequency, f , and the products are then summed. That is, we evaluate the product fx for each x- value and then add all the values o f fx. As we are denoting addition by '=' , this sum can be written =fx . The formula for the sample mean from a frequency distribution is thus:

$$\bar{x} = \frac{\Sigma fx}{\Sigma f}$$

The denominator of this expression, =f , is simply the sum of the frequencies, which is, of course, the same as n in the earlier expression for x.

The median

So far we have dealt with the most commonly used average, the mean. We now consider another widely used average, the median. In Example 4.2.4, we computed a mean weekly wage of £ 190.60 which the personnel manager could quote in the wage negotiations. An impartial commentator could argue (and the manager might agree) that this is a rather high figure for a supposedly representative average. As 98 out of the sampled 170 (i.e. 58 per cent) actually earn less than £ 190 per week, it may well be that in excess of 60 per cent of the workforce earn less than the 'average' of £ 190.60 per week. If we look at this wage distribution, shown in Figure 4.1 , it is easy to see the cause of this phenomenon. The two highest frequencies occur at the lowest wage classes and then the frequencies decrease slowly as the wages increase. The relatively small number of large wages has caused the mean value to be so large. Distributions of this type are said to have a long tail or to be skewed. It is a criticism of the mean as an average that very skewed distributions can have mean values that appear unrepresentative, in that they are higher or lower than a great deal of the distribution. To address this problem, we introduce another measure of average, the median. This is defined as the middle of a set of values, when arranged in ascending (or descending) order. This overcomes the above problem, since the median has half the distribution above it, and half below. We leave the wage distribution for now, and look at a simpler example.

Example

Shop A's weekly takings are given by the following sample over six weeks. The sample has an arithmetic mean of £1,050.

£1,120 £990 £1,040 £1,030 £1,105 £1,015

A prospective purchaser of the business notices that the mean is higher than the takings in four of the 6 weeks. Calculate the median for him.

Solution

First of all, we arrange the takings figures in ascending order:

£990 £1,015 £1,030 £1,040 £1,105 £1,120

The question now is: what is the middle number of a list of six? With a little thought, you can see that there are two 'middle' values, the third and fourth. The median is thus taken to be the mean of these two values.

$$\text{Median} = \frac{(1{,}030 + 1{,}040)}{2} = 1{,}035$$

Hence, the median weekly takings figure that the prospective purchaser could quote is £1,035.

After this example, it is clear that, in the case of an odd number of values, the determination of the median is even easier, as there is a clear single middle item in an odd number of values. In general, if there are n observations, the position of the median is given by (n = 1) /2. With six observations, this gives 7/2 = 3.5, which is the position halfway between the third and fourth observations. In the case of frequency distributions, the determination of the median is not as straightforward, but can be illustrated by returning to the earlier wage distribution.

Example

(a) Calculate the median of the following data:

25 52 18 43 27

(b) Calculate the median of the data on staff absences (hint: use cumulative frequencies).

No. of employees absent	*No. of days (f)*
2	2
3	4
4	3
5	4
6	3
7	3
8	3

Solution

(a) First write the data in order of magnitude:

18 25 27 43 52

The median is in the third position [check: (5 + 1)/2 = 3] and is therefore 27.

(b) Find cumulative frequencies:

No. of employees absent	*No. of days (f)*	*Cumulative frequency*
2	2	2
3	4	6
4	3	9
5	4	13
6	3	16
7	3	19
8	3	22

There are 22 observations, so the position of the median is given by (22 + 1)/2 = 11.5, that is, the median is midway between the eleventh and twelfth observations. From the cumulative frequencies it is clear that both the eleventh and twelfth observations have value 5, so the median is 5.

The mode

The mode or modal value of a data set is that value that occurs most often, and it is the remaining most widely used average. The determination of this value, when you have raw data to deal with, consists simply of a counting process to find the most frequently occurring value, and so we do not dwell on this case here, but move on to look at frequency distributions.

Example

Find the mode for the following distributions:

(a)

Complaints per week	*No. of weeks*
0	5
1	12
2	7
3	2
4	1

(b)

Weekly wage (£)	No. of employees
180–under 185	41
185–190	57
190–195	27
195–200	23
200–205	15
205–210	7

Solution

(a) The mode is the value with the highest frequency, so here the mode is one complaint per week.
(b) The frequency distribution in the second case shows that the *modal class* (that one with the highest frequency) is £185 to under £190. Figure 4.4 shows a way of finding a single value to represent the mode: basically, the construction shown weights the mode towards the second most frequent class neighbouring the modal class, the one to the left in this case. The figure shows that *the modal weekly wage is approximately £186.60*.

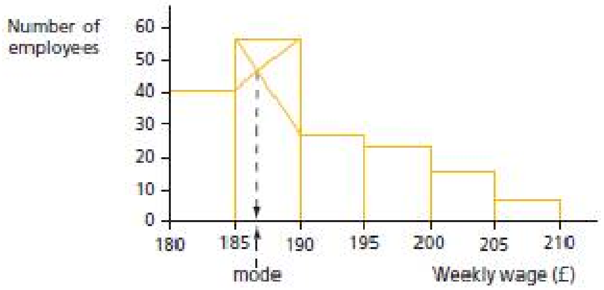

A comparison of the three averages

As our earlier discussion of the weekly wage distribution indicates, it is often just as important to use a measure of average appropriate to the situation as it is to evaluate the measure accurately. In this section we further discuss the relative merits and demerits of the three averages encountered in this chapter.

It is arguable that the mode is the least useful and important of the three. There are many distributions that have no definite single 'peak' in their histograms, and so it is difficult to attribute any sensible meaning to a modal value in such cases. Further, the mode is often unrepresentative of the whole data set: it may occur at one extremity of a skewed distribution or, at the very least, it takes no account of a high proportion of the data, only representing the most common value. All in all, it is fair to say that, in the vast majority of cases, the mode renders only a general description of one feature of a distribution, and is a relatively unimportant average when compared with the median or the mean.

The mean, on the other hand, has a number of features that usually make it the most appropriate and representative measure. First of all, it has the great advantage of being what most people recognise as 'the average' . It is therefore most easily communicated to non-specialists. Second, the mean is the only one of the three that takes account of all the data: like the mode, the median ignores some of a distribution by concentrating only on the middle part. The mean is thus arguably the most representative of all of a distribution. Finally, the mean is the measure that is most useful for further statistical analysis. Having said that, there are some circumstances in which one might consider the median to be more appropriate than the mean. We have already encountered one important such occasion in the skewed wage distribution of Example 4.3.3. Data is said to be positively skewed when the tail lies in the positive direction (i.e. to the right) and negatively skewed when the tail lies to the left. The mean is always pulled towards the tail, the mode towards the opposite steep slope, and the median lies between the two. So, for positively skewed data such as the wages example (and wage distributions generally) the mean will tend to overstate the average, and the mode will underestimate it.

For negatively skewed data the opposite is usually true, with the mean being too small and the mode too large.

In general, the median can often be argued to be the most representative average of a highly skewed distribution.

Another instance in which we may doubt the suitability of the mean is when we are dealing with discrete data. In Example 4.2.2 we saw an absenteeism distribution which has

Mean 3.07	absentees/day
Median 3	absentees/day
Mode 2	absentees/day

It is impossible to have 3.07 absentees, whereas the other two values are attainable. This is a common problem when dealing with means of discrete distributions, which sometimes leads to the median being used.

There is no hard and fast rule here, and each such case must be treated on its merits. Rounding the mean off to a 'possible' value of three absentees per day only represents a very small (around 2 per cent) change in its value. As this now agrees with the median, this common value can be accepted as the most appropriate measure. The decision would not be so easy with a discrete variable with a sample mean of 1.4 and median and mode both 3: rounding the mean to 1 would involve a large (almost 30 per cent) change, and so you would have to accept an unattainable mean value of 1.4 or the common median/ modal value of 3 as the 'average' .

A final category in which you might not use the mean is illustrated below.

An estate agent wishing to quote the average regional house price in his advertising brochures collates the following data on the houses he has helped to sell in the last 6 months:

House price (£'000)	*Number of houses*
Under 40	5
40–under 45	9
45–under 50	20
50–under 60	25
60–under 80	18
80–under 120	9
120 and over	6

Solution

We have already seen that, in such grouped frequency distributions, all three measures of average are of necessity only approximations. The two open-ended classes in this distribution, however, impose an extra source of inaccuracy in the evaluation of the mean: in order to obtain *x*-values for use in the formula for $\bar{x}$. We have to assume closing values for these classes. Plausible examples of these values might be:

30–under 40
120–under 150

Of course, the values chosen affect the *x*-values (the mid-points of the various classes) and thus the value of $\bar{x}$. The estimates of the median and modal values are unaffected by this choice as they both occur towards the centre of this distribution.

Again, there is no rule for resolving this dilemma. Most people would be willing to overlook the extra inaccuracy in the mean value, provided that plausible closing values for the open-ended classes are available (as is the case here). If, however, there was some doubt or debate over the closing values, then the median would arguably be the better measure.

Measures of spread

Having obtained an average value to represent a set of data, it is natural to question the extent to which the single value is representative of the whole set. Through a simple example we shall see that part of the answer to this lies in how 'spread out' the individual values are around the average. In particular, we shall study six measures of spread, or dispersion, as it is sometimes called:

- the range;
- the interquartile range;
- the quartile deviation;
- the mean absolute deviation;
- the standard deviation;
- the coefficient of variation.

The range

The measure of spread that is most associated with the mode is the range, since both statistics are relatively quick and easy to obtain, so they are well suited to initial exploration of the data. As we shall see later, neither of them are very useful statistics in most other circumstances.

The range is defined as the highest value minus the lowest value – but this can be misleading. Where the data is arranged in classes:

$$\text{Range} = \text{Upper most interval limit} - \text{Lowest interval limit}$$

Where the data is not grouped, the range is best viewed as the number of values from the very bottom to the very top and is given by:

$$\text{Range} = \text{Highest value} - \text{Lowest value} + 1$$

These apparently different definitions amount in practice to the same thing. If we regard the highest value as being at the centre of an interval of unit width, then the uppermost interval limit is given by the highest value plus 0.5. Similarly, the lowest interval limit will be given by the lowest value minus 0.5. Consequently, the value of the range is the same whichever method is used. The following example will illustrate the calculation of the range, and will demonstrate why such a measure may be needed.

Example

A recently retired couple are considering investing their pension lump sums in the purchase of a small shop. Two suitably sited premises, A and B, are discovered. The average weekly takings of the two shops are quoted as £1,050 and £1,080 for A and B, respectively. Upon further investigation, the investors discover that the averages quoted come from the following recent weekly takings figures:

Shop A:	£1,120	£990	£1,040	£1,030	£1,105	£1,015
Shop B:	£1,090	£505	£915	£1,005	£2,115	£850

Advise the couple.

Solution

You can easily check that the 'averages' quoted are, in fact, the means of the two samples. Based on these two figures alone, it might seem sensible for the couple to prefer shop B to shop A, but a glance at the actual data casts doubt on this conclusion. It is clear that the values for shop B are far more spread out than those for shop A, thereby making the mean for shop B arguably less representative. This difference is illustrated well by the ranges of the two sets:

$$\text{Range of } A = \text{Highest} - \text{lowest} + 1 = 1{,}120 - 990 + 1 = £131$$

$$\text{Range of } B = 2{,}115 - 505 + 1 = £1{,}611$$

It can be seen that the much larger range in the latter case is almost entirely due to the single value '£2,115'. The retired couple would therefore be well advised to look at larger samples of weekly takings figures to see if this value is some sort of freak and whether shop B does indeed generate higher weekly takings on average.

The interquartile range; the quartile deviation

The measure of spread most associated with the median is the quartile deviation, which we shall now consider. To do this, we begin by defining:

- the third quartile (denoted Q_3) as the value which has 75 per cent of the data below it; and
- the first quartile (Q_1) as the value with 25 per cent of the data below it.

The required measure, the interquartile range, is then the range of the middle 50 per cent of the data, or $Q_3 = Q_1$.

It will be noted that, if we referred to the second quartile (Q_2), we should simply be dealing with the median.

Example

After receiving complaints from trade union representatives concerning the disparity between higher- and lower-paid workers in his company, the personnel manager of the company asks for information on the current wage structure. He is given the following data:

Basic weekly wage (£)	*Number of employees*
under 200	16
200–under 225	153
225–under 250	101
250–under 275	92
275–under 300	68
300 and over	50

The manager decides to calculate a statistical measure of the spread of these data. Perform this calculation.

Solution

If we had the raw data, it would be a relatively simple counting process to find the wage figures 25 per cent and 75 per cent the way along the list, and thus to find the range or the interquartile range. As the data is presented, however, the range is unsuitable, as there are 50 employees in the open-ended upper class, any one of which could seriously distort the measure. Also, the interquartile range cannot be found by a process of mere counting.

In fact, to determine the interquartile range, we adopt the same approach as we did for the median. First of all, we assume that the wage values are evenly spread throughout their classes, and draw the ogive. The necessary cumulative frequency distribution is:

Basic weekly wage (£) (less than)	*Cumulative frequency*
200	16
225	169
250	270
275	362
300	430

It will be noted that it is unnecessary to close the final class in order to draw the ogive, and so we do not do so. The ogive is shown in Figure 4.6. Alternatively, a sensible closing value, such as £325, could be selected and an extra point, with cumulative frequency 480, added to the ogive.

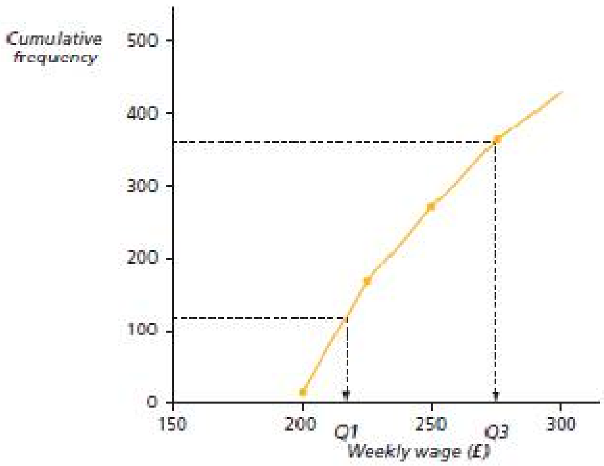

There is a very closely related measure here, the quartile deviation, which is half the interquartile range. In the above example, the quartile deviation is £ 28.50. In practice, the quartile deviation is used rather more than the interquartile range. If you rearrange $Q_3 = Q_1$ as $(Q_3 = M) = (M = Q_1)$ you will see that the two expressions in brackets give the distances from the quartiles to the median and then dividing by two gives the average distance from the quartiles to the median. So we can say that approximately 50 per cent of the observations lie within = one quartile deviation of the median.

Example

Using the data on the output of product Q (see Example 4.3.4), find the quartiles, the interquartile range and the quartile deviation from the ogive (Figure 4.3).

Solution

The total frequency = 22, so the cumulative frequency of Q_1 is $22/4 = 5.5$, and the cumulative frequency of Q_3 is $(3 \times 22/4) = 16.5$

From the ogive, $Q_1 = 362.5$kg and $Q_3 = 383.5$kg

Hence, the interquartile range $= 383.5 - 362 = 21.5$kg, and the quartile deviation $= 21.5 \div 2 = 10.75$kg.

Deciles

Just as quartiles divide a cumulative distribution into quarters, deciles divide a cumulative distribution into tenths. Thus:

The first decile has 10 per cent of values below it and 90 per cent above it, the second decile has 20 per cent of values below it and 80 per cent above it and so on. The use and evaluation of deciles can best be illustrated through an example.

Example

As a promotional example, a mail-order company has decided to give free gifts to its highest-spending customers. It has been suggested that the highest-spending 30 per cent get a gift, while the highest-spending 10 per cent get an additional special gift. The following distribution of a sample of spending patterns over the past year is available:

Amount spent (£)	*Number of customers spending this amount*
under 50	37
50–under 100	59
100–under 150	42
150–under 200	20
200–under 300	13
300 and over	9

To which customers should the gift and the additional special gift be given?

Solution

The cumulative frequency distribution (ignoring the last open-ended class) is:

Amount spent (less than, £)	*Cumulative frequency*
50	37
100	96
150	138
200	158
300	171

The ogive is shown in Figure 4.7.

Histogram Payment record of 100 customers

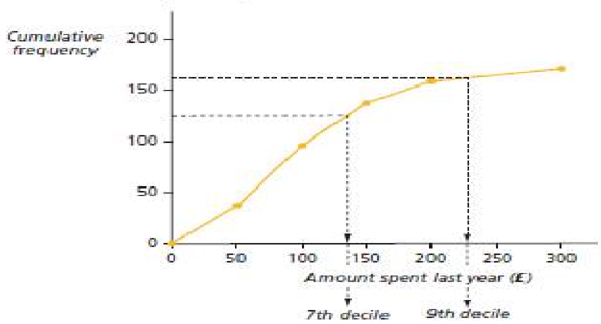

The ninth decile will correspond to a cumulative frequency of 162 (90 per cent of the total frequency, 180). From the ogive, this is: £230.

Similarly, the seventh decile corresponds to a cumulative frequency of 70 per cent of 180, that is 126. From the ogive, this is: £135.

Hence, in order to implement the suggestion, the company should give the free gift to those customers who have spent over £135 in the past year, and the additional free gift to those who have spent over £230.

The mean absolute deviation

If the mean is the average being used, then one very good way of measuring the amount of variability in the data is to calculate the extent to which the values differ from the mean. This is essentially the thinking behind the mean absolute deviation and the standard deviation (for which, see Section 4.11).

£1,120 £990 £1,040 £1,030 £1,105 £1,015

Example

Measure the spread of shop A's weekly takings (Example 4.7.1), given the following sample over 6 weeks. The sample has an arithmetic mean of £1,050.

Solution

A simple way of seeing how far a single value is from a (hopefully) representative average figure is to determine the difference between the two. In particular, if we are dealing with the mean, x , this difference is known as the deviation from the mean or, more simply, the deviation. It is clear that, for a widely spread data set, the deviations of the individual values in the set will be relatively large. Similarly, narrowly spread data sets will have relatively small deviation values. We can therefore base our measure on the values of the deviations from the mean. In this case:

$$\text{Deviation} = x - \bar{x}$$

In this case, the values of $(x - \bar{x})$ are:

£70, −£60, −£10, −£20, £55, −£35

The obvious approach might now be to take the mean of these deviations as our measure. Unfortunately, it can be shown that this always turns out to be zero and so the mean deviation will not distinguish one distribution from another. The basic reason for this result is that the negative deviations, when summed, exactly cancel out the positive ones: we must therefore remove this cancellation effect.

One way to remove negative values is simply to ignore the signs, that is, to use the absolute values. In this case, the absolute deviations are:

$|x - \bar{x}|$: £70, £60, £10, £20, £55, £35

The two vertical lines are the mathematical symbol for absolute values and are often referred to as 'modulus' , or 'mod' , of (x = x) in this case. The mean of this list is now a measure of the spread in the data. It is known as the mean absolute deviation. Hence the mean absolute deviation of weekly takings for shop A is:

$$\frac{70 + 60 + 10 + 20 + 55 + 35}{6} = £41.67$$

Thus, our first measure of the spread of shop A's weekly takings is £41.67. In your exam you cannot be asked to draw the ogive so you just have to know how to obtain the quartiles and percentiles from it. It is possible to calculate these statistics but this is not required in your syllabus and the formulae are not given.

The standard deviation

In the preceding section, we solved the problem of negative deviations cancelling out positive ones by using absolute values. There is another way of 'removing' negative signs, namely by squaring the figures. If we do that, then we get another, very important, measure of spread, the standard deviation.

Example

Evaluate the measure of the spread in shop A's weekly takings (Example 4.7.1), using this new approach.

Solution

Recall that we have the deviations:

$x - \bar{x}$: £70, −£60, −£10, −£20, £55, −£35

so, by squaring, we get:

$(x - \bar{x})^2$: 4,900, 3,600, 100, 400, 3,025, 1,225

The mean of these squared deviations is:

$$\frac{13{,}250}{6} = 2{,}208.3$$

This is a measure of spread whose units are the square of those of the original data, because we squared the deviations. We thus take the square root to get back to the original units (£). Our measure of spread is therefore:

$$\sqrt{2208.3} = £46.99$$

This is known as the *standard deviation*, denoted by 's'. Its square, the intermediate step before square-rooting, is called the *variance*, s^2.

The formula that has been implicitly used here is:

$$s = \sqrt{\frac{\Sigma(x - \bar{x})^2}{n}}$$

Applying the same series of steps to the data in a frequency distribution will give us the corresponding formula in this case:

- square the deviations: $(x - \bar{x})^2$
- find the mean of the $(x - \bar{x})^2$ values occurring with frequencies denoted by f.

$$\frac{\Sigma f(x - \bar{x})^2}{n} (= s^2)$$

- Take the square root:

$$\sqrt{\frac{\Sigma(x - \bar{x})^2}{\Sigma f}} (= s)$$

In practice, this formula can turn out to be very tedious to apply. It can be shown that the following, more easily applicable, formula is the same as the one above:

$$s = \sqrt{\frac{\Sigma fx^2}{\Sigma f} - \left(\frac{\Sigma fx}{\Sigma f}\right)^2}$$

The coefficient of variation

The coefficient of variation is a statistical measure of the dispersion of data points in a data series around the mean. It is calculated as follows:

$$\text{Coefficient of variation} = \frac{\text{Standard deviation}}{\text{Expected return}}$$

The coefficient of variation is the ratio of the standard deviation to the mean, and is useful when comparing the degree of variation from one data series to another, even if the means are quite different from each other.

In a financial setting, the coefficient of variation allows you to determine how much risk you are assuming in comparison to the amount of return you can expect from an investment. The lower the ratio of standard deviation to mean return, the better your risk return tradeoff. Note that if the expected return in the denominator of the calculation is negative or zero, the ratio will not make sense.

In Example, it was relatively easy to compare the spread in two sets of data by looking at the standard deviation figures alone, because the means of the two sets were so similar. Another example will show that it is not always so straightforward.

Example

Government statistics on the basic weekly wages of workers in two countries show the following. (All figures converted to sterling equivalent.)

Country V:	$\bar{x} = 120$	$s = £55$
Country W:	$\bar{x} = 90$	$s = £50$

Can we conclude that country V has a wider spread of basic weekly wages?

Solution

By simply looking at the two standard deviation figures, we might be tempted to answer 'yes'. In doing so, however, we should be ignoring the fact that the two mean values indicate that wages in country V are inherently higher, and so the deviations from the mean and thus the standard deviation will tend to be higher. To make a comparison of like with like we must use the *coefficient of variation:*

Solution

By simply looking at the two standard deviation figures, we might be tempted to answer 'yes'. In doing so, however, we should be ignoring the fact that the two mean values indicate that wages in country V are inherently higher, and so the deviations from the mean and thus the standard deviation will tend to be higher. To make a comparison of like with like we must use the coefficient of variation:

$$\text{Coefficient of variation} = \frac{s}{\bar{x}}$$

$$\text{Coefficient of variation of wages in country V} = \frac{55}{120} = 45.8\%$$

$$\text{Coefficient of variation of wages in country W} = \frac{50}{90} = 55.6\%$$

Hence we see that, in fact, it is country W that has the higher variability in basic weekly wages.

A comparison of the measures of spread

Like the mode, the range is little used except as a very quick initial view of the overall spread of the data. The problem is that it is totally dependent on the most extreme values in the distribution, which are the ones that are particularly liable to reflect errors or one-off situations. Furthermore, the range tells us nothing at all about how the data is spread between the extremes.

The standard deviation is undoubtedly the most important measure of spread. It has a formula that lends itself to algebraic manipulation, unlike the quartile deviation, and so, along with the mean, it is the basis of almost all advanced statistical theory. This is a pity because it does have some quite serious disadvantages. If data is skewed, the standard deviation will exaggerate the degree of spread because of the large squared deviations associated with extreme values. Similarly, if a distribution has open intervals at the ends, the choice of limits and hence of mid-points will have a marked effect on the standard deviation. The

quartile deviation, and to a lesser extent the interquartile range, is the best measure of spread to use if the data is skewed or has open intervals. In general, these measures would not be preferred to the standard deviation because they ignore much of the data and are little
Known
Finally, it is often the case that data is intended to be compared with other data, perhaps nationwide figures or previous year's figures, etc. In such circumstances, unless you have access to all the raw data, you are obliged to compare like with like, regardless perhaps of your own better judgement.
All the averages give a typical or expected value for the distribution. The mean is the total shared out equally, the median is the halfway value and the mode is the most common value.
All the measures of spread tell you how variable the data is. If the measures are relatively large, it means that the data is very variable. Both the mean deviation (directly) and the standard deviation (by a circuitous route of squaring and then square-rooting) find the average distance of the observations from the mean. In other words, they measure average variability about the mean.
The quartile deviation gives the average distance of the quartiles from the median. It is a measure of the variability of the central 50 per cent of observations, as is the interquartile range. Finally, the range itself measures the spread of the data from the very bottom to the very top.
Aside from understanding and being able to explain what the various statistics mean, there are other points of relevance to interpretation:
1. Can you rely on the data – was the sample large, representative and randomly taken?
2. Is the data skewed or open-ended? If so, are you using the right statistics?
3. Are you comparing like with like?
4. If you have to compare variability in two samples that have markedly different means, use the coefficient of variation rather than the standard deviation.
5. Finally, always remember to interpret statistics in their proper context. Give them units and do not simply interpret them in an abstract manner.
6. Think about which averages (measures of location) are used with which measures of spread, for example, the mean is used with the standard deviation (or variance).

Formulae definitions

- The mean, $\bar{x} = \Sigma x/n$ or $x = \Sigma fx/\Sigma f$ for frequency distributions.
- The median is the middle value when the data is arranged in ascending or descending order. It can be evaluated directly except in grouped frequency distributions, when it can be estimated from an ogive as the x - value corresponding to half the total frequency.
- The mode is the most commonly occurring value.
- The standard deviation,

$$s = \sqrt{\frac{\Sigma(x - \overline{x})^2}{n}}$$

or

$$s = \sqrt{\frac{\Sigma fx^2}{\Sigma f} - \left(\frac{\Sigma fx}{\Sigma f}\right)^2}$$

for frequency distributions.

- Coefficient of variation = $s/\overline{x}$
- Mean absolute deviation = $\Sigma|x - \overline{x}|/n$
- The interquartile range = $Q_3 - Q_1$
- The quartile deviation = $\frac{1}{2}(Q_3 - Q_1)$

  where Q_3 is the third quartile (which has 75 per cent of the values below it) and Q_1 is the first quartile (which has 25 per cent of the values below it).
- The values of the quartiles can be estimated from the ogive, in the case of a frequency distribution.
- Deciles divide a distribution into tenths, so that the first decile has 10 per cent of values below it, the second decile 20 per cent of values below it and so on. The values of deciles can be estimated from an ogive.
- Range = Highest interval limit – Lowest interval limit.

REVIEW QUESTIONS

1. A company is investigating the cost of absenteeism within its production department. Computer records revealed the following data:

Days absent last year	*Number of people*
0	94
1–5	203
6–10	105
11–20	68
21–30	15
31–40	10
41+	5
Total	500

A. Complete the following cumulative frequency distribution table:

No. of days absent	*Cumulative number of people*
0	94
1 and <6	297
6 and <11	402
A	470
21 and <31	B
31 and <41	495
41 and <51	C

B. The following chart is the ogive (cumulative 'less than ' frequency diagram) of this data. State the y -coordinates of the points labelled D, E and F.

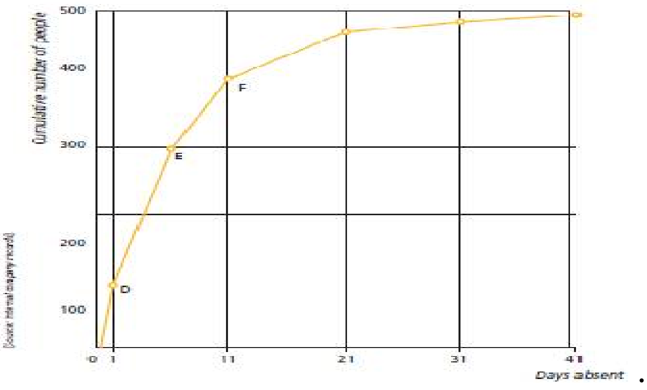

C. Mark estimates of the following statistics on the ogive, clearly showing the constructions you have used to obtain them. You need not attempt to estimate the values of the statistics but should clearly label any cumulative frequencies you use in estimating them.

(A) median

(B) upper quartile

(C) highest decile.

D. If the median were 5 days, the lower quartile 2 days and the upper quartile 10 days, calculate the quartile deviation.

E. If the median were 5 days, lower quartile 2 days and upper quartile 10 days, which of the following comments would be correct?

(A) Only a quarter of the staff had less than 2 days absence.

(B) The high value of the upper quartile shows that absenteeism has increased.

(C) Half the staff had less than 5 days absence.

(D) The low value of the lower quartile shows that absenteeism is not very variable from one person to the next.

(E) Half the staff had more than 12 days absence.

(F) Three-quarters of the staff had less than 8 days absence.

(G) A quarter of the staff had more than 10 days absence.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-5 INDEX NUMBERS

CONTENTS

INTRODUCTION

We conclude the study of averages by looking at a special category, index numbers, which measure how a group of related commercial quantities vary, usually over time. As we shall see, most well-known index numbers are averages, but they have the extra property that they relate the quantities being measured to a fixed point or base period.

Definitions

If a series of values relating to different times are all expressed as a percentage of the value for a particular time, they are called index numbers with that time as the base.

$$\text{Index number} = \frac{\text{Value in any given year}}{\text{Value in base year}} \times 100$$

We shall generally refer to years but monthly data might have a particular month as base and so forth, so strictly speaking we should say 'time point' rather than 'year' .

Example

Express the following data with 1995 as the base:

Year	1995	1996	1997	1998	1999
Value	46	52	62	69	74

Solution

We have to express each value as a percentage of the value for 1995. That means we must divide each by the 1995 value (i.e. by 46) and multiply by 100.

Year	1995	1996	1997	1998	1999
1995 = 100	100	113	135	150	161

A few points of note:

A few points of note:

- The base year did not have to be the first of the series. We could have chosen any year.

• Expressions such as '1995 = 100' tell us that the associated values are index numbers with base 1995. The index number for the base year (1995 in this case) will always be 100.
• We have rounded to the nearest whole number simply to avoid cluttering the text, while you get used to the idea of index numbers. In fact, they could be expressed to any degree of accuracy.

Interpretation of index numbers

An index of 113 tells us that there has been a 13 per cent increase since the base year. In Example 5.2.1 we can see that values increased by 13 per cent from 1995 to 1996, by 35 per cent over the 2 years from 1995 to 1997, by 50 per cent over 3 years and by 61 per cent over the 4 years 1995–99.

It is essential to realise that the percentage changes always refer back to the base year. It is not possible to derive the percentage increase from one year to the next by subtracting index numbers.

Example

Find the percentage increase from 1998 to 1999 for the data in Example 5.2.1.

Solution

$$\frac{161}{150} \times 100 = 107.3$$

Therefore, the percentage increase is 7.3 per cent.

If we were to subtract the index numbers we would get 161-150 = 11 points, which is clearly not the same as 11 per cent. Were the Financial Times 250 Share Index (which, at the time of writing, is at about 6,000) to rise by 50 points, that represents a rise of only $100 \times \frac{50}{6000} = 0.8$ per cent. This 50 point rise must not be confused with a 50 per cent rise! To derive the percentage increase from year A to year B, the easiest method is to index the year B figure with base year A and then subtract 100.

If values have declined, the index number will be less than 100, and when you interpret them by subtracting 100 the resulting negative tells you there has been a decline.

For example, an index of 94 means there has been a decline of 6 per cent (94-100= - 6) since the base year.

Finally, some index numbers become very large, like the FT100 index mentioned above. Its 6,250 level means that there has been a 6,150 per cent increase in share prices since the base year = not a very meaningful interpretation, even to those of us who are quite numerate. A much better interpretation is that share prices are now 62.5 times what they were in the base year.

Hence, there are two ways of interpreting an index number:
• subtracting 100 gives the percentage increase since the base year;
• dividing by 100 gives the ratio of current values to base-year values.

Year	1991	1992	1993	1994	1995	1996	1997
Profits (£m)	1.2	1.5	1.8	1.9	1.6	1.5	1.7

(a) Express the profits figures above as index numbers with:
(i) base 1991;
(ii) base 1994.
(b) Interpret both the index numbers for 1995.
(c) Find the percentage increase from 1996 to 1997.
(d) Interpret the index number 2,500 with 1989 = 100.

Solution

	Year	1991	1992	1993	1994	1995	1996	1997
(a)	(i) 1991 = 100	100	125	150	158	133	125	142
	(ii) 1994 = 100	63	79	95	100	84	79	89

(b) The index of 133 means there has been a 33 per cent increase in profits from 1991 to 1995. The index of 84 means that profits in 1995 are 16 per cent below their level in 1994.
(c) 100 × (1.7/1.5) = 113 so there has been a 13 per cent increase from 1996 to 1997.
(d) An index of 2,500 means that values now are 2,500/100 = 25 times their value in the base year, 1989.

Choice of base year

Although a particular base may be satisfactory for several years, it becomes less meaningful as time passes and eventually it is necessary to shift to a new base this is called rebasing. The only requirements of a suitable base year are, first, that it should be a fairly typical year. For example, if prices are being indexed then a year should be chosen in which prices were neither specially high nor specially low. Second, it should be sufficiently recent for comparisons with it to be meaningful. For example, it might be useful to know that production had changed by a certain percentage over the last year or two or perhaps over 10 years, but an index with base, say, 50 years ago would not be very relevant. It is also the case, as we shall see shortly, that many index numbers span a wide range of popular goods and, reflecting what people actually buy, they are very different now from 20 or 30 years ago. The base year has to move in order to keep up with the composition of index numbers to some extent.

Change of base year

If at all possible you should return to the original data and recalculate the index numbers with the new base year. However, if only the index numbers are available, they can be indexed as if they were the original data. The only problem is that sometimes rounding errors will build up. As the following example shows, changing the base using index numbers instead of raw data can give very good results.

Example

(a) Express the following as index numbers with base 1990.
(b) Using the original data, change the base to 1995.
(c) Using the index with base 1990, change the base to 1995 and compare your results with those in part (b).

Express your answers to one decimal place throughout.

Year	1990	1991	1992	1993	1994	1995	1996	1997	1998
Staff	8	9	9	12	20	22	24	25	27

Solution

Year	1990	1991	1992	1993	1994	1995	1996	1997	1998
Staff	8	9	9	12	20	22	24	25	27
(a) 1990 = 100	100.0	112.5	112.5	150.0	300.0	275.0	300.0	312.5	337.5
(b) 1995 = 100	36.4	40.9	40.9	54.5	90.9	100.0	109.1	113.6	122.7
(c) 1995 = 100	36.4	40.9	40.9	54.5	90.9	100.0	109.1	113.6	122.7

The index labelled (b) has been obtained from the original data, by dividing by 22 and multiplying by 100. The index labelled (c) has been obtained from the 1990 = 100 index numbers by dividing by 275 and multiplying by 100. As you can see, the results are identical when rounded to one decimal place in this case.

Combining series of index numbers

When a series of index numbers is subject to a change of base or perhaps a small change of composition, you will find in the series a year with two different index numbers and the change of base will be shown in the series. The technique involved in combining two series into a single one is called splicing the series together.

The price index below changed its base to 1983 after many years with base 1970. Recalculate it as a single series with base 1983. By how much have prices risen from 1981 to 1985?

Year	Price index
	(1970 = 100)
1980	263
1981	271
1982	277
1983	280
	(1983 = 100)
1984	104
1985	107

Solution

The index numbers from 1983 onwards already have 1983 = 100, so nothing need be done to them. What we have to do is to change the base of the original series, so it too is 1983. In this series the value for 1983 is 280, so we must divide the index numbers for 1980–82 by 280 and multiply by 100.

Year	Price index	Price index	
	(1970 = 100)	(1983 = 100)	
1980	263	94	= 100 × (263/280)
1981	271	97	= 100 × (271/280) etc.
1982	277	99	
1983	280	100	
	(1983 = 100)		
1984	104	104	
1985	107	107	

Now that we have a single series spanning both 1981 and 1985, we can compare the two:

$$100 \times \left(\frac{107}{97}\right) = 110$$

You may notice that we rounded to the nearest whole number in this example. This is because the original index numbers had plainly been rounded to the nearest whole number and at best we can hope that our results will be accurate to that same extent. You cannot acquire increased accuracy in the course of calculating.

Chain-base index numbers

So far we have dealt with index numbers that have the same base for several years. These are called fixed-base index numbers and, unless you are informed to the contrary, it is reasonable to assume that index numbers are of this type. However, it is often of more interest to know the annual increase. A chain-base index number (or simply a chain index) expresses each year's value as a percentage of the value for the previous year.

$$\text{Chain-base Index} = \frac{\text{This year's value}}{\text{Last year's value}} \times 100$$

Example

Average exam marks of 40 per cent, 55 per cent and 58 per cent with weights of 2, 2 and 1.

Solution

$$\text{Weighted average} = \frac{(2 \times 40) + (2 \times 55) + (1 \times 58)}{2 + 2 + 1} = \frac{248}{5} = 49.6\%$$

Weights are a measure of the importance that we allocate to each item. In the above example, the first two exams are rated as twice as important as the third one. In the arithmetic mean, values are weighted by the frequency with which they occur; in price indices, similar weighting systems operate.

We cannot calculate the chain index for 1995 because we do not have the 1994 figure ('n/a' means 'not available'). The interpretation is the same as for fixed-base index numbers, except that the percentage change is each time over the previous year. In this example, the results tell us that values rose by 13.0 per cent from 1995 to 1996, by 19.2 per cent from 1996 to 1997, by 11.3 per cent the next year and by 7.2 per cent from 1998 to 1999. Fixed-base index numbers can easily be changed into chain-base indices by treating them as if they were the original data. The only problem is to not pick up spurious accuracy. Try the next example.

Composite index numbers

In practice, most price indices cover a whole range of items and so there are two processes involved in the construction of the index number. One is indexing = comparing current values with those of the base year – and the other is averaging or somehow combining together the items under consideration.

We shall begin by using the method of combining individual price indices by means of a weighted average. In Chapter 4, we used the formula $\Sigma fx/\Sigma f$ for the average or arithmetic mean. The formula for a weighted average is the same as this but with weights, denoted by w, instead of frequencies. Hence:

$$\text{Weighted average} = \frac{\sum wx}{\sum w}$$

where x denotes the values being averaged and w denotes the weights.

Average exam marks of 40 per cent, 55 per cent and 58 per cent with weights of 2, 2 and 1.

Solution

Average exam marks of 40 per cent, 55 per cent and 58 per cent with weights of 2, 2 and 1.

Solution

$$\text{Weighted average} = \frac{(2 \times 40) + (2 \times 55) + (1 \times 58)}{2+2+1} = \frac{248}{5} = 49.6\%$$

Weights are a measure of the importance that we allocate to each item. In the above example, the first two exams are rated as twice as important as the third one. In the arithmetic mean, values are weighted by the frequency with which they occur; in price indices, similar weighting systems operate.

Relative price indices

The notation commonly used for the construction of index numbers is as follows: the subscripts '0' and '1' are used, respectively, for the base year and the year under consideration, usually called the current year. Hence, for any given item:

P_0 = price in base year $\qquad P_1$ = price in current year

Q_0 = quantity in base year $\qquad Q_1$ = quantity in current year

$V_0 = P_0Q_0$ = value in the base year $\qquad V_1 = P_1Q_1$ = value in current year

where value means the total expenditure on the item, and other sorts of weights are denoted by w as before.

For a given item, the price index = 100 = (P_1/P_0), but quite often we work with the ratio called the price relative = P_1/P_0 and leave the multiplication by 100 to the end of the calculation.

The usual formula for a relative price index is therefore:

Relative price index

$$\text{Relative price index} = \frac{\sum[w \times (P_1/P_0)]}{\sum w} \times 100$$

This formula is given in your exam. The weights could be base-year quantities (i.e. Q_0) or values (i.e. P_0Q_0), or current-year quantities or values (i.e. Q_1 or P_1Q_1), or they could simply be decided on some other basis such as the weighting of exam marks.

The index will be called base-weighted or current-weighted, depending on whether it uses base or current weights.

Example

A grocer wishes to index the prices of four different types of tea, with base year 1990 and current year 1995. The available information is as follows:

	1990		1995	
	Price (£)	Quantity (crates)	Price (£)	Quantity (crates)
Type	P_0	Q_0	P_1	Q_1
A	0.89	65	1.03	69
B	1.43	23	1.69	28
C	1.29	37	1.49	42
D	0.49	153	0.89	157

Calculate the base-weighted relative price index using as weights (a) quantities; and (b) values (i.e. revenue for each item).

Solution

	Price relative (Rel)	Base-year quantity (Q_0)	Base-year value (V_0)	$Rel \times Q_0$	$Rel \times V_0$
A	1.157	65	57.85	75.22	66.95
B	1.182	23	32.89	27.19	38.88
C	1.155	37	47.73	42.74	55.13
D	1.816	153	74.97	277.85	136.15
Total		278	213.44	423.00	297.11

Base-weighted relative price indices are:

$$\text{Weighted by quantity: } \frac{\sum(Rel \times Q_0)}{\sum Q_0} \times 100 = \frac{423}{278} \times 100 = 152.2$$

$$\text{Weighted by value: } \frac{\sum(Rel \times V_0)}{\sum V_0} \times 100 = \frac{297.11}{213.44} \times 100 = 139.2$$

The first index tells us that prices have risen on average by 52 per cent; the second that they have risen by 39 per cent. Why might this be so? The really big price rise is D’s 82 per cent. The size of the index will be very strongly influenced by the weight given to D. In the first case, the quantity 153 is bigger than all the other quantities put together. D gets more than half of the total weight and so the index strongly reflects D’s price rise and is very high. However, when we use value for weighting, D’s value is only about one-third of the total because its price is low, so the price index is rather smaller.

Now it is your turn to see what happens if we use current quantities and values as weights.

Example

Using the data of Example, calculate the current-weighted relative price index with weights given by

(a) quantities; and

(b) values.

Solution

Aggregative price indices

In the previous section we indexed first (except that we did not multiply by 100) and subsequently combined the indices together in a weighted average. In this section we shall construct index numbers the other way round. We shall first combine the prices to give the total cost of a notional shopping basket and only then shall we index the cost of the basket at current prices compared with the cost at base-year prices. These indices are called 'aggregative' because the first step is to aggregate, or combine together, all the items under consideration. The first step in calculating a relative price index is to calculate all the price relatives.

If we want to compare prices now with those in the base year, we must compare like with like – the composition of the two baskets must be identical in every respect apart from price. If you buy goods with prices given by P in quantities given by Q , the total cost of the shopping basket is σ PQ . The general form of an aggregative price index is therefore:

$$\frac{\sum_w P_1}{\sum_w P_0} \times 100$$

where the weights will be the quantities chosen.

The formula for aggregative price indices is not given in your exam. It is less likely to be required than the relatives method.

Example

Find the aggregative price index for the following data, using the weights given:

Item	P_0	P_1	w
A	35	36	2
B	23	27	3
C	15	21	5

Solution

Item	wP_0	wP_1
A	70	72
B	69	81
C	75	105
Total	214	258

$$\text{Aggregative price index} = 100 \times \frac{\sum wP_1}{\sum wP_0}$$

$$= 100 \times \frac{258}{214}$$

$$= 120.6$$

Choice of base weighting or current weighting

- In general, current weighting will seem better because it remains up to date.
- In particular, current-weighted indices reflect shifts away from goods subject to high price rises.
- Base-weighted indices do not do this and hence exaggerate inflation.
- However, current quantities can be very difficult to obtain – some considerable time may elapse after a year ends before a company knows what quantity it sold, whereas base quantities are known and remain steady for the lifetime of the index.
- So, current-weighted price indices are usually much more costly and time-consuming to calculate than are base-weighted ones.

- The stability of base weights means that the index for each year can be compared with that of every other year which, strictly speaking, a current-weighted index cannot.
- There can be no general guidance on the choice: it depends on the resources available and on the degree to which prices and quantities are changing. The only other consideration is that, as always, you must compare like with like. The retail price index (RPI), as we shall see, is a current-value-weighted relative index weighted by (almost) current values, and that method of construction should be used if at all possible if comparison with the RPI is a major function of the index being constructed.

Quantity indices

Although it is the most important and most frequently encountered, price is not the only financial factor measured by index numbers. Quantity indices constitute another category. They show how the amounts of certain goods and commodities vary over time or location. They are of importance when one is considering changes in sales figures, volumes of trade and so on.

When considering price indices, quantities emerged as the best weighting factor. Here, the converse is true: prices are considered the most appropriate weights when calculating quantity indices. Accordingly:

A relative quantity index will take the form:

$$\frac{\sum[w \times (Q_1/Q_0)]}{\sum w} \times 100$$

An aggregative quantity index will take the form:

$$\frac{\sum wQ_1}{\sum wQ_0} \times 100$$

where in both cases the weights could be prices, either base year or current, or values or some other measure of the importance of the items. P_0 , P_1 , Q_0 and Q_1 have the same meanings as earlier in this chapter.

The calculation of quantity indices and their application involve no new arithmetical techniques, as the following example illustrates.

Example

A company manufactures two products, A and B. The sales figures over the past 3 years have been as follows:

	A Sales '000	B Sales '000
1993	386	533
1994	397	542
1995	404	550
Weight	22	19

Using 1993 as a base, compute aggregative indices with the weights given for the combined sales of A and B

in 1994 and 1995, and interpret their values.

Solution

For 1994, with 1993 as the base, we have:

Product A:	$Q_0 = 386$	$Q_1 = 397$	$w = 22$
Product B:	$Q_0 = 533$	$Q_1 = 542$	$w = 19$

and so the quantity index is:

$$100 \times \frac{\sum Qw}{\sum Q_0 w} = 100 \times \frac{(397 \times 22 + 542 \times 19)}{(386 \times 22 + 533 \times 19)} = 102.22$$

In the same way, the quantity index for 1995 is 103.86.

These figures show that the volumes of sales in 1994 and 1995 were 2.22 and 3.86 per cent, respectively, higher than in the base year, 1993.

Notice that there is a considerable difference between these results, with the aggregative method telling us that the quantity sold has risen by 18 per cent compared with the 27 per cent given by the relative method. The choice of construction method and of weighting system will often have a major impact on the value of the resulting index number. There is, therefore, considerable scope for selecting the method that gives the results that you want!

The construction of the UK retail price index

The RPI is compiled and published monthly. It is as far as possible a relative index with current expenditures as weights. In the press it is mainly published as an annual increase, but it is in fact a fixed-base index number and currently the base is January 1987. In the region of 130,000 price quotations are obtained monthly from a representative sample of retail outlets. It is representative in terms of geographic location within the United Kingdom, type and size of shop and timing over the month. Some 600 goods and services, called price indicators , are selected to be surveyed.

The selection of items for inclusion in the index and the weighting given to each itemis largely determined by a massive government survey called the Family Expenditure Survey (FES). In the course of each year, a representative sample of households is asked to keep a diary for 2 weeks listing all their spending. They are interviewed about major items such as housing and services. Decisions about items to include as price indicators will be based on the findings of the FES, and similarly the weights, expressed as parts per £ 1,000, are determined by the expenditure on each item by the average household. The weights are the expenditures of the previous year, which is probably as current as is practical. The items covered by the index are combined into 80 sections such as ‘men’s outerwear’ and ‘rail fares’ , which are in turn combined to give five broad groups before being averaged together in the overall index. The groups are:

- food and catering;
- alcohol and tobacco;
- housing and household expenditure;
- personal expenditure;
- travel and leisure.

There are some omissions from the RPI that may perhaps be thought important. The FES does not cover the very rich (top 4 per cent of the population) or those living on state benefits. In particular, a separate quarterly index is produced to monitor the prices of items purchased by pensioners. Life insurance, pension contributions and the capital element of mortgage repayments are not included; nor is tax. There is another index that takes taxation into account.

Using the RPI

Index numbers are widely used. The RPI, for example, is the measure of price inflation. Whenever you read that inflation is currently running at 3 per cent (or whatever) this means that, during the previous twelve months, the value of the RPI has increased by 3 per cent. In this section, we look at a number of applications of index numbers that, though not quoted as often in the media as the above example, are still of great commercial importance.

The first of these is index linking. In an attempt to protect people's savings and the incomes of some of the more vulnerable sections of society against the effects of inflation, many savings schemes, pensions and social benefits have, at various times in the past, been linked to the RPI in a way illustrated below.

Example

At the start of a year, the RPI stood at 340. At that time, a certain person's index-linked pension was £4,200 per annum and she had £360 invested in an index-linked savings bond. At the start of the following year, the RPI had increased to 360. To what level would the pension and the bond investment have risen?

Solution

First of all, the RPI has risen by 20 from 340. As a percentage, this is

$$\frac{20}{340} \times 100 = 5.88\%$$

The pension and the investment, being index-linked, would increase by the same percentage. The pension thus increases by 5.88% of £4,200 = £247 (nearest £), and the investment by 5.88% of £360 = £21.17. Hence, at the start of the year in question, the pension would be £4,447 per annum, and the investment would stand at £381.17.

Although the RPI is by far the most common index used in linking, there are others. For example, some house insurance policies have premiums and benefits that are linked to the index of house rebuilding costs, a far more suitable index than the RPI, which relates to general retail prices. At periods of low inflation, the practice of linking pay, pensions, savings and so on to the RPI is relatively harmless. When inflation is high, the automatic (and high) rises produced by linking tend to increase costs to industry and commerce, which therefore have to increase their prices. This in turn induces a rise in inflation that triggers off further index-linked rises in pay (and so on). Inflation is therefore seen to be exacerbated by this practice. Indeed, this has been observed in some countries with hyperinflation and where widespread index-linking has been introduced. The second common use is in the deflation of series, or the removal of inflation from a series of figures, so that they can be compared. An example will illustrate the simple arithmetical process involved.

The main criticism of deflation, and indeed another problem of index linking, is that we are applying an average figure (for price rises or whatever) to a particular set of people who may or may not be 'average' . To illustrate the effects this could have, consider the cases of: • an old-age pension for a single person being linked to the RPI;

• a brewery in the north of England deflating its profit figures by the RPI. In the first instance, the single pensioner cannot be considered 'average' in at least two senses. The RPI measures price rises for the average family, which a single pensioner certainly is not, and the income of a pensioner is generally considerably below average. The effect of this latter factor is that a pensioner will spend a considerably larger portion of his/her income on heating and lighting, and on food, than most other people. Thus, at times when prices of these commodities are rising faster than others (a situation that has occurred in the past), the RPI will underestimate the average price rises in a pensioner's 'basket' of goods and so linking to this index will leave him/her worse off. Indeed, a 'pensioners ' price index' has been introduced to overcome this.

The RPI measures price rises for all commodities, not just one type such as beer and related products that might be rising in price at a different rate from the average. The brewery in the latter example would therefore be advised to deflate by the 'alcoholic drinks' section of the RPI. Even then, prices might be rising at a faster or slower rate in the north of England, compared with this average UK figure. The real profit figures would then be too high or too low, respectively.

The technique of using the RPI to compare the purchasing power of wages etc. over several years is very important and it is a popular question with examiners.

REVIEW QUESTIONS

Q1-The managers of the catering division of a hospital wish to develop an index number series for measuring changes in food prices. As an experiment, they have chosen four items in general use that are summarised below, along with weights that reflect the quantities currently being bought.

	Prices per unit		
	1996	*1997*	*Weight*
Flour (kg)	0.25	0.30	10,000
Eggs (boxes)	1.00	1.25	5,000
Milk (litres)	0.30	0.35	10,000
Potatoes (kg)	0.05	0.06	10,000

A Write the formula for the aggregative price index denoting weight by w , 1996 price by P0 and 1997 price by P1 .

B. Calculate the aggregative price index for 1997 with base 1996, giving your answer to two decimal places.

C. Some of the following comments are advantages of current weights over base weights and some are advantages of base weights compared to current weights. Delete answers accordingly.

	Advantage of
(A) Current weights are expensive to obtain	current/base/incorrect
(B) Base-weighted indexes are preferable in times of high inflation.	current/base/incorrect
(C) Current-weighted indexes remain up to date	current/base/incorrect
(D) Current weights may be very difficult to obtain	current/base/incorrect
(E) Base weights are always out of date	current/base/incorrect
(F) Current-weighted indexes reflect changes in demand following price rises	current/base/incorrect
(G) Base weights need only be obtained once	current/base/incorrect
(H) Bases-weighted indexes can be meaningfully compared from one year to the next	current/base/incorrect

D. Which of the following are reasons why it is usual to weight the constituent parts of a price index?

(A) Weights reflect the prices of the constituent parts.
(B) Prices are always the price per some weight or other.
(C) Weights reflect the relative importance of the constituent parts.
(D) Weights must be known so that the entire ' shopping basket ' will weigh 100.

2. A company buys and uses five different materials. Details of the actual prices and quantities used for 1995, and the budgeted fi gures for 1996, are as follows:

Material	Weight	Actual 1995 Unit price £	Budgeted 1996 Unit price £
A	21	11	12
B	56	22	26
C	62	18	18
D	29	20	22
E	31	22	23

A. Calculate the missing values in the following table, taking 1995 as the base year and giving your answers correct to two decimal places where appropriate.

Material	Price relative	Price relative × Weight
A	109.09	2,290.89
B	118.18	6,618.08
C	100	F
D	110	G
E	104.55	3,241.05
Total		H

B If the total of the price relatives weights was 21,000, calculate the relative price index for the above data with base 1995, giving your answer to the nearest whole number.

C. If the relative price index calculated in 3.2 were 105, which of the following statements would be correct?
(A) Prices are budgeted to be 1.05 times their current levels.
(B) Turnover is budgeted to rise by 5 per cent.
(C) Quantities sold are budgeted to rise by 5 per cent.
(D) Prices are budgeted to rise by 5 per cent.
(E) Prices are budgeted to rise by 105 per cent.
(F) Prices are budgeted to rise by £ 5 on average.

D. Would you describe the price index in Part 3.2 as base weighted or current weighted or neither?
Delete as appropriate base/current/neither

3.The data below refers to average earnings index numbers in Great Britain for different sectors of industry, 1988 = 100, and the Retail Price Index, 1987=100.

Date	Whole economy	Production industries	Service industries	Retail Price Index
1988	100	100	100	107
Feb 89	104.6	104.9	104.4	111.5
May 89	107.5	108.1	107.2	115.0
Aug 89	109.1	109.2	108.7	115.8
Nov 89	112.8	112.9	112.7	118.5
Feb 90	114.0	114.3	113.7	120.2
May 90	118.5	118.2	118.6	126.2
Aug 90	120.9	119.7	121.1	128.1
Nov 90	123.8	123.7	123.0	130.0
Feb 91	124.7	125.2	123.8	130.9
May 91	128.1	129.2	127.1	133.5
Aug 91	130.8	130.2	130.4	134.1
Nov 91	130.8	131.8	129.7	135.6

I. Using 1988 as the base throughout, defl ate the production industries index at the points Nov 89 (A), Nov 90 (B) and Nov 91 (C), giving your answers to one decimal place.

II. Which of the following correctly states what a deflated production industries index of 105 would mean?

(A) Five per cent more goods and services can be bought by average earnings in Production Industries compared to the base year.

Average earnings in Production Industries are 5 per cent more than in the base year.

(C) Real earnings in the base year were 5 per cent more than in the year in question.

(D) Average earnings in Production Industries were 5 per cent more than the average for Great Britain.

III. A retired person from the service industries had a pension of £ 5,000 a year starting in May 1989 and updated each November in line with the average earnings index for that sector. Find the value of the pension (to the nearest £) in Nov 1989 (A), Nov 1990 (B) and Nov 1991 (C).

IV. If the answers to 4.3 were D 5,000, E 5,800 and F 6,000 fi nd the real values (in constant May 1989 prices) of the pension in Nov 89 (A), Nov 90 (B) and Nov 91 (C) to the nearest £ .

V. If the answer to (C) in Question 4.4 (i.e. the real value of the pension in Nov 91) was 5,100, which of the following statements would be correct?

(A) From May 89 to Nov 91 the purchasing power of the pension has increased by 100 per cent.

(B) From May 89 to Nov 91 the purchasing power of the pension has increased by £ 100 at Nov 91 prices.

(C) From May 89 to Nov 91 the purchasing power of the pension has increased by £ 100 at May 89 prices.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-6 FINANCIAL MATHEMATICS

CONTENTS

INTRODUCTION

Perhaps the most familiar use of mathematics in finance concerns interest calculations and other topics related to investments. In this chapter we shall develop formulae involving interest payments and equivalent rates of interest, but we shall also cover highly important but less well-known concepts such as present value. The practical applications dealt with in this chapter include loans, mortgages and regular saving plans.

Simple interest

One of the most basic uses of mathematics in finance concerns calculations of interest, the most fundamental of which is simple interest. Suppose £P is invested at a fixed rate of interest of r per annum (where r is a proportion) and that interest is added just once at the end of a period of n years. The interest earned each year is calculated by multiplying the rate of interest r by the amount invested, £P, giving an amount £rP. After n years the sum of £rPn will be credited to give a total at the end of the period, £V, of:

$$V = P + rPn$$

or:

$$V = P(1 + rn)$$

This well-known formula is often referred to as the simple interest formula.

Example

An amount of £5,000 is invested at a rate of 8 per cent per annum. What will be the value of the investment in 5 years' time, if simple interest is added once at the end of the period?

Solution

The interest rate in the formula needs attention: it is assumed that r is a proportion, and so, in this case, we must convert $r = 8$ per cent into a proportion:

$$r = 0.08$$

Also, we have

$$P = 5{,}000 \text{ and } n = 5$$

So

$$V = P(1 + rn) = 5{,}000(1 + 0.08 \times 5) = 5{,}000 \times 1.4 = 7{,}000$$

Thus, the value of the investment will be £7,000.

Compound interest

In practice, simple interest is not used as often as compound interest. Suppose £P is invested at a fixed rate of interest of r per annum and that interest is added at the end of each year; that is, it is compounded annually. After 1 year, the value of the investment will be the initial investment £P, plus the interest accrued, £rP , and so will be

$$P + rP = P(1 + r)$$

During the second year, the interest accrued will be r times the amount at the end of the first year, and so will be $rP(1 + r)$. The value at the end of the second year will be $P(1+r) + rP(1+r) = P(1+r)(1+r) = P(1+r)^2$

Proceeding in this way, after n years the value, £V, will be given by $V = P(1+r)^n$

As you will see, in financial mathematics we work with an annual ratio denoted by 1+ r rather than with the rate of interest.

Example

An amount of £5,000 is invested at a fixed rate of 8 per cent per annum. What amount will be the value of the investment in 5 years' time, if the interest is compounded:

(a) annually?
(b) every 6 months?

Solution

(a) The only part of this type of calculation that needs particular care is that concerning the interest rate. The formula assumes that r is a proportion, and so, in this case:

$$r = 0.08$$

In addition, we have $P = 5{,}000$ and $n = 5$, so:

$$V = P(1 + r)^5 = 5{,}000 \times (1 + 0.08)^5 = 5{,}000 \times 1.469328 = 7{,}346.64$$

Thus, the value of the investment will be £7,346.64.

It will be noted, by comparing this answer with that of Example 6.2.1, that compound interest gives higher values to investments than simple interest.

(b) With slight modifications, the basic formula can be made to deal with compounding at intervals other than annually. Since the compounding is done at 6-monthly intervals, 4 per cent (half of 8 per cent) will be added to the value on each occasion. Hence, we use $r = 0.04$. Further, there will be ten additions of interest during the five years, and so $n = 10$. The formula now gives:

$$V = P(1 + r)^{10} = 5{,}000 \times (1.04)^{10} = 7{,}401.22$$

Thus, the value in this instance will be £7,401.22.

In a case such as this, the 8 per cent is called a *nominal* annual rate, and we are actually referring to 4 per cent per 6 months.

Equivalent rates of interest

Value at the end of 1 year =1=1.04=1.04 =1.0816 This is the annual ratio that results not from 8 per cent per annum but from 8.16 per cent. Hence, the effective annual rate of interest is 8.16 per cent in this case.

Example

An investor is considering two ways of investing £20,000 for a period of 10 years:

- option A offers 1.5 per cent compounded every 3 months;
- option B offers 3.2 per cent compounded every 6 months. Which is the better option?

Solution

We have, for option A, $P = 20{,}000$; $n = 10 \times 4 = 40$; $r = 0.015$ and so:

$$V = 20{,}000(1 + 0.015)^{40} = £36{,}280.37$$

For option B, $P = 20{,}000$; $n = 10 \times 2 = 20$; $r = 0.032$ and so:

$$V = 20{,}000(1 + 0.032)^{20} = £37{,}551.21.$$

Hence, option B is the better investment.

In this case, $P = 20{,}000$ was given but it is not necessary to be given an initial value because £1 can be used instead.

Depreciation

The same basic formula can be used to deal with depreciation, in which the value of an item goes down at a certain rate. We simply ensure that the rate of 'interest' is negative.

Example

A company buys a machine for £20,000. What will its value be after 6 years, if it is assumed to depreciate at a fixed rate of 12 per cent per annum?

Solution

We have $P = 20{,}000$; $n = 6$; $r = -0.12$, hence:

$$V = P(1 + r)^6 = 20{,}000(1 - 0.12)^6 = 20{,}000 \times 0.4644041 = 9{,}288.08$$

The machine's value in 6 years' time will therefore be £9,288.08.

Example

A piece of capital equipment is purchased for £120,000 and is to be scrapped after 7 years. What is the scrap value if the depreciation rate is 8 per cent per annum?

Solution

$P = 120{,}000$, $n = 7$, $r = -0.08$, so

$$V = 120{,}000(1 - 0.08)^7 = 120{,}000(0.92)^7 = £66{,}941.59$$

More complex investments

We can return now to the evaluation of investments, but now considering situations where there are several different investments spread over a period of time.

Example

A man invests £3,000 initially and then £1,800 at the end of the first, second and third years, and finally £600 at the end of the fourth year. If interest is paid annually at 6.5 per cent, find the value of the investment at the end of the fifth year.

Solution

The diagram shows when the investments and evaluation take place.

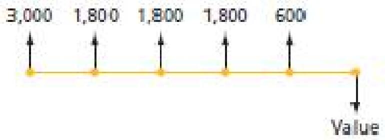

The £3,000 is invested for 5 years and grows to

$$3,000(1.065^5) = 4,110.26$$

The three sums of £1,800 are invested for 4, 3 and 2 years, and grow in total to

$$1,800(1.065^4 + 1.065^3 + 1.065^2) = 6,531.55$$

Finally, the £600 is invested for just 1 year and grows to

$$600 \times 1.065 = 639$$

The total value at the end of 5 years is £11,280.81

A *sinking fund* is a special type of investment in which a constant amount is invested each year, usually with a view to reaching a specified value at a given point in the future. Questions need to be read carefully in order to be clear about exactly when the first and last instalments are paid.

Geometric progressions

We worked out the final values for the sinking fund questions simply by using calculators in the usual fashion. However, a sinking fund could easily run for 20 years or more – in fact, the endowment element of some mortgages is a very common example of a sinking fund that would typically run for 20 – 25 years. So it is useful to digress briefly to discuss geometric progressions and how they can help with all this arithmetic. A geometric progression (GP) is a series of numbers of the form

$A, AR, AR^2, AR^3, \ldots$

Where A and R are numbers.

The particular feature that defines a GP is that, after an initial term, A, each term in the progression is a constant multiple (R) (or ratio) of the preceding one. We shall need to know the sum of the first n terms of such a series. Denoting this by S

$$S = \frac{A(R^n - 1)}{(R - 1)}$$

Example

If six annual instalments of £800 are made, starting immediately, at 5 per cent per annum, the value of the investment immediately after the sixth instalment is given by the following expression. Use GP theory to evaluate it.

$$£800(1.05^5 + 1.05^4 + 1.05^3 + 1.05^2 + 1.05 + 1)$$

Solution

The series in the brackets, viewed back to front, is a GP with $A = 1$, $n = 6$ and $R = 1.05$, so its sum is:

$$S = 1 \times \frac{1.05^6 - 1}{1.05 - 1} = 6.8019$$

Hence, the value of the fund is £800 × 6.8019 = £5,442 (to the nearest £).

Notice that *n* is given by the number of terms, not by the greatest power of *R*.

Present values

The present value of a sum of money to be paid or received in the future is its value at present, in the sense that it is the sum of money that could be invested now (at a certain rate of interest) to reach the required value at the subsequent specified time. Some examples will make this clearer and illustrate two ways of calculating present values.

Example

Find the present value of:

(a) £200 payable in 2 years' time, assuming that an investment rate of 7 per cent per annum, compounded annually, is available;

(b) £350 receivable in 3 years' time, assuming that an annually compounded investment rate of 6 per cent per annum, is available.

Solution

(a) From the definition, we need to find that sum of money that would have to be invested at 7 per cent per annum and have value £200 in 2 years' time. Suppose this is £X, then the compound interest formula gives:

$$V = P(1 + r)^n$$

Thus:

$$200 = X(1 + 0.07)^2$$

$$X = \frac{200}{1.1449} = 174.69$$

Thus, the present value is £174.69: that is, with an interest rate of 7 per cent, there is no difference between paying £174.69 now and paying £200 in 2 years' time.

(b) Using the compound interest formula again:

$$350 = X(1 + 0.06)^3$$

$$X = \frac{350}{1.191016} = 293.87$$

The present value is thus £293.87.

This method of calculation is said to be from first principles. The present value of a quantity, £V, discounted at 100 r per cent for n years is given by:

$$\frac{V}{(1 + r)^n}$$

Alternatively, present value tables that are provided in your exam give a present value factor (or discount factor) of 0.873 for n = 2 years at r = 7 per cent. This means that the present value of £ 1 for this combination is £ 0.873; hence the present value (PV) of £ 200 is:

$$200 \times 0.873 = £174.60$$

Similarly, the PV factor for $n = 3$ and $r = 6$ per cent is 0.840; so the PV of £350 is:

$$350 \times 0.840 = £294$$

Before leaving the PV tables, we note that they simply give the approximate values of

$$\frac{1}{(1 + r)^n}$$

Thereby simplifying the calculations considerably.

It can be seen that use of the tables loses some accuracy. When there are many such calculations, however, their use is considerably faster, and so tables are generally preferred.

However, their use is not always possible, since you will note that there are 'gaps' in the tables. For instance, the combination n = 2.5 years and r = 4.5 per cent does not appear in the tables, and so first principles would have to be used in an example involving these values.

Example

Calculate the present values of the following amounts. Use PV tables where you can, and first principles otherwise:

(a) £12,000 payable in 6 years' time at a rate of 9 per cent;
(b) £90,000 payable in 8 years' time at 14 per cent;
(c) £80,000 payable in 5 years' time at 6.3 per cent;
(d) £50,000 payable in 4 years and 3 months' time at 10 per cent.

Solution

(a) From tables, discount factor at 9 per cent for 6 years is 0.596:

$$PV = 12{,}000 \times 0.596 = £7{,}152$$

(b) Discount factor at 14 per cent for 8 years is 0.351:

$$PV = 90{,}000 \times 0.351 = £31{,}590$$

(c) The tables cannot be used for 6.3 per cent. $(1 + r) = 1.063$ and $n = 5$:

$$PV = \frac{80{,}000}{1.063^5} = £58{,}941.84 \text{ (to two d.p.)}$$

(d) The tables cannot be used for part years. To use a calculator, convert 4 years, 3 months into the decimal 4.25. So $n = 4.25$ and $(1 + r) = 1.1$:

$$PV = \frac{50{,}000}{1.1^{4.25}} = £33{,}346.56 \text{ (to two d.p.)}$$

Net present values – practical examples

In many situations, there are a number of financial inflows and outflows involved, at a variety of times. In such cases, the net present value (npv) is the total of the individual present values, after discounting each, as above.

Example

A company can purchase a machine now for £10,000. The company accountant estimates that the machine will contribute £2,500 per annum to profits for five years, after which time it will have to be scrapped for £500. Find the NPV of the machine if the interest rate for the period is assumed to be 5 per cent. (Assume, for simplicity, that all inflows occur at year ends.)

Solution

We set out the calculations in a systematic, tabular form:

After year	*Total inflow (£)*	*Discount factor*	*Present value (£)*
0	−10,000	1.000	−10,000
1	2,500	0.952	2,380
2	2,500	0.907	2,267.50
3	2,500	0.864	2,160
4	2,500	0.823	2,057.50
5	3,000	0.784	2,352
			1,217

Hence, the NPV is £1,217.

The fact that the middle four inflow values are the same (£2,500) means that the cumulative present value table (provided in your exam) can be used to calculate the total pv arising from an inflow of £2,500 at a constant interest rate of 5 per cent for each of 4 years, starting at the end of the first year:

After year	Total inflow (£)	Discount factor	Present value (£)
0	−10,000	1.000	−10,000
1–4	2,500 per year	3.546	8,865
5	3,000	0.784	2,352
			1,217

This table gives the NPV as £1,217, exactly as before.

The fact that the npv is positive means that the investment is more profitable than investing the original £ 10,000 at 5 per cent. In fact, you would need to invest £ 11,217 at 5 per cent in order to generate this particular set of positive cash flows, so this investment is worthwhile.

Problems using NPV in practice

One of the major difficulties with present values is the estimation of the 'interest rates' used in the calculations. Clearly, the appropriate rate(s) at the start of the time period under consideration will be known, but future values can be only estimates. As the point in time moves further and further into the future, the rates become more and more speculative. For this reason, the NPVs of investments A and C in Example 6.9.3 are so close as to be indistinguishable, practically speaking.

Many situations in which NPV might be involved are concerned with capital investments, with the capital needing to be raised from the market. For this reason, the 'interest rate(s)' are referred to as the cost of capital, since they reflect the rate(s) at which the capital market is willing to provide the necessary money.

Another problem with calculating net present value is the need to estimate annual cash flows, particularly those that are several years in the future, and the fact that the method cannot easily take on board the attachment of probabilities to different estimates. Finally, it is a usual, although not an indispensable, part of the method to assume that all cash flows occur at the end of the year, and this too is a potential source of errors. With easy access to computers it is now possible to calculate a whole range of NPVs corresponding to worst-case and best-case scenarios as well as those expected, so to some extent some of the problems mentioned above can be lessened.

Annuities

An annuity is an arrangement by which a person receives a series of constant annual amounts. The length of time during which the annuity is paid can either be until the death of the recipient or for a guaranteed minimum term of years, irrespective of whether the annuitant is alive or not. In other types of annuity, the payments are deferred until sometime in the future, such as the retirement of the annuitant.

$$\frac{1}{r} - \frac{1}{r(1+r)^t}$$

When two or more annuities are being compared, they can cover different time periods and so their net present values become relevant. In your exam you will be given the following formula for the NPV of a £ 1 annuity over t years at interest rate r , with the first payment 1 year after purchase.

The cumulative present value tables can also be used.

Example

An investor is considering two annuities, both of which will involve the same purchase price. Annuity A pays £5,000 each year for 20 years, while annuity B pays £5,500 each year for 15 years. Both start payment 1 year after purchase and neither is affected by the death of the investor.

Assuming a constant interest rate of 8 per cent, which is the better?

Solution

Using tables, the cumulative PV factors are 9.818 for A and 8.559 for B.

Hence PV of annuity A = 5,000 × 9.818 = £49,090, and

PV of annuity B = 5,500 × 8.559 = £47,074.5

You will only be able to use the tables given in your exam if the period of the annuity is twenty years or less and if the rate of interest is a whole number. It is, therefore, essential that you learn to use the formula as well. You will notice that there is some loss of accuracy, due to rounding errors, when tables are used.

Using the above formula:

$$\text{Factor for the NPV of A} = \frac{1}{0.08} - \frac{1}{0.08(1 + 0.08)^{20}} = 9.818147$$

and so the NPV of A is:

5,000 × 9.818147 = £49,091 (to the nearest £)

Similarly:

$$\text{Annuity factor for the NPV of B} = \frac{1}{0.08} - \frac{1}{0.08(1 + 0.08)^{15}} = 8.559479$$

and so the NPV of B is:

5,500 × 8.559479 = £47,077

From the viewpoint of NPVs, therefore, annuity A is the better choice. As we have already seen, however, there are two further considerations the investor may have. Assuming constant interest rates for periods of 15 or 20 years is speculative, so the NPVs are only approximations: they are, however, the best that can be done and so this point is unlikely to affect the investor's decision. More importantly, although any payments after the investor's death would go to their estate, some people may prefer more income 'up front' during their lifetime. Unless the investor is confident of surviving the full 20 years of annuity A, they may prefer annuity B = especially as the two NPVs are relatively close to each other. A further example will demonstrate an NPV being expressed as an equivalent annuity.

Finally, there is the concept of perpetuity. As the name implies, this is the same as an annuity except that payments go on forever. It is therefore of interest to those who wish to ensure continuing payments to their descendants, or to some good cause. It must be recognised, however, that constant payments tend to have ever decreasing value, owing to the effects of inflation, and so some alternative means of providing for the future may be preferred.

PV of a perpetuity

As t becomes very large, the second term in the formula for the PV of an annuity gets smaller and smaller, to the point where it becomes zero, and the factor for the NPV of a perpetuity simplifies considerably to:

$$\frac{1}{r}$$

Example

Consider the position of A and B in Example 6.11.1 if they were perpetuities.

Solution

Using the formula:

$$\text{NPV of perpetuity A} = \frac{1}{0.08} \times 5{,}000 = £62{,}500$$

$$\text{NPV of perpetuity B} = \frac{1}{0.08} \times 5{,}500 = £68{,}750$$

The NPV of annuity B is therefore higher. The assumption of a constant interest rate of 8 per cent for ever is clearly highly unlikely to materialise but, as before, it is all that can be done. In purchasing a perpetuity, an investor is not interested particularly in the income during their lifetime, so the latter consideration of Example is not pertinent here: perpetuity B is unequivocally the better.

You can check this easily. At 6 per cent, the interest on £200,000 is £12,000 per annum, so the annuity can be paid indefinitely without touching the capital.

Loans and mortgages

Most people will be aware that, when a mortgage is taken out on a property over a number of years, there are several ways of repaying the loan. We shall concentrate here on repayment mortgages, because they are among the most popular, and because they are the only ones that involve complex mathematical calculations. The features of a repayment mortgage are:

- a certain amount, £M, is borrowed to be repaid over n years;
- interest (at a rate r) is added to the loan retrospectively at the end of each year; and
- a constant amount, £P, is paid back each year by the borrower, usually in equal monthly instalments.

Viewed from the standpoint of the lender, a repayment mortgage is an annuity. The lender pays the initial amount (M) for it and in return receives a series of constant annual payments (P) for n years. The relationship between these variables is given by putting M equal to the present value of the annuity, using either tables or formula as appropriate.

Example

(a) A £30,000 mortgage is taken out on a property at a rate of 12 per cent over 25 years. What will be the gross monthly repayment?

(b) After 2 years of the mortgage, the interest rate increases to 14 per cent: recalculate the monthly repayment figure.

Solution

(a) Equating present values gives:

$$30{,}000 = P\left(\frac{1}{0.12} - \frac{1}{0.12 \times 1.12^{25}}\right) = 7.843139P$$

giving $P = 30{,}000/7.843139 = £3{,}825$ per annum (nearest £) and a monthly repayment of £318.75 (to two d.p.).

(b) After 2 years, immediately after the second annual repayment, the amount still owing is:

$$30{,}000 \times 1.12^2 - 3{,}825 \times 1.12 - 3{,}825 = £29{,}523$$

The mortgage now has 23 years to run and at 14 per cent interest we have:

$$29{,}523 = P\left(\frac{1}{0.14} - \frac{1}{0.14 \times 1.14^{23}}\right) = 6.792056P$$

giving $P = 29{,}523/6.792056 = £4{,}346.70$ per annum and a monthly repayment of £362.22 (two d.p.).

Internal rate of return

We have seen that if NPV is positive it means that the project is more profitable than investing at the discount rate, whereas if it is negative then the project is less profitable than a simple investment at the discount rate.

If NPV is zero the project is identical in terms of profit to investing at the discount rate, and hence this rate of interest gives us the rate of return of the project.

The internal rate of return (IRR) is the discount rate at which NPV is zero. It is obtained generally by a trial and error method as follows.

1. Find a discount rate at which NPV is small and positive;
2. Find another (larger) discount rate at which NPV is small and negative;
3. Use linear interpolation between the two to find the point at which NPV is zero.

Example

Find the IRR for the following project.

Time	*Cash flow (£'000)*
0	(80)
1	40
2	30
3	20
4	5

Solution

The question offers no guidance as to what discount rates to try, so we will select 5 per cent randomly. Since 5 per cent turns out to give a positive NPV we now randomly select 10 per cent in the hope that it will give a negative NPV.

Time	*Cash flow £'000*	*PV (5%) £'000*	*PV (10%) £'000*
0	(80)	(80.000)	(80.000)
1	40	38.095	36.364
2	30	27.211	24.793
3	20	17.277	15.026
4	5	4.114	3.415
Net present value:		6.697	(0.402)

We can now use either (a) a graphical method (Figure 6.1) or (b) a calculation based on proportions.

(a) *Graphical method*

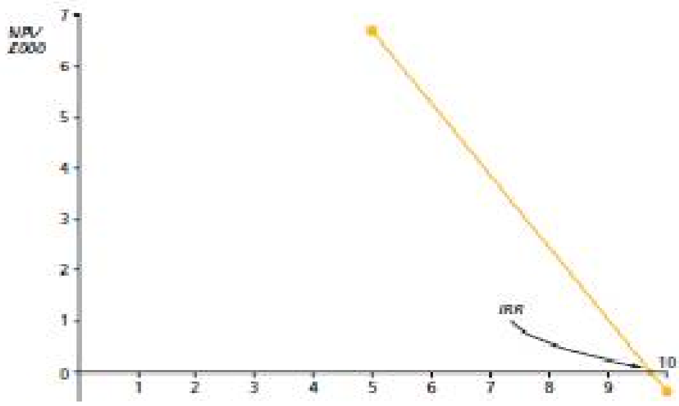

Graph of NPV on discount rate

REVIEW QUESTIONS

1. A company is planning a new product for which a 10-year life is anticipated. The product is expected to follow a typical life cycle of growth, maturity and decline with a cash flow of £56,000 in year 1. Estimates of cash flows expected from years 2 – 10 are as follows:

Year	Percentage rate of change expected on the previous year's cash flow
2	+2
3	+5
4	+10
5	+10
6	+10
7	+5
8	−1
9	−3
10	−5

Assume all cash flows arise at year ends. Work throughout to the nearest £.

I. Calculate the cash flow expected in the tenth year.

II. Supposing that the expected cash flows are as follows, complete the following table, to calculate the net present value of the expected cash flows, by filling in the appropriate numerical values in the spaces indicated by the letters. Use a discount rate of 8 per cent per annum and use

Year	Cash flow (£)	Discount factor	Present value (£)
1	56,000	...	...
2	57,000	...	A
3	60,000	...	...
4	66,000	...	...
5	72,000	...	B
6	79,000	...	...
7	84,000	...	...
8	82,000	...	C
9	80,000	...	...
10	75,000	...	...
Net present value			

III. If the net present value was £450,000 what is the maximum amount that the company could invest now in the product if it is to meet a target of an 8 per cent return?

IV. If the net present value was £450,000 and, if the company needs to borrow at 8 per cent in order to fi nance the project, which of the following statements is/are correct?

(A)£450,000 is the profi t that the company expects to make if they can borrow at 8 per cent.

(B) £450,000 is the maximum profi t that the company might make if they borrow at 8 per cent.

(C) £450,000 is the maximum that the company should borrow if they wish to make a profit.

(D) £450,000 is the present value of the profit that the company expects to make if they borrow at 8 per cent.

V. The following list includes three valid comments about the use of discounting in appraising investment decisions. Indicate which comments are correct.

a. Discounting takes account of the time value of cash flows;
b. Discounting ignores risk;
c. Discounting is more accurate because present value tables are available;
d. Discounting does not involve making estimates;
e. Discounting is of use only if an appropriate discount rate can be obtained.

2. To carry out identical tasks, a company uses several machines of the same type, but of varying ages. They have a maximum life of five years. Typical financial data for a machine are given below:

Time	*Now*	*After 1 year*	*After 2 years*	*After 3 years*	*After 4 years*	*After 5 years*
Initial cost	£10,000	–	–	–	–	–
Maintenance + service costs	£1,000	£1,500	£2,000	£2,500	£3,000	£5,000
Resale value if sold	–	£7,000	£5,000	£3,500	£2,500	£2,000

The rate of interest is 15 per cent.
These machines are assumed to produce flows of revenue that are constant.

I. Complete the following table, to calculate the net present value of the costs if the machine is kept for 5 years, by filling in the appropriate numerical values in the spaces indicated by the letters. Work to the nearest £ throughout.

Time (years)	*Outflows (£)*	*Discount factor*	*Present value (£)*
0	A	D	...
1	B	...	...
2	...	E	...
3	...	...	F
4	...	...	...
5	C	...	...

II. The present values of the costs of keeping the machine for between 1 and 4 years are given below. Convert each of these into an annual equivalent amount and write them in the final column of the table, for one mark each.

Year of scrapping	*Present value of cost (£)*	*Annuity equivalent (£)*
1	6,215	...
2	10,037	...
3	13,159	...
4	15,748	...
5	18,669	...

III. If the annual equivalent amounts calculated above were as follows, what would be the most economical age at which to replace the machines?

Year of scrapping	*Annual equivalent (£)*
1	7,000
2	6,000
3	5,800
4	5,500
5	5,600

3. A company is planning capital investment for which the following cash flows have been estimated:

Time	Net cash flow (£)
Now	(10,000)
At the end of year 1	500
At the end of year 2	2,000
At the end of year 3	3,000
At the end of year 4	4,000
At the end of year 5	5,000
At the end of year 6	2,500
At the end of year 7	2,000
At the end of year 8	2,500

The company has a cost of capital of 15 per cent.

I. Complete the following table by filling in the appropriate numerical values in the spaces indicated by the letters. Work to the nearest £ throughout. Cash flows are all at the ends of years unless stated otherwise.

Year end	Net cash flow (£)	Discount factor at 15%	Present value (£)
Now	(10,000)	A	C
1	500	B	...
2	2,000	...	...
3	3,000	...	...
4	4,000	...	...
5	5,000	...	...
6	2,500	...	...
7	2,000	...	...
8	2,500	...	D

II. Which of the following defines the ‘ internal rate of return ’ .
- (A) The current discount rate used by a company.
- (B) The discount rate at which net present value is zero.
- (C) The discount rate recommended by a trade association or similar.
- (D) The discount rate at which cash flows total zero.

III. 6.4.3 The net present value is £(543) when the discount rate is 20 per cent and is £3,802 when it is 10 per cent. Which of the following statements about the value of the internal rate of return (IRR) is correct in this case?
- (A) The IRR must be below 10 per cent.
- (B) The IRR must lie between 10 per cent and 15 per cent.
- (C) The IRR must lie between 15 per cent and 20 per cent.
- (D) The IRR must be greater than 20 per cent.
- (E) None of the above statements is correct.

IV. Calculate the approximate internal rate of return (to the nearest whole per cent point) of this investment without calculating any further net present values.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-7 CORRELATION AND REGRESSION

CONTENTS

INTRODUCTION

In the next two chapters we look at one of the major applications of statistics, namely forecasting. Although there are a number of ways of producing forecasts that involve little or no mathematics, we shall concentrate here on two of the most important quantitative approaches, causal and extrapolative. A causal approach is based on the assumption that changes in the variable that we wish to forecast are caused by changes in one or more other variables. With an extrapolative approach, we examine past data on the variable that is to be forecast, in order to determine any patterns the data exhibits. It is then assumed that these patterns will continue into the future: in other words, they are extrapolated. Let us begin with the casual approach, and, to simplify matters, we shall deal with the case in which the variable to be forecast (the dependent variable, y) depends on only one other variable (the independent variable, x). Such data is called bivariate because in each situation we have two pieces of information, denoted x and y . The dependent variable, y, is also referred to as the response variable and the independent variable, x, as the influencing variable. So x is assumed to influence a response in y. Before actually looking at how to produce forecasts in such situations, we must consider the question as to how we know that changes in y are caused by changes in x (alternatively: how we know that y depends on x). The answer to this involves the study of correlation.

Correlation

Two variables are said to be correlated if they are related to one another, or, more precisely, if changes in the value of one tend to accompany changes in the other. Now, we have already used the (x, y) notation of Chapter 2, and this initially suggests a graphical approach: if there are pairs of data available on the variables x and y, then these can be plotted as points against a set of x- and y -axes. The result is known as a scatter diagram , scatter graph or sometimes a scatter plot .

Example

A company is investigating the effects of its advertising on sales. Consequently, data on monthly advertising and sales in the following month are collated to obtain:

Advertising expenditure in month (£'000)	*Total sales in following month (£'000)*
1.3	151.6
0.9	100.1
1.8	199.3
2.1	221.2
1.5	170.0

Plot these data on a scatter diagram.

Solution

Since the company is interested in how advertising affects sales, it is clear that sales should be the dependent variable, y, and advertising the independent, x. The scatter diagram is shown in Figure 7.1.

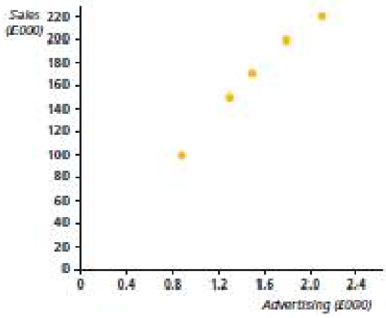

There are many key terms in italics in these introductory paragraphs.

Since, unlike in Chapter 2, we have not been told that y is a function of x here, the points have not been joined up. In particular, although the points appear to be close to lying on a straight line, they do not lie exactly on a line; we do not know that a linear function is involved and so none has been drawn in.

The scatter diagram in the above example seems to show a case where the two variables are related to one another. Further, the relationship seems to be of an approximately linear nature: it is an example of linear correlation.

Since the approximation is so good, and the points are close to a straight line, we talk of strong linear correlation. Finally, as the gradient of the 'line' is positive, it is positive linear (or direct) correlation.

Scatter diagrams in Excel

The data in the above example can be plotted onto a scatter diagram in Excel. To do this the data must first be entered into the spreadsheet. To create the chart select the range a 3: b 7 and click on the chart icon. Select X-Y scatter and choose the first chart option. The data and the resulting chart are shown in Figure 7.2 .

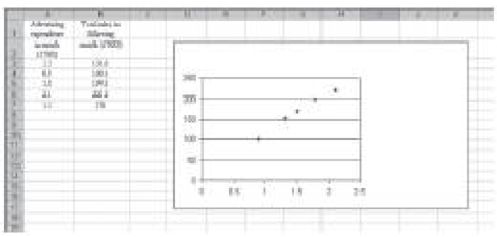

Example

Most of the examples in this chapter relate to the following table. A company owns six sales outlets in a certain city. The sales last year of two of its key products are given below, together with the sizes of each outlet:

Most of the examples in this chapter relate to the following table. A company owns six sales outlets in a certain city. The sales last year of two of its key products are given below, together with the sizes of each outlet.

Outlet	Floor space m^2	Sales of L '000 units	Sales of M '000 units
A	75	22.4	30.7
B	60	21.1	12.9
C	108	29.6	47.1
D	94	27.1	38.8
E	92	27.0	41.5
F	130	36.9	79.0

Did you plot size of outlet as x in both cases? We are investigating the effect of size on sales, so sales must be the dependent variable (see Figures 7.3 and 7.4).

Figure 7.3 shows further examples of:

- very weak positive linear correlation, in which y shows a slight tendency to increase as x does;
- strong negative linear (or inverse) correlation, in which y shows a strong tendency to decrease as x increases; and
- non-linear correlation, in which x and y are clearly related, but not in a linear fashion.

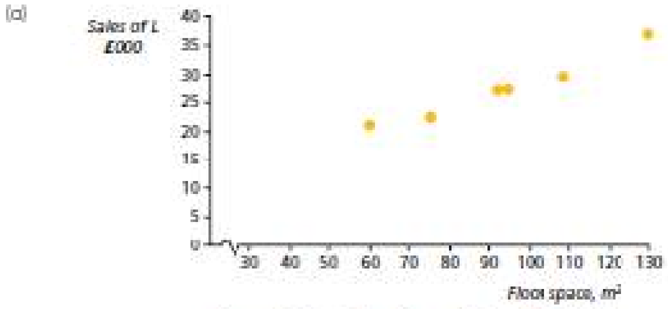

Figure 7.3 Scatter diagram (Example 7.2.2(a))

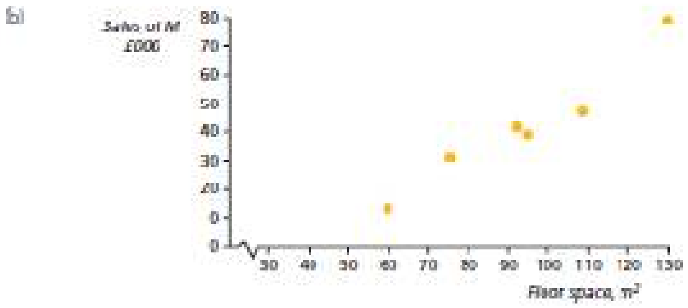

Although such scatter diagrams are useful for getting a feel for the presence or otherwise of correlation, it is often difficult to judge what is 'weak' and what is 'strong', and, indeed, whether a large number of points on a diagram constitute any correlation at all. Therefore, an objective measure of correlation is needed.

Example

From your own experience, try to think of pairs of variables that might have the different degrees of correlation, from weak to strong and from negative to positive.

Solution

These are just some examples:

- costs probably have a strong positive correlation with the number of units produced;
- number of deaths on the roads probably has a middling positive correlation with traffic levels;

● the level of street crime is often thought to relate to the level of visible policing, so the correlation would be negative but probably not strong;
● a strong negative correlation would probably be found if almost any measure of bodily function, such as the condition of the heart, were compared with age in adults, although the graph is unlikely to be perfectly linear.

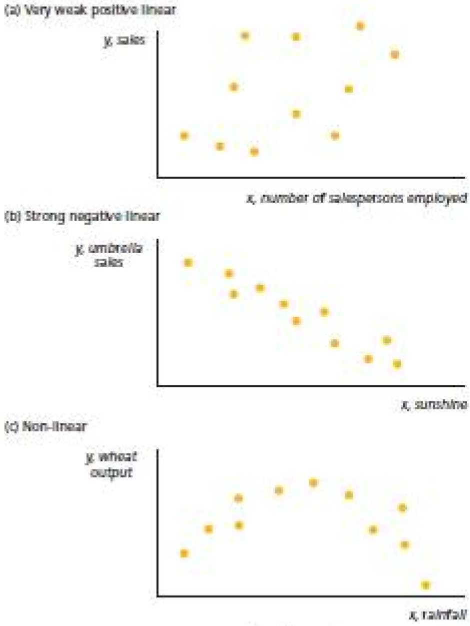

Figure 7.5 Examples of correlation

Pearson's correlation coefficient

The statistician Pearson developed a measure of the amount of linear correlation present in a set of pairs of data. Pearson's correlation coefficient, denoted r , is defined as:

$$r = \frac{n\Sigma xy - \Sigma x \Sigma y}{\sqrt{(n\Sigma x^2 - (\Sigma x)^2)(n\Sigma y^2 - (\Sigma y)^2)}}$$

where n is the number of data points.
This measure has the property of always lying in the range = 1 to = 1, where:

- r =+1 denotes perfect positive linear correlation (the data points lie exactly on a straight line of positive gradient);
- r = - 1 denotes perfect negative linear correlation (again the data points lie on a straight line but with a negative gradient); and
- r = 0 denotes no linear correlation.

The strength of a correlation can be judged by its proximity to -1 or +1: the nearer it is (and the further away from zero), the stronger is the linear correlation. A common error is to believe that negative values of r cannot be strong. They can be just as strong as positive values except that y is decreasing as x increases.

Example

Evaluate Pearson's correation coefficient for the data on sales and advertising spend in Excmple 7.2.1, and interpret its value.

Solution

As with previous calculations involving summations, we facilitate the calcuations by setting them out in columns:

x	y	x^2	y^2	xy
1.3	151.6	1.69	22,982.56	197.08
0.9	100.1	0.81	10,020.01	90.09
1.8	199.3	3.24	39,720.49	358.74
2.1	221.2	4.41	48,929.44	464.52
1.5	170.0	2.25	28,900.00	255.00
7.6	842.2	12.40	150,552.50	1,365.43

$$\text{Thus, } r = \frac{(5 \times 1365.43) - (7.6 \times 842.2)}{\sqrt{[(5 \times 12.4) - 7.6^2][(5 \times 150{,}552.5) - 842.2^2]}} = \frac{426.43}{\sqrt{4.24 \times 43{,}451.66}} = 0.993$$

The value of Pearson's correlation coefficient in this case is 0.993. The arithmetic in such a calculation can be seen to be potentially very tedious. It is worthwhile investigating the availability of any computer packages or special functions on a calculator in order to ease the computation of correlation coefficients. Note that a simple check is that your calculated value for the correlation coefficient must be between -1 and 1.

This formula is given in your exam so you do not need to worry about remembering it.

The value of the coefficient in this case is clearly very close to the value 1, indicating a very strong positive linear correlation, and reflecting the close proximity of the points in Figure to a straight line.

Example

Using the data on floor space and sales from Example 7.2.2, evaluate Pearson's correlation coefficients for:

(a) sales of L and size;
(b) sales of M and size.

Solution

(a) The necessary summations are $n = 6$; $\Sigma x = 559$; $\Sigma y = 164.1$; $\Sigma x^2 = 55{,}089$; $\Sigma y^2 = 4{,}648.15$; $\Sigma xy = 15{,}971.2$. Hence:

$$r = \frac{(6 \times 15{,}971.2) - (559 \times 164.1)}{\sqrt{(6 \times 55{,}089 - 559^2)(6 \times 4{,}648.15 - 164.1^2)}}$$
$$= \frac{4{,}095.3}{\sqrt{18{,}053 \times 960.09}} = \frac{4{,}095.3}{\sqrt{(17{,}332{,}504.77)}} = \frac{4{,}095.3}{4{,}163.23} = 0.984$$

This is a very strong positive correlation between outlet size and sales of L.

(b) For M, the summations are $n = 6$; $\Sigma x = 559$; $\Sigma y = 250$; $\Sigma x^2 = 55{,}089$; $\Sigma y^2 = 12{,}796$; $\Sigma xy = 25{,}898.5$

These result in r = 0.974, which similarly shows a very strong positive correlation between outlet size and sales of M.

Many students initially find these calculations very difficult. Even if you got the right answer, you may find it useful to run through the calculations once more. Correlation is a very important topic in business mathematics.

Interpreting correlation coefficients

In general, it is not always as straightforward to interpret a value of r as in the above case. Although it would be inappropriate for the purpose of this text to go into detailed theory, it must be noted that the sample size (n) has a crucial effect: the smaller the value of n, the 'easier' it is for a large value of r to arise purely by accident.

Very rough guidelines are that, with a sample of ten data points, a minimum correlation of about 0.6 is needed before you can feel confident that any sort of linear relationship holds. With twenty data points, the minimum correlation needed is about 0.4.

Extrapolation is a further danger in the interpretation of r. If your x-values range from 0.9 to 2.1, then r = 0.993 tells you that there is a near-perfect linear relationship between x and y in that range. However, you know nothing at all about the relationship outside that range. It may or may not continue to be linear. The process of drawing conclusions outside the range of the data is called extrapolation. It often cannot be avoided but it leads to unreliable conclusions.

It is possible that an apparently high correlation can occur accidentally or spuriously between two unconnected variables. There is no mathematical way of checking when this is the case, but common sense can help. In the case under discussion, it seems plausible that sales and the advertising spend are connected, and so it would seem reasonable to assume that this is not an accidental or spurious correlation.

More importantly here, two variables can be correlated because they are separately correlated to a hidden third variable. The size of the region could well be such a variable: larger regions would tend to have larger sales figures and the management of larger regions would tend to have larger advertising budgets. It is therefore possible that this high correlation coefficient may have arisen because the variable 'sales' is highly correlated with size of region, advertising expenditure is highly correlated with size of region, but sales and advertising spend are not themselves directly connected.

Even if this third variable effect is not present, we still cannot conclude that y depends on x. The strong correlation lends support to the assumption that this is so, but does not prove it. Correlation cannot be used to prove causation.

In your assessment, interpreting correlation is as important as calculating the coefficient.

Rank correlation: Spearman's coefficient

There are occasions when the degree of correlation between two variables is to be measured but one or both of them is not in a suitable quantitative form. In such circumstances, Pearson's coefficient cannot be used, but an alternative approach = rank correlation =

might be appropriate. The most common measure of this type is Spearman's rank correlation

coefficient, R:

$$R = 1 - \frac{6\,\Sigma d^2}{n(n^2 - 1)}$$

where d denotes the difference in ranks, and n the sample size. You do not need to remember this formula because it will be given in your exam. The arithmetic involved in calculating values of this coefficient is much easier than that for Pearson's coefficient, as the following example illustrates.

Example

As part of its recruitment procedures, a company awards applicants ratings from A (excellent) to E (unsatisfactory) for their interview performance, and marks out of 100 for a written test. The results for five interviewees are as follows.

Interviewee	*Interview grade*	*Test score*
a	A	60
b	B	61
c	A	50
d	C	72
e	D	70

Calculate the Spearman's rank correlation coefficient for this data, and comment on its value.

Solution

In order to apply the formula, the grades and scores are ranked, with the best scores given a rank of 1. Notice how interviewees a and c share the best interview grade. They therefore share the ranks 1 and 2 to give 1.5 each.

Interviewee	*Rank of interview grade*	*Rank of test score*	*d*	d^2
a	1.5	4	−2.5	6.25
b	3	3	0	0.00
c	1.5	5	−3.5	12.25
d	4	1	3	9.00
e	5	2	3	9.00
				36.50

Hence:

$$R = 1 - \frac{6\,\Sigma d^2}{n(n^2-1)} = 1 - \frac{6 \times 36.50}{5(25-1)} = -0.825$$

The high negative value (near to = 1) indicates that interview grades and test scores almost totally disagree with each other – good interview grades go with the lowest test scores and vice versa. This should concern the company, as it may mean that one or both methods of judging applicants is faulty. The interpretation of R -values is similar to that for r. Warnings similar to those in Section 7.4 also apply when judging values of R.

There seems to be some positive correlation between price and taste, with the more expensive wines tending to taste better. Given the sample size the result is not really reliable and it cannot be extrapolated to wines costing more than £4. Had the taste rankings been allocated in the opposite order, the correlation would be =0.59. Students often find this calculation difficult and it is worth running through it again if you had problems.

Probably the most common errors are either forgetting to subtract from 1 or subtracting the numerator from 1 prior to dividing by the denominator.

Tied rankings can also be difficult. B, D and H all cost £2.99. Had they been marginally different they would have been ranked 2, 3 and 4. Since they are identical, they each have the rank of 3 (the average of 2, 3 and 4). Similarly F and G share the ranks 7 and 8 by giving them an average 7.5 each.

The d column is obtained by subtracting rank of taste minus rank of price, but it would be equally correct the other way round.

Which correlation coefficient to use

If the data have already been ranked, there is no option but to use the rank correlation coefficient (R). Where actual values of x and y are given, Pearson's coefficient (r) should generally be used since information is lost when values are converted into their ranks. In particular, Pearson's coefficient must be used if you intend to use regression for forecasting (see later). The only advantages in converting actual data into ranks and using Spearman's coefficient are:

1. that the arithmetic is easier, but this is a minor point given computers and scientific calculators;
2. that Spearman checks for a linear relationship between the ranks rather than the actual figures. If you simply want to confirm, say, that the variables increase together but have no concern about the linearity of the relationship, you might prefer to use the rank correlation coefficient.

Regression

The preceding sections give us a way of checking on whether it may be valid to assume that one variable, y, depends on another, x. We now proceed to consider how, after making such an assumption, y can be forecast from x.

For simplicity, we restrict our attention to instances of linear correlation. Thus, we are interested in situations where the dependence is in the form of a straight line. As we saw in Chapter 2, this involves equations of the type $y = a+bx$

where a and b are numbers. We are, therefore, initially concerned with determining suitable straight line(s) for the particular problem.

The least-squares criterion

The approach is illustrated through an example.

A company has the following data on its sales during the last year in each of its regions and the corresponding number of salespersons employed during this time:

Region	Sales (units)	Salespersons
A	236	11
B	234	12
C	298	18
D	250	15
E	246	13
F	202	10

Develop a linear model for forecasting sales from the number of salespersons.

Solution

The linear correlation coefficient between these two variables can be shown to be 0.948. This high value encourages us to assume that sales, y, might depend on the number of salespersons, x, in a linear way.

The scatter diagram for the data is shown in Figure 7.6. For convenience of drawing, the scales on the axes do not start from zero. However, this has the effect of exaggerating the divergences from linearity. A truer impression would be obtained from a graph containing the origin, but this would not be so easy to draw.

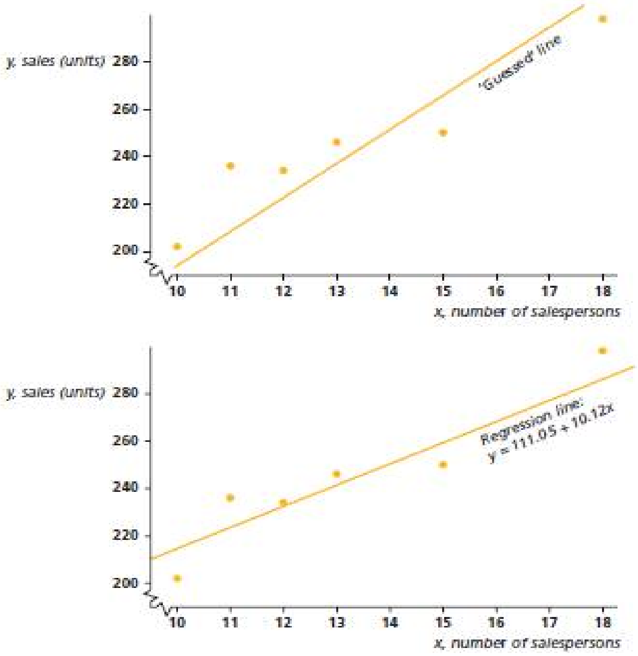

In the upper part of the figure, a straight line has been gauged or 'guessed' by using a ruler to draw a line that appears to be 'close' to all five data points. We have deliberately fitted a very poor guessed line so that the errors are clear. If you do have to fit a line 'by eye', the aim is to follow the slope of the points and to draw the line as far as possible through the centre of the points with roughly equal numbers either side. The lower part of Figure 7.5 shows the best possible fitted line.

This approximate approach may well be accurate enough in many instances, but it is certainly arbitrary. A number of different, equally plausible, lines could be drawn in: the question is, how can you judge whether one line is 'better' than another? Indeed, which is the 'best' ?

If we look at the 'guessed' line, it is clear that there are discrepancies between actual y-values and those obtained from the line. There are y-errors present, and the sizes of these enable us to judge one line against another. Examples of y -errors in this instance are:

$x = 13$: actual $y = 246$
y from line = 239 (approximately)
y-error in line = −7

$x = 15$: actual $y = 250$
y from line = 266
y-error in line = +16

Some errors are positive and some negative. Simply adding the errors to judge the 'goodness' of the line, therefore, would not be a sensible idea, as positive errors would tend to be cancelled out by negative ones. To eliminate this effect, we square the errors, and deem one line 'better' than another if its sum of squared errors is lower. The 'best' line is thus the one with the least sum of squared errors: the so-called least-squares regression line of y on x . Without going through the theory, this can be shown to have equation

$$y = a + bx$$

where:

$$b = \frac{n\Sigma xy - (\Sigma x)(\Sigma y)}{n\Sigma x^2 - (\Sigma x)^2}$$

and

$$a = \overline{y} - b\overline{x} \quad (\overline{y}, \overline{x}\text{: are the means of } y \text{ and } x \text{, respectively})$$

The calculation of a and b is set out in a familiar tabular form:

x	y	x^2	xy
11	236	121	2,596
12	234	144	2,808
18	298	324	5,364
15	250	225	3,750
13	246	169	3,198
10	202	100	2,020
79	1,466	1,083	19,736

$$b = \frac{(6 \times 19{,}736) - (79 \times 1{,}466)}{(6 \times 1{,}083) - 79^2} = \frac{2{,}602}{257} = 10.12$$

$$\bar{x} = \frac{79}{6} = 13.17$$

$$\bar{y} = \frac{1{,}466}{6} = 244.33$$

and so:

$$a = 244.33 - (10.12 \times 13.17) = 111.05$$

Thus, the least-squares regression line in this case is $y = 111.05 + 10.12x$

This line has been plotted on the lower scatter diagram of Figure 7.6 by calculating the coefficients of any two points on the line – for example:

- when $x = 10$, $y = 111.05 + 10.12 \times 10 = 212.25$
- when $x = 17$, $y = 111.05 + 10.12 \times 17 = 283.09$

These points should then be plotted on the graph and joined by a straight line.

Interpreting a and b

You may remember from Chapter 2 that in the equation of a straight line, y = a + bx , a is the intercept on the y -axis and b is the gradient or slope of the line. Expressed slightly differently, that means that a is the value of y when x = 0, and b is the increase in y for each unit increase in x. The b -value of 10.12 tells us that each extra salesperson generates an extra 10.12 sales (on average), while the a -value of 111.05 means that 111.05 units will be sold if no salespeople are used. The latter conclusion may well be nonsensical because x = 0 is outside the range of the data, but we return to this later.

It should be noted that, unlike Pearson's correlation coefficient, these calculations do not use Σy^2 , and so no time has been wasted evaluating it. Also, it will be appreciated that calculations such as these can involve potentially large numbers, and so it might be worthwhile to use an available computer package or statistical function on a calculator.

The interpretation of a and b is a frequent exam question.

Example

Using the data given in Example 7.2.2 on sales and floor space, find the least-squares regression lines for:

(a) sales of L against size;
(b) sales of M against size.

Interpret the values of a and b in your answers.

Solution

(a) Using the summations calculated in Example 7.3.2:

$$b = \frac{n\Sigma xy - (\Sigma x)(\Sigma y)}{n\Sigma x^2 - (\Sigma x)^2} = \frac{(6 \times 15{,}971.2) - (559 \times 164.1)}{(6 \times 55{,}089) - 559^2} = \frac{4{,}095.3}{18{,}053} = 0.2268$$

$$a = \bar{y} - b\bar{x} = \frac{164.1}{6} - 0.2268 \times \frac{559}{6} = 6.2198$$

Rounding to two decimal places gives the regression line of sales of L against floor space:

$$y = 6.22 + 0.23x$$

(b) Similarly, the regression line of sales of M against floor space can be shown to be:

$$y = -39.05 + 0.87x$$

In part (a), $b = 0.23$ means that if the floor space were increased by one square metre, sales would increase by an average of 0.23 ('000 units), that is, by 230 units. The a-value of 6.22 gives the absurd result (due to extrapolation) that an outlet with zero floor space would on average sell 6,220 units.

In part (b), the corresponding results are that sales can be expected to increase by 870 units for each extra square metre of floor space and the nonsensical sales of −39,050 units would occur if an outlet had zero size.

Forecasting

Once the equation of the regression line has been computed, it is a relatively straightforward process to obtain forecasts.

Example

In the situation of Example 7.8.1, forecast the number of sales that would be expected next year in regions that employed (a) 14 salespersons; and (b) 25 salespersons.

Solution

As we have the 'best' line representing the dependence of sales on the number of salespersons we shall use it for the forecasts. The values could be read off the line drawn on the scattergraph, but it is more accurate to use the equation of the line.

(a) The regression line is $y = 111.05 + 10.12x$ so, when $x = 14$:

$$y = 111.05 + 10.12 \times 14 = 252.73$$

Rounding this to a whole number, we are forecasting that 253 units will be sold in a region employing 14 sallespersons.

(b) Substituting $x = 25$ into the formula:

$$y = 111.05 + 10.12 \times 25 = 364.05$$

Hence the forecast is sales of 364 units in a region employing 25 salespersons.

Have you noticed that x = 25 is well outside the range of the data? What does this tell you about the reliability of the estimated 364? If you are in any doubt about its unreliability, look back to the absurd consequences of putting x = 0 in Example 7.9.1. We give one more example to illustrate the complete process of forecasting from paired samples.

Which variable to denote by y

When calculating the correlation coefficient it does not matter which variable you call x and which y, as the result will be the same either way. However, the regression equation and subsequent forecasts will be totally changed if you change the designation of the variables. This is because the regression line only minimises the sum of squares of the y -errors, and this is only equivalent to minimising x -errors in the case of perfect correlation, that is, when r= - 1. It is therefore essential that you stop and think about which variable to call y at the start of any exercise on regression.

Variable y is the dependent variable and sometimes it is very clear which that is. However, there are occasions when the dependency could perhaps work either way, in which case the following may be of assistance:

- if you wish to forecast a particular variable, that variable must be denoted by y;
- if the question asks for the regression of a first variable on or against a second variable, the first variable is denoted by y . This is nothing to do

with the order in which the variables are tabulated in the question, which could easily have x first.

For example:

- in an investigation of the downwards trend of sales over time, the independent variable x = time and the dependent variable y = sales;
- in an investigation of the fall in cinema takings as sales of videos increase, possibly the decline in the cinema mirrors the increased use of videos (in which case cinema sales = y) , but perhaps it is the other way round. Perhaps cinemas were closed, in a property boom say, and people buy videos because there is no longer a convenient local cinema (in which case video sales = y). However, if the question referred to the regression of cinema sales against video sales or asked for an estimate of cinema sales for a known level of video sales then the doubt would be removed and cinema sales would have to be denoted by y .

Judging the validity of forecasts

When we have made forecasts, obvious questions to be asked are 'how accurate are they?' and 'what validity do they have?'

Such queries can be addressed in a number of ways.

The importance of using the correlation coefficient as a check on the validity of the assumption of causality has already been stressed. In addition, you should bear in mind the caveats mentioned in earlier parts of the chapter. In particular, is there a hidden third variable in the problem? Thus, in Examples, sales might not depend on the number of salespersons at all, but on the size of the region, as we mentioned when first discussing this problem. If this is the case, then simply increasing the number of salespersons within a region would not in itself increase sales. Even if this is not the case, have we got the causation the right way round? In Example it might be that, as profits increase, the company feels able to spend more on advertising, so that advertising expenditure depends on profits, contrary to the implicit assumption we made when forecasting profits. If this is the case, increasing the advertising would not necessarily increase profits. Before leaving the correlation coefficient, we mention another, closely related, measure, the coefficient of determination, r^2 . The value of this measure, when expressed as a percentage, shows the percentage of variations in the variables that can be explained by the regression analysis. The remaining variation is due to factors omitted from the analysis.

Example

Evaluate the coefficients of determination for the situations in (a) Example 1 and (b) Example2, and interpret their values.

Solution

(a) We have seen that $r = 0.948$, so $r^2 = 0.948^2 = 0.899$

Hence 89.9 per cent of the variations in sales from one region to the next can be explained by the corresponding differences in the number of salespersons. Only about 10 per cent of the differences in regional sales appear to be due to factors other than staffing levels.

(b) From $r = 0.936$ we get $r^2 = 0.936^2 = 0.876$

Thus, 87.6 per cent of the variations in profits from one year to the next can be explained by the corresponding variations in advertising expenditure, leaving a surprisingly low 12.4 per cent apparently due to other factors.

Consider now the two forecasts made in Example 7.10.1. The second one is distinctly different from the first, in that we have taken the regression line far beyond the upper data point (x = 18 salespersons) to twenty-five salespersons. The forecast is an extrapolation beyond the range of the data. This is an uncertain step to take, as the sales within a region at a certain time must have a ceiling: there must come a point where extra salespersons will generate no further sales. If this point has been passed with twenty-five salespersons, then our forecast will be an overestimate. The first case, by contrast, is an interpolation within the range of the data, and so can be considered more valid. In the same way, the profit forecast of Exampl2 is a slight extrapolation and so should be treated with some caution.

Extreme cases of extrapolation have already been seen when interpreting values of the coefficient a in earlier regression equations. In doing this, we are effectively extrapolating to the x -value of zero, and so we should not be surprised if the result seems implausible. The approach we have adopted is, of course, a considerable simplification of reality. Profits, sales, and so on, depend on a number of factors, not all of them quantifiable, whereas we have assumed here that they depend on just one other quantitative variable. We have studied only simple regression.

There is an extension to the topic, known as multiple regression, that enables a variable to be forecast in terms of any number of other variables. This is beyond the scope of this text.

All the forecasts made in this chapter have been for ‘next year’ , whereas the data comes, of course, from the past. There is, therefore, an implicit assumption that conditions that obtained in the past still obtain now and, more importantly, will continue to obtain during the period of the forecast. There is no mathematical way of checking that this is so, but the forecaster will have qualitative knowledge of the particular company and its market, and so will be able to form a judgement. If, for example, a new company was known to be making a big push in the market of the company in Example1, you might doubt the forecast of next year’s profit figures.

In conclusion, this section has looked at a number of considerations that should be borne in mind when judging the validity of a regression-based forecast. We shall summarise these in the next section.

Example

Comment on the likely reliability of your forecasts in Example.3.

Solution

The high values of the coefficients of variation suggest reliable forecasts, with the higher value for sales of L indicating the more reliable forecasts here. Forecasts for outlet G are interpolations and so are more reliable than those for outlet H, which are extrapolations. Finally, there is the possibility of other factors that might affect sales that are not covered in the analysis (advertising budgets at the various outlets, their relatively advantageous or disadvantageous locations, etc.).

REVIEW QUESTIONS

An ice-cream supplier has recorded some sales data that he believes shows a relationship between temperature and sales. The results shown below are for ten sample days in the summer:

Temperature (°C) x	Cartons sold y
13	10
16	11
17	14
19	15
20	16
21	19
23	24
26	25
27	26
28	27

I. Using the intermediate totals given below, calculate the coeffi cient of correlation giving your answer correct to two d.p.

$\Sigma x = 210 \quad \Sigma y = 210$
$\Sigma x^2 = 4{,}634 \quad \Sigma y^2 = 3{,}865$
$\Sigma xy = 4{,}208$

II. If the correlation coefficient was 0.95 calculate the coefficient of determination, giving your answer to the nearest whole number.

III. If the correlation coefficient was 0.95, which of the following statements would be correct?

IV. The positive sign tells us that there is a strong relationship between temperature and sales.

(B) The positive sign tells us that as temperature rises, so do sales.

(C) The value of the correlation coefficient tells us that there is a strong linear relationship between temperature and sales.

(D) The value of the correlation coefficient tells us that for each increase of 1 degree in temperature, sales increase by 0.95 cartons.

(E) The value of the correlation coefficient tells us that for each decrease of 1 degree in the temperature, sales decrease by 5 per cent.

(F) The value of correlation coefficient tells us that high temperatures cause high sales.

V. If the coefficient of determination was 85 per cent, which of the following statements would be correct?

(A) When temperature increases by 1°C, sales increase by 85 per cent.

(B) When temperature increases by 1°C, sales increase by 15 per cent.

(C) On 85 per cent of days it is possible to accurately predict sales if an accurate prediction of temperature exists.

(D) 85 per cent of the changes in sales from one day to the next can be explained by corresponding changes in temperature.

VI. The following graph displays the data. What type of graph is it?

(A) Scattergram
(B) Histogram
(C) Pictogram
(D) Ogive

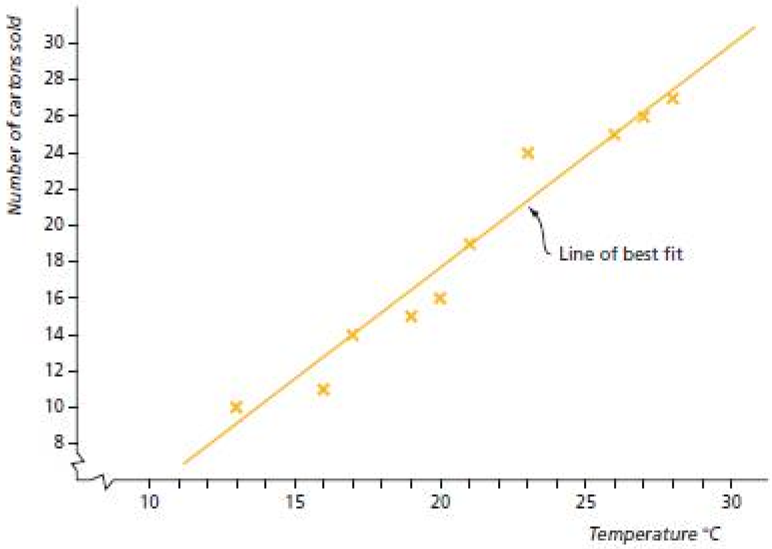

VII. (A) A freehand line of best fit has been fitted to the graph of the data. Estimate the likely sales when the temperature is 15°C, giving your answer to the nearest whole number.

(B) Estimate the likely sales when the temperature is 30°C, giving your answer to the nearest whole number.

VIII. Which of the following statements about the reliability of the estimates made in are correct, assuming that the correlation is 0.95.

(A) The estimate for a temperature of 15°C should be reliable because it involves interpolation.

(B) Both estimates are less reliable than they otherwise would be because the sample is small.

(C) The estimate for 30°C should be reliable because it involves extrapolation.

(D) Both estimates are more reliable than they would otherwise be because the correlation is high.

IX. Using Excel, enter the data above and draw a scatter diagram to examine the relationship between the temperature and the number of cartons of ice cream sold.

2. A travel agency has kept records of the number of holidays booked and the number of complaints received over the past ten years. The data is as follows:

Year	1	2	3	4	5	6	7	8	9	10
Number of holidays booked	246	192	221	385	416	279	343	582	610	674
Number of complaints received	94	80	106	183	225	162	191	252	291	310

The agency suspects there is a relationship between the number of bookings and the volume of complaints and wishes to have some method of estimating the number of complaints, given the volume of bookings.

I. Denoting number of holidays by X and number of complaints by Y, the following summations are given:

$$\Sigma X = 3{,}948,\ \Sigma Y = 1{,}894,\ \Sigma X^2 = 1{,}828{,}092,\ \Sigma Y^2 = 417{,}96,\ \Sigma XY = 869{,}790.$$

Calculate the value of the regression coefficient ' b ' , giving your answer correct to three d.p.

II. If the value of b is taken to be 0.4 calculate the value of the regression coefficient 'a ' , giving your answer correct to two d.p.

III. If the regression equation was y 31 ———————————— 0.4 x forecast the likely number of complaints if 750 holidays are booked, giving your answer to the nearest whole number.

IV. Which of the following methods could be used to check whether there is in fact a linear relationship between the variables.
   a. Scatter diagram
   b. Times series analysis
   c. Coefficient of variation
   d. Regression analysis
   e. Correlation coefficient

V. Which of the following comments about the likely reliability of the estimate of

VI. Complaints arising from 750 holidays is/are correct?
   a. The estimate is likely to be reliable because the value of ‘ a ’ is positive.
   b. The estimate is likely to be reliable because it lies outside the range of the data.
   c. The estimate is likely to be unreliable because the sample is small.
   d. The estimate is not likely to be reliable because the value of ‘ b ’ is not close to 1.
   e. The estimate is likely to be unreliable because it was obtained by extrapolation.

VII. Using Excel, enter the data above and produce a chart to show the relationship between the number of holidays booked and the number of complaints received, and then plot the least-squared line.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-8 TIME SERIES

CONTENTS

INTRODUCTION

There are many situations in which there are no plausible or available independent variables from which a dependent variable can be forecast. In such cases, approaches alternative to regression have to be adopted. One of these consists of using past values of the variable to be forecast, a so-called time series, and looking for patterns in them. These patterns are then assumed to continue into the future, so that an extrapolative forecast is produced. The first task is thus to discuss the various patterns that time series data displays.

Components and models of time series

There are considered to be four components of variation in time series:

- the trend, T;
- the seasonal component, S;
- the cyclical component, C; and
- the residual (or irregular, or random) component, R.

The trend in a time series is the general, overall movement of the variable, with any sharp fluctuations largely smoothed out. It is often called the underlying trend, and any other components are considered to occur around this trend. There are a number of basic trend patterns that business variables tend to follow, as shown in Figure 8.1 . The simplest (Figure 8.1 (a)) is a linear trend, in which the variable is basically growing (or declining) at a steady rate. A logistic trend (Figure 8.1 (b)) is typically followed by the sales figures of a product after its introduction: the level plateau is the market saturation figure that the sales eventually reach. A compound interest (or exponential) trend (Figure 8.1 (c)), as the name suggests, is a relatively steeply rising curve followed by variables whose values are compounded on earlier values: for instance, investments subject to compound interest.

The seasonal component accounts for the regular variations that certain variables show at various times of the year. Thus, a newly formed ice-cream manufacturing company may have sales figures showing a rising trend. Around that, however, the sales will tend to have peaks in the summer months and troughs in the winter months. These peaks and troughs around the trend are explained by the seasonal component. In general, if a variable is recorded weekly, monthly or quarterly, it will tend to display seasonal variations, whereas data recorded annually will not.

The cyclical component explains much longer-term variations caused by business cycles. For instance, when a country's economy is in a slump, most business variables will be depressed in value, whereas when a general upturn occurs, variables such as sales and p rofits will tend to rise. These cyclical variations cover periods of many years and so have little effect in the short term.

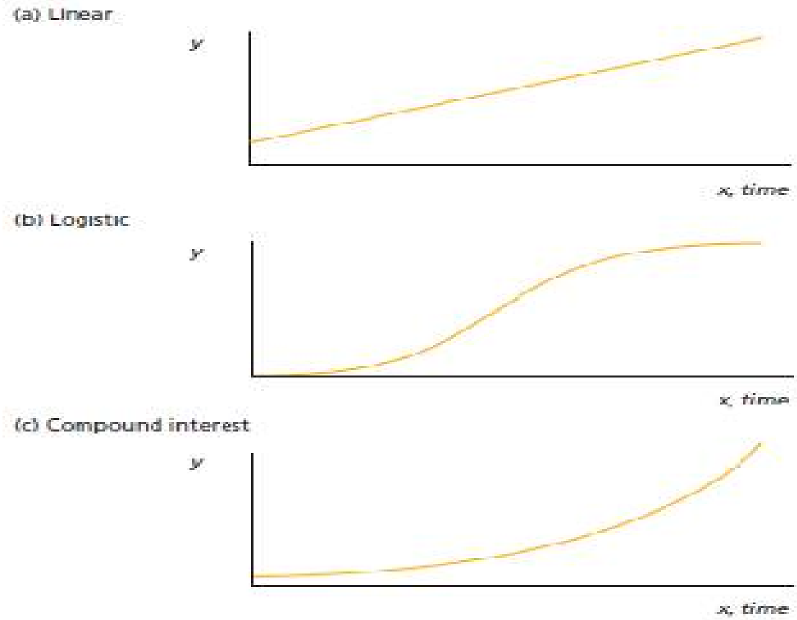

The residual component is that part of a variable that cannot be explained by the factors mentioned above. It is caused by random fluctuations and unpredictable or freak events, such as a major fire in a production plant. If the first three components are explaining the variable's behaviour well, then, subject to rare accidents, the irregular component will have little effect.

The four components of variation are assumed to combine to produce the variable in one of two ways: thus we have two mathematical models of the variable. In the first case there is the additive model, in which the components are assumed to add together to give the variable, Y :

Y =T+S+C+R

The second, multiplicative, model considers the components as multiplying to give Y :

Y=TX SXCXR

Thus, under the additive model, a monthly sales figure of £ 21,109 might be explained as follows:

- the trend might be £ 20,000;
- the seasonal factor: £ 1,500 (the month in question is a good one for sales, expected to be £ 1,500 over the trend);
- the cyclical factor: £ 800 (a general business slump is being experienced, expected to depress sales by £ 800 per month); and
- the residual factor: £ 409 (due to unpredictable random fluctuations). The model gives:

$$Y = T + S + C + R$$
$$21{,}109 = 20{,}000 + 1{,}500 + (-800) + 409$$

Notes

The multiplicative model might explain the same sales figures in a similar way:

- trend: £ 20,000;
- seasonal factor: 1.10 (a good month for sales, expected to be 10 per cent above the trend);
- cyclical factor: 0.95 (a business slump, expected to cause a 5 per cent reduction in sales); and
- residual factor: 1.01 (random fluctuations of = 1 per cent).

The model gives:

$$Y = T \times S \times C \times R$$
$$21{,}109 = 20{,}000 \times 1.10 \times 0.95 \times 1.01$$

It will be noted that, in the additive model, all components are in the same units as the original variable (£ in the above example). In the multiplicative model, the trend is in the same units as the variable and the other three components are just multiplying factors.

Forecasting linear trends

There are many ways of forecasting time series variables. To give a flavour of extrapolative forecasting we shall concentrate here on just one. The method consists of forecasting each component separately, and then combining them through one of the models to form a forecast of the variable itself. We begin with the trend, initially by assuming the simplest case of linear trends. In this case, there is no need for any new theory since we can find the trend as a linear regression line.

Example

The following table gives the quarterly sales figures of a small company over the last 3 years. Forecast the next four values of the trend in the series.

	Time period	Sales £'000
1992	quarter 1 (t = 1)	42
	quarter 2 (t = 2)	41
	quarter 3 (t = 3)	52
	quarter 4 (t = 4)	39
1993	quarter 1 (t = 5)	45
	quarter 2 (t = 6)	48
	quarter 3 (t = 7)	61
	quarter 4 (t = 8)	46
1994	quarter 1 (t = 9)	52
	quarter 2 (t = 10)	51
	quarter 3 (t = 11)	60
	quarter 4 (t = 12)	46

Solution

The graph of these data, the *time series graph*, is shown in Figure 8.2. This shows that the company's sales are following an upward trend, of a more or less linear shape, and that there is a definite seasonal pattern: each third quarter is a peak and each fourth quarter is a trough. The approach and model being used here are therefore appropriate.

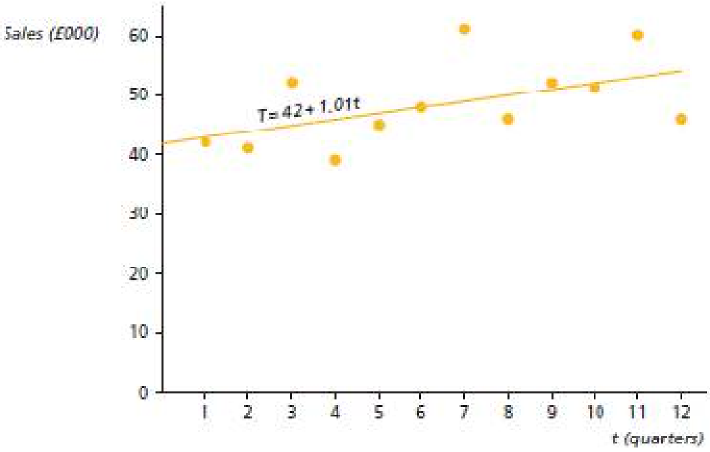

Times series graph and trend line (Example)

It will be noted that the twelve quarters for which we have data have been numbered from one to twelve, for ease of reference and to facilitate the computation of the regression line. It is left as an exercise for you to verify that this has equation:

T = 42.0+ 1.01t

where T is the assumed linear trend in sales (£'000) and t is the number of the quarter (1992, quarter 1: t = 1, and so on). This line has been superimposed on the graph in Figure. The process of calculating the trend, whether by regression or by moving averages (see later), is often described as 'smoothing the data'. As you can see from the above graph, the original ups and downs of the data have been smoothed away.

Forecasting seasonal components

Up to now, we have not had to concern ourselves with the choice of model. Since the nature of the seasonal component is so different in the two models, we now have to make a choice. The multiplicative model is usually considered the better, because it ensures that seasonal variations are assumed to be a constant proportion of the sales. The additive model, in contrast, assumes that the seasonal variations are a constant amount, and thus would constitute a diminishing part of, say, an increasing sales trend. Because there is generally no reason to believe that seasonality does become a less important factor, the multiplicative model is adopted more frequently, as demonstrated here.

The arithmetic involved in computing seasonal components is somewhat tedious but essentially simple. Assuming a very simple model in which there are no cyclical or residual variations:

$$\text{Actual value, } Y = T \times S$$

$$\text{so } S = \frac{Y}{T}$$

The seasonal component, S, is therefore found as the ratio of the actual values to the trend, averaged over all available data (so as to use as much information as possible). For forecasting purposes, the same degree of seasonality is assumed to continue into the future, and so the historical seasonal components are simply projected unaltered into the future.

Example

Calculate the seasonal components from the sales data and trend of Example 8.3.1.

Solution

The first, tedious step is to calculate the ratio of sales trend for each of the twelve quarters given. We show the first and last here, leaving the intermediate ten calculations as exercises:

$t = 1$: $T = 42.0 + (1.01 \times 1) = 43.01$

$$S = \frac{Y}{T} = \frac{42}{43.01} = 0.9765$$

$\vdots \qquad \vdots$

$t = 12$: $T = 42.0 + (1.01 \times 12) = 54.12$

$$S = \frac{Y}{T} = \frac{46}{54.12} = 0.8500$$

The complete set of ratios, arranged by quarter, is:

	Quarter 1	*Quarter 2*	*Quarter 3*	*Quarter 4*
1992	0.9765	0.9314	1.1548	0.8471
1993	0.9564	0.9988	1.2431	0.8185
1994	1.0178	0.9789	1.1297	0.8500
Total	2.9507	2.9091	3.5276	2.5156
Mean	0.9836	0.9697	1.1759	0.8719

When arranged like this the averaging process for each quarter is facilitated. The resulting values constitute the mean seasonal component for each quarter from the given data: they show that, on average in the past, quarter 1 sales have been 98 per cent (approximately) of the trend, quarter 2 sales 97 per cent of the trend, and so on. These values are now adopted as the required forecast seasonal components (denoted S). In this case the forecasts for the four quarters of 1995 are thus:

0.9836, 0.9697, 1.1759 and 0.8719, respectively.

As the four seasonal components under this model should, on average, cancel out over a year, an extra step is often taken here, to ensure they add up to 4 (an average of 1 each). The arithmetic is straightforward:

$$\text{Total} = 0.9836 + 0.9697 + 1.1759 + 0.8719 = 4.0011$$

To reduce this to 4, we will have to subtract from each one:

$$\frac{4.0011 - 4}{4} = 0.0003 \text{ (to four d.p.)}$$

This gives the seasonal components as:

0.9833, 0.9694, 1.1756 and 0.8716, respectively.

In this instance, the adjustment has had scarcely any effect and so can be ignored. In fact, the original data seems to have been rounded to three s.f. so giving the seasonal components to four d.p. cannot really be justified. They would be better rounded to 0.98, 0.97, 1.18 and 0.87

We have used arithmetic averaging to find the average seasonal variation and to adjust the averages so that our estimated components add to 4. An alternative method that is more mathematically 'correct' is to use geometric means and to adjust the average ratios so they multiply to 1. However, in practice it makes virtually no difference and the arithmetic mean is easier.

Producing the final forecast

We must now consider the final two components of variation. Isolating the cyclical component of time series has proved to be a controversial area in economics and statistics. There is no consensus on an approach to the problem. Also, as we have already mentioned, cyclical variations have little effect in the short term. For these reasons, we shall omit the factor C from this first treatment.

The residual component is by nature unpredictable. The best that we can do is to hope that any random fluctuations are small and that no freak events occur, so that the factor R has no overall effect.

For a component to be omitted or to have no effect, it must have the value 1 in the multiplicative model, since multiplying anything by 1 leaves it unchanged. We have thus simplified our model, for the purposes of forecasting, to

$$\hat{Y} = \hat{T} \times \hat{S}$$

Example

In the example under discussion here, forecast the sales during 1995.

Solution

We have already found values for $\hat{T}$ and $\hat{S}$, and so it is now a matter of pulling these values together to find $\hat{Y}$:

1995 quarter 1: $\hat{Y} = \hat{T} \times \hat{S} = 55.1 \times 0.9833 = 54.18$

1995 quarter 2: $\hat{Y} = 56.1 \times 0.9694 = 54.38$

1995 quarter 3: $\hat{Y} = 57.2 \times 1.1756 = 67.24$

1995 quarter 4: $\hat{Y} = 58.2 \times 0.8716 = 50.73$

The forecast sales for the four quarters of 1995 are thus £54,000, £54,000, £67,000 and £51,000, respectively (to the nearest £'000).

Seasonal adjustment

Before proceeding we digress slightly to look at a closely related topic, seasonal adjustment. This is important, because we are often presented with a single figure for weekly revenue, monthly profit, or whatever, and it is difficult to make judgements without some idea of the extent to which the figure has been distorted by seasonal factors and consequently does not give a good indication of the trend. One approach is to deseasonalise or remove the seasonal effects from the figure. In the multiplicative model, in which the factor S multiplies with all the other components, seasonal adjustment consists of dividing by S. In other words, from

$$Y = T \times S$$

we estimate:

$$T = \frac{Y}{S}$$

Effectively, the seasonally adjusted figure is an estimate of the trend.

Example

The company of Example 8.3.1 reports sales of £50,000 during the fourth quarter of a certain year. Seasonally adjust this figure.

Solution

We saw earlier that the seasonal component for the fourth quarter in this series is 0.8716. Dividing by this:

$$\frac{50,000}{0.8716} = 57,365$$

we see that the seasonally adjusted sales for the quarter in question are £57,365.

Moving average trends

The above approach is based on an assumption of a linear trend. Although this may appear plausible or 'appropriate' , there are many occasions where such an assumption might not be made. An alternative approach that does not depend on linearity, but that also has some relative disadvantages discussed later, involves using moving averages as the trend. The arithmetic involved in this approach is still voluminous but essentially simpler than that of regression analysis, and can just as easily be computerised. To illustrate the method, we continue to look at the example discussed above.

Example

In the example under discussion, compute the trend as a centred four-point moving average.

Solution

In the table below, the 'four-quarterly total' column is simply the sum of each set of four consecutive quarterly sales figures. The first is thus:

$$42 + 41 + 52 + 39 = 174$$

The second is:

$$41 + 52 + 39 + 45 = 177$$

and so on. The important question is where these totals should go. As they are to represent the four-quarterly period, the usual convention is to place them in the middle of the period, that is, between Q2 and Q3 for the first one, between Q3 and Q4 for the second, and so on. You will find that the table looks neater and is easier to read if you leave an empty line between the quarters, but there is often insufficient space to do this.

A small problem now arises because we wish each value of the trend to be eventually associated with a specific quarter. To overcome this, the figures are 'centred' – that is, each pair of values is added to give the 'centred eight-quarterly totals':

174 + 177 = 351 opposite 1992 Q3
177 + 184 = 361 opposite 1992 Q4 ... and so on

Dividing by 8 now gives the trend values shown:

		Sales (£'000)	Four-quarterly total	Centred eight-quarterly total	Moving averge (T)
1992	Q1	42			
	Q2	41			
	Q3	52	174	351	43.88
	Q4	39	177	361	45.13
1993	Q1	45	184	377	47.13
	Q2	48	193	393	49.13
	Q3	61	200	407	50.88
	Q4	46	207	417	52.13
1994	Q1	52	210	419	52.38
	Q2	51	209	418	52.25
	Q3	60	209		
	Q4	46			

We now complete the process of forecasting from these trend values. There are no new techniques involved, as the steps of Examples 8.4.1 and 8.5.1 are being followed with new values for $\hat{T}$.

Other types of data

Before moving on, it will be noted that we have centred quarterly data here. In order to deal with weekly data, for example, a centred 104-point moving average would be needed for the trend (and, incidentally, there would be fifty-two seasonal components, one for each week). Monthly data would lead to a 24-point moving average trend and twelve seasonal components.

The moving average approach can also be used for trends from annual data. If the data has a clear cycle of highs and lows spanning, say, 5 years, then non-centred five-point moving averages would be used. If there is no cyclical pattern the choice is arbitrary, but three or five-point moving averages are often used to smooth the data because no centering is needed when an odd number of figures are averaged. Cyclical components are estimated by averaging Y = T values in the same way as seasonal components.

Example

Using the data of Example 8.3.2, calculate the trend for the sales of article B as a centred four-point moving average.

Solution

Year	Quarter	Sales (Y)	Four-point moving total	Eight-point moving total	Four-point moving ave. trend (T)
1993	1	24.8			
	2	36.3			
	3	38.1	146.7	299.8	37.4750
	4	47.5	153.1	311.9	38.9875
1994	1	31.2	158.8	322.9	40.3625
	2	42.0	164.1	336.6	42.0750
	3	43.4	172.5	353.8	44.2250
	4	55.9	181.3	369.4	46.1750
1995	1	40.0	188.1	386.8	48.3500
	2	48.8	198.7	410.6	51.3250
	3	54.0	211.9	438.5	54.8125
	4	69.1	226.6	462.2	57.7750
1996	1	54.7	235.6	477.5	59.6875
	2	57.8	241.9	483.6	60.4500
	3	60.3	241.7		
	4	68.9			

Judging the validity of forecasts

As in the preceding chapter, we now have to consider how valid are these and other extrapolative forecasts. First of all, as the name implies, they are extrapolations, and so there is the possibility of error, as discussed earlier. In particular, you should monitor background circumstances to detect any changes that might invalidate the assumption that these are constant.

Further, assumptions made about the trend can be critical. The adoption of a linear trend may appear plausible but it is sometimes difficult to check its validity. For example, the moving average trend shown in Figure 8.3 may indicate that the sales are following a logistic form (see Figure 8.1 (b)), and that the linear regression approach may be extrapolating the middle portion of the graph beyond the 'market saturation' plateau. Each such successive step into the future becomes increasingly less reliable. This 'plateauing out' is reflected in the forecasts of Example. The moving average approach is not without its problems either. The method of calculating T meant that there were two existing quarters (1994 Q3 and Q4) through which any trend extrapolations had to extend before getting into the future. There was no guidance as to where the trend line should be extended: it had to be done 'by eye' , using 'judgement ' , and so the additional two quarters cast further doubt on the reliability of the trend forecasts.

There are refinements to these basic methods that can remove the necessity to make such assumptions or to assert such judgements on the trend, and can deal with non constant seasonal components. These are beyond the scope of this text.

The methods of this chapter, and any amendments to them, depend on the assumptions that a time series has a certain number of components of variation, and that these combine in a certain way ('the model'). One way of checking on these assumptions is to assess the values of the

Notes

residuals from past data. To do this, we reintroduce R into our model:

$$Y = T \times S \times R$$

$$\therefore R = \frac{Y}{T \times S}$$

Thus, in 1992 quarter 1 of the time series under discussion here (linear trend):

$Y = 42$; $T = 43.01$ (evaluated in Example 8.4.1); $S = 0.9836$

so that:

$$\therefore R = \frac{Y}{T \times S} = \frac{42}{43.01 \times 0.9836} = 0.9928$$

Alternatively, for the other approach to the trend, using the figures of Examples 8.7.1 and 8.7.3:

$$1992\ Q3: R = \frac{Y}{T \times S} = \frac{52}{43.88 \times 1.1920} = 0.9942$$

Proceeding in this way, all past values of the residual component can be found (you might complete the calculations of these values for practice):

		R, linear regression trend	*R, moving average trend*
1992	Q1	0.9928	
	Q2	0.9605	
	Q3	0.9820	0.9942
	Q4	0.9715	0.9895
1993	Q1	0.9724	0.9805
	Q2	1.0300	1.0004
	Q3	1.0572	1.0058
	Q4	1.0535	1.0104
1994	Q1	1.0348	1.0195
	Q2	1.0095	0.9995
	Q3	0.9607	
	Q4	0.9748	

Ideally, the residuals should be having little effect and so should be close to 1. All the above values are fairly near to 1, which gives some support to the validity of the forecasts. Direct comparisons of the eight quarters possible, however, shows that the right-hand column is always closer to 1 than the left. The moving average approach therefore appears more valid, and this, in turn, possibly reflects the fact that the moving averages have dealt with a 'plateauing out' of sales, whereas the linear regression has extrapolated beyond it. Against this must be set the fact that we had to 'guess' where the moving average trend goes next, whereas it is known where a regression line goes.

Further, if we had more quarterly sales figures we could inspect the R - values for patterns: as an assumedly random component, there should not be any. If, for example, the values gradually moved away from 1, the model would be getting progressively less reliable, so casting doubts on any forecasts from it. Similarly, if there was a seasonal pattern in R, this would cast doubt on the underlying assumption of constant seasonality.

Finally, you will have noticed that there is a great amount of arithmetic involved in producing the forecasts of this chapter. It is therefore highly advantageous to use one of the many available computer packages that deal with such extrapolative models.

Computations involving the additive model

Although it has been stated that the multiplicative model is the more often applicable, the additive model may occasionally be used, and so we give an example of this latter model. As before, the computation of the trend does not depend on the model chosen, and so any form of trend can be applied.

REVIEW QUESTIONS

1. The managers of a company have observed recent demand patterns of a particular product line in units. The original data, which has been partially analysed, is as follows:

Year	Quarter	Data	Sum of fours	Sum of twos
1993	2	31		
	3	18	94	190
	4	20	96	193
1994	1	25	97	195
	2	33	98	197
	3	19	99	198
	4	21	99	198
1995	1	26	99	199
	2	33	100	201
	3	19	101	
	4	22		
1996	1	27		

You have been commissioned to undertake the following analyses and to provide appropriate explanations. (Work to three d.p.)

I. In the following table, fi nd the missing values of the underlying four-quarterly moving average trend.

Year	Quarter	Sum of twos	Moving average
1993	4	190	A
1994	1	193	B
	2	195	24.375
	3	197	24.625
	4	198	24.750
1995	1	198	24.750
	2	199	24.875
	3	201	C

II. Calculate the seasonally adjusted demand (to three d.p.) for the four quarters of 1994 based on the multiplicative model if the seasonal factors are as follows:
Quarter 1 1.045
Quarter 2 1.343
Quarter 3 0.765
Quarter 4 0.847

III. Which of the following statements about seasonal adjustment is/are correct?
(A) Seasonally adjusted data has had the seasonal variations removed from it. Correct/incorrect
(B) Seasonally adjusted data has had the seasonal variations included in it. Correct/incorrect
(C) Seasonal adjustment is the process by which seasonal components are adjusted so that they add to zero.
Correct/incorrect
(D) Seasonal adjustment is the process by which estimates of the trend can easily be obtained.
Correct/incorrect

IV. If the seasonally adjusted values are increasing, which of the following would you deduce?
(A) The trend is upwards.
(B)The trend is downwards.

(C) No deductions about the trend are possible from the information given.

(D) Seasonal variability is increasing.

(E) Seasonal variability is decreasing.

(F) No deductions are possible about seasonal variability.

V. If A denotes the actual value, T the trend and S the seasonal component, write down the formula for the seasonally adjusted value if an additive model is being used.

2. You are assisting the management accountant with sales forecasts of two brands – Y and Z – for the next three quarters of 1993. Brand Y has a steady, increasing trend in sales of 2 per cent a quarter and Brand Z a steadily falling trend in sales of 3 per cent a quarter. Both brands are subject to the same seasonal variations, as follows:

Quarter	*Q1*	*Q2*	*Q3*	*Q4*
Seasonality	−30%	0	−30%	+60%

The last four quarter's unit sales are shown below:

	1992 Q2	*1992 Q3*	*1992 Q4*	*1993 Q1*
Brand Y	331	237	552	246
Brand Z	873	593	1,314	558

I. Which of the following statements about the seasonal variations is/are correct?

A. Actual sales are on average 30 per cent below the trend in the third quarter.

B. Actual sales in the fi rst and third quarters are identical on average.

C. Average sales in the second quarter are zero.

D. Actual sales in the fourth quarter are on average 60 per cent above the trend.

E. Actual sales in the fi rst quarter are 1.3 times the trend.

F. Actual sales in the fourth quarter are 1.6 times the trend.

II. Seasonally adjust the sales fi gures for 1993 Q1, giving your answers to one d.p.

III. Forecast the trend for brand Y for 1993 Q4, giving your answer to one d.p.

IV. If the trend forecast in 3.3 was 370, forecast the actual sales of brand Y for 1993 Q4, giving your answer to the nearest whole number.

V. Forecast the trend for brand Z for 1993 Q3, giving your answer to one d.p.

VI. If the trend forecast in 3.5 was 770, forecast the actual sales of brand Z for 1993 Q3, giving your answer to the nearest whole number.

VII. Which of the following are assumptions on the basis of which time series forecasts are made?

A. That there will be no seasonal variation.

B. That the trend will not go up or down.

C. That there will be no change in the existing seasonal pattern of variability.

D. That the model being used fi ts the data.

E. That there will be no unforeseen events.

3. The quarterly sales of a product are monitored by a multiplicative time series model. The trend in sales is described by

Y =100 + 5X

where Y denotes sales volume and X denotes the quarterly time period. The trend in sales for the most recent quarter (fi rst quarter 1991, when X = 20) was 200 units. The average seasonal variations for the product are as follows

Quarter	First	Second	Third	Fourth
Seasonal effect	0	−20%	+40%	−20%

The price of a unit was £ 1,000 during the fi rst quarter of 1991. This price is revised every quarter to allow for infl ation, which is running at 2 per cent a quarter.

I. Forecast the trend in the number of units sold for the remaining three quarters of 1991.

II. Forecast the actual number of units sold (to the nearest whole number) for the remaining three quarters of 1991.

III. Forecast the price per unit for the remaining quarters of 1991, giving your answers correct to two d.p.

IV. If the prior forecasts were as follows, forecast the sales revenue for the remaining quarters of 1991, giving your answers to the nearest £.

Quarter of 1991	Forecasts Numbers sold	Price per unit (£)
2	150	1010
3	300	1030
4	170	1050

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-9 PROBABILITY

CONTENTS

INTRODUCTION

Most people have some intuitive conception of uncertainty or chance, as the following typical statements illustrate:

- 'On past evidence, there seems to be a 50/50 chance of this project succeeding.'
- 'I reckon that, if we stay on this course, we will have only a one in ten chance of making a profit next year.'
- 'The consultants' report says that our project launch has a 60 per cent chance of success.' Each of the above sentences contains a term attempting to quantify the degree of uncertainty in a business situation. In this chapter we introduce the formal study of such quantify action, or probability, initially looking at several different approaches to the subject. More detail of uncertainty and risk is given in Section 9.14.

Definitions of probability

Probability is a branch of mathematics that calculates the likelihood of a given event occurring, and is often expressed as a number between 0 and 1. It can also be expressed as a proportion or as a percentage. If an event has a probability of 1 it can be considered a certainty: for example, the probability of spinning a coin resulting in either 'heads' or 'tails' is 1, because, assuming the coin lands flat, there are no other options. So, if an event has a probability of 0.5 it can be considered to have equal odds of occurring or not occurring: for example, the probability of spinning a coin resulting in 'heads' is 0.5, because the spin is equally likely to result in 'tails.' An event with a probability of 0 can be considered an impossibility: for example, the probability that the coin will land flat without either side facing up is 0, because either 'heads' or 'tails' must be facing up. In real life, very few events are ever given a probability of zero as there is always an unknown element in human affairs. As mentioned above one way of viewing/understanding a probability is as a proportion, as the following simple example will illustrate.

Example

An ordinary six-sided dice is rolled. What is the probability that it will show a number less than three?

Solution

Here it is possible to list all the possible equally likely outcomes of rolling a dice, namely the whole numbers from one to six inclusive:

1, 2, 3, 4, 5, 6

The outcomes that constitute the 'event' under consideration, that is, 'a number less than three' are:

1, 2

Hence the proportion of outcomes that constitute the event is 2/6 or 1/3, which is therefore the desired probability.

An ordinary six-sided dice is rolled. What is the probability that it will show a number less than three?

Solution

Here it is possible to list all the possible equally likely outcomes of rolling a dice, namely the whole numbers from one to six inclusive:

1, 2, 3, 4, 5, 6

The outcomes that constitute the 'event' under consideration, that is, 'a number less than three' are:

1, 2

Hence the proportion of outcomes that constitute the event is 2/6 or 1/3, which is therefore the desired probability.

Note that this answer agrees with the intuitive statements you might make about this situation, such as 'the chances are one in three' . In situations like this, where it is possible to compile a complete list of all the equally likely outcomes, we can define the probability of an event, denoted P (event), in a way that agrees with the above intuitive approach:

$$P(\text{event}) = \frac{\text{Total number of outcomes which constitute the event}}{\text{Total number of possible outcomes}}$$

This is known as exact probability because it involves having a complete list of all possible outcomes and counting the exact number that constitute the event. This definition, however, is not always practical for business purposes, as you can rarely state all the possible outcomes. To illustrate this, and to demonstrate a way of overcoming the problem, we consider the following.

Example

A quality controller wishes to specify the probability of a component failing within 1 year of installation. How might she proceed?

Solution

To find this probability from an exact approach would necessitate obtaining a list of the lifetimes of *all* the components, and counting those of less than 1 year. It is clearly impossible to keep such a detailed record of every component, after sale.

An alternative, feasible approach is to take a *sample* of components, rather than the whole population, and test them under working conditions, to see what proportion fail within one year. Probabilities produced in this way are known as *empirical* and are essentially approximations to the true, but unobtainable, exact probabilities. In this case, the quality controller may choose to sample 1,000 components. If she then finds that 16 fail within 1 year:

$$P(\text{component failing within 1 year}) = \frac{16}{1{,}000} \text{ or } 0.016$$

For this approximation to be valid, it is essential that the sample is representative. Further, for a more accurate approximation, a larger sample could be taken, provided that the time and money are available.

We make two comments before moving on. First of all, since we are defining probabilities as proportions, probabilities will lie in the range 0–1, with 0 denoting an impossibility and 1 denoting a certainty. Second, there are many practical instances in which a suitable sample is unavailable, so an empirical probability cannot be found. In such cases a subjective probability could be estimated, based on judgement and experience. Although such estimates are not entirely reliable, they can occasionally be useful, as we shall see later. The second quotation in the Introduction to this chapter is an example of the use of judgement to estimate a subjective probability.

Addition rules of probability

In principle, it is possible to find any probability by the methods discussed above. In practice, however, there are many complex cases that can be simplified by using the so-called rules of probability. We shall develop these via examples.

Example

According to personnel records, the 111 employees of an accountancy practice can be classified by their work-base (A, B or C) and by their professional qualifications thus:

	Office A	*Office B*	*Office C*	*Total*
Qualified	26	29	24	79
Not qualified	11	9	12	32
Total	37	38	36	111

What is the probability that a randomly selected employee will:

(a) work at office A or office B?
(b) work at office A or be professionally qualified or both?

Solution

(a) There are 37 people working in A and 38 in B, making 75 out of 111 who work in either A or B. Hence, *P*(A or B) = 75/111.

(b) Examining the table, we can apply our earlier rule: 37 are employed at office A, and 79 are qualified, making a total of 116. It is clear, however, that we have 'double counted' the 26 employees who both work at office A and are qualified. Subtracting this double-counted amount, we see that

$$116 - 26 = 90$$

employees have the desired property. Hence:

$$P(\text{employed at office A or professionally qualified}) = \frac{90}{111}$$

We can generalise from the above results. We have:

(a) $P(\text{A or B}) = \frac{75}{111} = \frac{37}{111} + \frac{38}{111} = P(\text{A}) + P(\text{B})$.

This is called the *special addition rule* of probability and it holds only when A and B cannot both be true. A and B are said to be *mutually exclusive* in that, if either is true, then the other is excluded. In other words, *P*(A and B) = 0.

(b) $P(\text{employed at office A or professionally qualified}) = \frac{90}{111} = \frac{37 + 79 - 26}{111}$

$$= \frac{37}{111} + \frac{79}{111} - \frac{26}{111}$$

= *P*(employed at office A) + *P*(professionally qualified) − *P*(office A and qualified)

This is an example of the *general addition law* of probability:

$$P(\text{X or Y}) = P(\text{X}) + P(\text{Y}) - P(\text{X and Y})$$

The last term in this law compensates, as we have seen, for double counting. If, however, there is no possibility of double counting – that is, if X and Y cannot occur together (i.e. *P*(X and Y) = 0) – then this term can be omitted and the law simplified to the 'special addition law':

$$P(\text{X or Y}) = P(\text{X}) + P(\text{Y})$$

These rules are given in your CIMA assessment in the forms:

$$P(A \cup B) = P(A) + P(B) \text{ and}$$
$$P(A \cup B) = P(A) + P(B) - P(A \cap B).$$

The probability of opposites

In the example we have been using, a person is either qualified or not qualified, with probabilities 79/111 and 32/111, respectively. Not surprisingly, P (Qualified) = P (Unqualified) = 111/111 = 1. It is certain that a person is either qualified or unqualified and 1 is the probability of certainty. In general:

1. If events A and B are mutually exclusive and together cover all possibilities then P (A) = P (B) = 1. This holds true for any number of such events.

2. P (A is not true) = 1 = P (A is true). This is in practice a remarkably useful rule. It often happens that it is easier to calculate the probability of the very opposite of the event you are interested in. It is certainly always worth thinking about if you are not sure how to proceed.

The multiplication rules of probability

Example

In the situation described in Example 9.3.1, what is the probability that a randomly selected employee will come from office B *and* not be qualified?

Solution

A reading from the table shows that nine of the 111 employees come under the required category. Hence:

$$P(\text{employed at office B and not qualified}) = \frac{9}{111}$$

Developing this, as above, to derive a general rule:

$$= \frac{38}{111} \times \frac{9}{38}$$
$$= P(\text{office B}) \times P(\text{not qualified, if from B})$$

This is an example of the general multiplicative law of probability:

$$P(X \text{ and } Y) = P(X) \cdot P(Y, \text{ if } X)$$

or

$$P(X \text{ and } Y) = P(X) \cdot P(Y|X)$$

In the latter form of this law we see the notation $P(Y|X)$ which is read as 'the probability of Y *if* (or *given*) X'. Such a probability is called a *conditional probability*. It is used because the fact that X occurs affects the probability that Y will occur. In the above, for example:

$$P(\text{not professionally quallified}) = \frac{32}{111}$$

Yet the value we must use in the calculation is

$$P(\text{not professionally qualified if from office B}) = \frac{9}{38}$$

On occasion, X and Y are statistically independent. That is, the fact that X occurs has no effect on the probability of Y occurring (and vice versa),

$$P(Y|X) = P(Y)$$

In this case, the rule can be simplified to the *special multiplication rule:*

$$P(X \text{ and } Y) = P(X) \cdot P(Y)$$

These rules are given in your CIMA assessment in the forms:

$$P(A \cap B) = P(A) \times P(B) \text{ and}$$
$$P(A \cap B) = P(A) \times P(B|A)$$

We now give two more worked examples to illustrate further the application of the above rules. They also serve to demonstrate typical applications of probability in finance and accountancy. The first is in the area of life assurance.

Before leaving this example, we point out two simplifications (and therefore two assumptions) we have made. The probabilities quoted in the question (0.69 and 0.51) are empirical, arising from the histories of many people in the past. To apply these values to the couple in the question required the assumption that the people are ‘typical’ of the sample from which the basic probabilities came. If, for example, either spouse had a dangerous occupation, this assumption would be invalid: in practice, actuaries would have data, and therefore empirical probabilities, to deal with such ‘untypical’ people.

Second, we have had to make the assumption that the life expectations of spouses are independent, which is probably not true.

It should be emphasised that these estimates are valid only if the test results are representative of the actual product performance and if this product does resemble the previous ‘similar items’ regarding return rates of faulty product (100 per cent) and faulty packaging (50 per cent).

Discrete probability distributions; expectations

We use the word discrete here in the same sense as in Chapter 3, namely to describe a variable that can assume only certain values, regardless of the level of precision to which it is measured. For example, the number of errors made on an invoice is a discrete variable as it can be only 0 or 1 or 2 or . . . and never 2.3, for example.

A discrete probability distribution is similar to a discrete frequency distribution (see Chapter 3) in that it consists of a list of all the values the variable can have (in the case of exact probabilities) or has had (in the case of empirical probabilities), together with the appropriate corresponding probabilities. A simple example will illustrate.

Example

The records of a shop show that, during the previous 50 weeks' trading, the number of sales of a certain item have been:

Number of sales/week	*Number of weeks*
0	4
1	16
2	22
3	6
4 or more	2

Construct the corresponding probability distribution

Solution

The variable here is clearly *discrete* (number of sales) and the probabilities are to be based on the *empirical* data given. Hence we shall have a discrete distribution of empirical probabilities. Now, using the definition of Section 9.2.

$$P(0 \text{ sales in a week}) = \frac{4}{50} = 0.08$$

Proceeding in this way, we can build up the distribution:

Number of sales/week	*P(number of sales/week)*
0	0.08
1	0.32
2	0.44
3	0.12
4 or more	0.04
	1.00

The *expected value* of a discrete probability distribution, $E(X)$, is defined as:

$$E(X) = \Sigma XP$$

where the summation is over all values of the variable, X, and P denotes the *exact* probability of the variable attaining the value X. At first sight, this appears to be an abstract concept, but an example will show that it has both a practical and an intuitively clear meaning.

Before looking at an important application, there is a special case of expected values worth mentioning. In the example above, on how many weeks would you expect there to be no sales during a trading period (4 weeks) and during a trading year (50 weeks)? The intuitive answers to these questions are that, since the probability of no sales in any one week is 0.08, we should 'expect':

in 4 weeks, no sales to occur in 4 = 0.08 = 0.32 weeks and in 50 weeks, no sales to occur in 50 = 0.08 = 4 weeks.

In fact, to fit in with these intuitive ideas, we extend the definition of expectation. If there are n independent repeats of a circumstance, and the constant probability of a certain outcome is P, then the expected number of times the outcome will arise in the n repeats is nP. In the above, when we have n = 4 weeks, our (assumedly independent) repeated circumstances, and the constant probability P = 0.08 (outcome of no sales in a week), then the expected value is:

nP = 4x0.08= 0.32 as intuition told us.

Example

If a coin is tossed three times, the resulting number of heads is given by the following probability distribution:

No. of heads (X)	*Probability* (P)
0	1/8
1	3/8
2	3/8
3	1/8
Total	1

Find:

(a) the expected number of heads in three throws of a coin;
(b) the expected number of heads in 30 throws of a coin.

Solution

(a)

No. of heads (X)	*Probability* (P)	*PX*
0	1/8	0
1	3/8	3/8
2	3/8	6/8
3	1/8	3/8
Total	1	12/8 i.e. 1.5

As expected, if a coin is tossed three times the expected (or average) number of heads is 1.5.

(b) If a coin is tossed 30 times we would intuitively expect 30/2 = 15 heads. Notice that 30 throws = 10 repeats of the three-throw trial and the expected number of heads is 10 × 1.5 = 15.

Expectation and decision-making

Many business situations require a choice between numerous courses of action whose results are uncertain. Clearly, the decision-maker's experience and judgement are important in making 'good' choices in such instances. The question does arise, however, as to whether there are objective aids to decision-making that, if not entirely replacing personal judgement, can at least assist it. In this section, we look at one such possible aid in the area of financial decision-making.

In order to introduce a degree of objectivity, we begin by seeking a criterion for classing one option as 'better' than another. One commonly accepted criterion is to choose the option that gives the highest expected financial return. This is called the expected value (EV) criterion.

Example

A decision has to be made between three options, A, B and C. The possible profits and losses are:

Option A: a profit of £2,000 with probability 0.5 or otherwise a loss of £500
Option B: a profit of £800 with probability 0.3 or otherwise a profit of £500
Option C: a profit of £1,000 with probability 0.8, of £500 with probability 0.1 or otherwise a loss of £400

Which option should be chosen under the EV criterion?

Solution

The expected value of each option is:

$EV(A) = (2000 \times 0.5) + (-500 \times 0.5) = £750$
$EV(B) = (800 \times 0.3) + (500 \times 0.7) = £590$
$EV(C) = (1000 \times 0.8) + (500 \times 0.1) + (-400 \times 0.1) = £810$

Thus, we would choose option C in order to maximise expected profit. However, it is arguable that a person or organisation that cannot afford a loss would opt for the 'safe' option B, which guarantees a profit in all circumstances.

Limitations of this approach

We are not advocating that the above approach is ideal: merely an aid to decision-making. Indeed, many texts that develop so-called decision theory further address at some length the limitations we shall discuss here. Suffice it to say, at this point, that attempts to overcome these problems meet with varying degrees of success, and so it is inconceivable that an 'objective' approach can ever replace the human decision-maker.

Another limitation that this approach shares with most other attempts to model reality is that the outcomes and probabilities need to be estimated. In Example 9.8.1 all the profits and losses are estimates, as are the probabilities attached to them. The subsequent analysis can never be more reliable than the estimations upon which it is based. There is also often a considerable degree of simplification with very limited discrete probability distributions being used when more complex ones or perhaps continuous distributions might be more appropriate. In Example 9.8.2 it is quite unbelievable that when times are good sales are exactly 50,000 – common sense tells us that they must be more variable than that. The probabilities in Examples 9.7.1, 9.7.2 and 9.7.3 are empirical, arising from past experience, and so have some degree of reliability unless demand patterns change dramatically.

In other cases (see the second quotation in Section 9.2), only subjective estimates of probabilities may be available, and their reliability may be open to question. There is therefore a doubt over this approach when subjective probabilities are used. Example 9.8.3 has another feature that will tend to lend some validity to this decision-making approach. It is a repeated decision, made every day. The expected values therefore have a commercial meaning: they are long-term average profits. Thus, if the store holder orders two units per day, she/he will average £ 16 per day profit. As this is higher than the average profits attainable from any other choice, this is a valid and sensible decision. In many cases, however, individuals or companies are faced with one-off decisions. An analysis of expected values would give the best average profits over a long run of many repeats of the decision, a circumstance that does not obtain in a one-off situation. One must question the use of the EV criterion in the latter case.

Finally, the comment made at the end of Example 9.8.1 demonstrates a serious deficiency in this approach: it takes no account of the decision-makers' attitude to risk. In the example, option B offered the lowest EV but also the least element of risk, and so an analysis of expected values does not give the whole picture. Particularly with one-off decisions, it can only give a guide to decision-makers.

Even with objective testing it is still important to be aware of the limitations of methods. The failure to take account of risk is a key criticism of this approach.

Characteristics of the distribution normal

In Section, the idea of a discrete probability distribution was introduced. The normal distribution is a continuous probability distribution. The values of probabilities relating to a normal distribution come from a normal distribution curve, in which probabilities are represented by areas: (Figure). An immediate consequence of probabilities being equated to areas is that the total area under the normal curve is equal to 1.

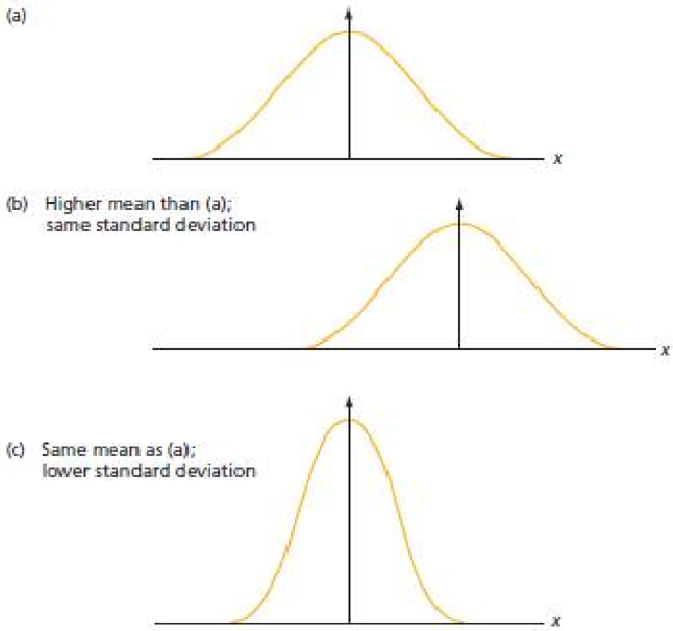

As Figure illustrates, there are many examples of the normal distribution. Any one is completely defined by its mean (μ) and its standard deviation (σ). The curve is bell shaped and symmetric about its mean, and, although the probability of a normal variable taking a value far away from the mean is small, it is never quite zero. The examples in the figure also demonstrate the role of the mean and standard deviation. As before, the mean determines the general position or location of the variable, whereas the standard deviation determines how spread the variable is around its mean.

Use of the tables of normal distribution

The preceding section describes the normal distribution but is insufficient to enable us to calculate probabilities based upon it, even though we know that the total area under the curve is one. To evaluate normal probabilities, we must use tables such as those given in your exam. These tables convert normal distributions with different means and standard deviations to a standard normal distribution, which has a mean of 0 a standard deviation of 1

This special distribution is denoted by the variable z . Any other normal distribution denoted x , with mean μ and standard deviation σ can be converted to the standard one (or standardised) by the following formula:

$$z = \frac{x - \mu}{\sigma}$$

Example

Use normal distribution tables to find the following:

(a) $P(0 < z < 1)$;
(b) $P(0 < z < 1.25)$;
(c) $P(z < 2.1)$;
(d) $P(0.7 < z < 1)$;
(e) $P(z > 1.96)$;
(f) $P(z < -1.96)$;
(g) $P(-1.96 < z < 1.96)$;
(h) $P(-1.2 < z < 2.8)$;
(i) $P(z > 3)$.

Solution

We shall use the abbreviation TE to mean 'table entry' so that, for example, TE(1) means the probability given in the table corresponding to the value $z = 1.00$. *Please note that this is not an abbreviation in standard usage.*

You will find it useful to look at the diagrams in Figure 9.3 while working through these solutions.

(a) $P(0 < z < 1) = TE(1.00) = 0.3413$. In the table this is the entry in row 1.0 and column 0.00.
(b) $P(0 < z < 1.25) = TE(1.25) = 0.3944$. In the table this is the entry in row 1.2 and column 0.05.
(c) $P(z < 2.1) = 0.5 + TE(2.1) = 0.5 + 0.4821 = 0.9821$. This probability includes all the negative values of z, which have a probability of 0.5, as well as those between 0 and 2.1 which are covered by the table entry.
(d) $P(0.7 < z < 1) = TE(l) - TE(0.7) = 0.3413 - 0.2580 = 0.0833$. This is given by the small area under the curve from 0 to 0.7 subtracted from the larger area from 0 to 1.
(e) $P(z > 1.96) = 0.5 - TE(1.96) = 0.5 - 0.475 = 0.025$. This tail-end area is given by the area under half the curve (i.e. 0.5) minus the area from 0 to 1.96.
(f) $P(z < -1.96) = P(z > 1.96) = 0.025$ by symmetry.
(g) $P(-1.96 < z < 1.96) = 1 - 2 \times 0.025 = 0.95$, which is the total area of 1 minus the two tail-ends. This symmetrical interval which includes 95 per cent of normal frequencies is very important in more advanced statistics.
(h) $P(-1.2 < z < 2.8) = TE(1.2) + TE(2.8) = 0.3849 + 0.4974 = 0.8823$. We have split this area into two. That from 0 to 2.8 is simply the table entry and that from −1.2 to 0 equals the area from 0 to +1.2 by symmetry, so it too is given by the table entry.
(i) $P(z > 3) = 0.5 - 0.49865 = 0.00135$. The method here is the standard one for tail-end areas but we wanted to make two points. The first is that virtually all normal frequencies lie between three standard deviations either side of the mean. The second is that, for symmetrical data, the standard deviation will be approximately one-sixth of the range.

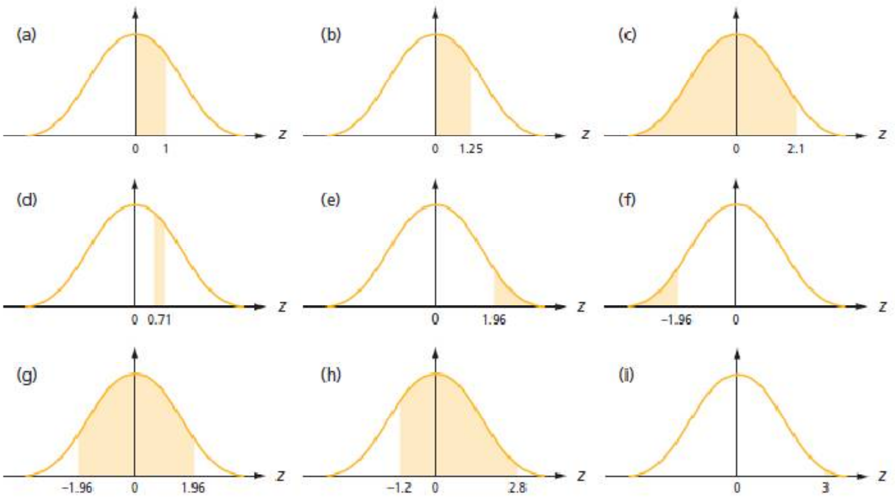

Venn diagrams

Venn diagrams were first developed in the 19th century by John Venn. They are a graphic device useful for illustrating the relationships between different elements or objects in a set. A Venn diagram is a picture that is used to illustrate intersections, unions and other operations on sets. Venn diagrams belong to a branch of mathematics called set theory. They are sometimes used to enable people to organise thoughts prior to a variety of activities. Using Venn diagrams enables students to organise similarities and differences in a visual way. A set of elements is a group which has something in common. Such a group could be all the children in a school. Such a Venn diagram is shown in Figure 9.6 .

A second Venn diagram (Figure 9.7) shows all the boys and the girls at the school.

Figure 9.6 Venn Diagram representing the set of all the children in a Village School

A second Venn diagram (Figure 9.7) shows all the boys and the girls at the school.

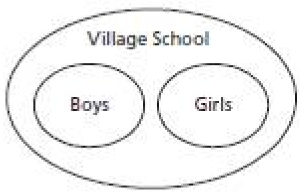

Figure 9.7 Venn Diagram representing the boys and girls in a Village School

A third Venn diagram (Figure 9.8) could be the children who play for the village foot-ball team.

Figure 9.8 Venn diagram representing the children from the village who play in the football team

A third Venn diagrame could be the children who play for the village football team. So, if we want to represent the children in a village school who also play football for the village team we would draw the following Venn diagram. Then the area of intersection between the two ellipses is a pictorial representation of the children from the Village School who are in the Village football team.

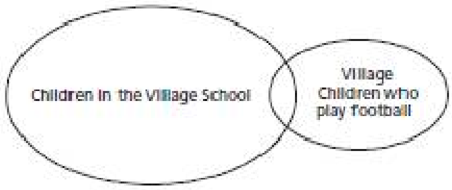

Using Venn diagrams to assist with probability

Venn diagrams can be a useful way of understanding calculations of probability. This will be explained through the use of an example.

Example

The probability that a woman drinks wine is 0.4.

The probability that a woman drinks gin and tonic is 0.7.

The probability that a woman does drink wine and does not drink gin and tonic is 0.1.

This example will show how a Venn diagram will assist in obtaining the probability that a woman selected at random drinks wine, and gin and tonic.

The first step is to draw a Venn diagram showing the two sets of women, that is those that drink wine and those that drink gin and tonic (Figure 9.10). The group that drink wine will be referred to as set W and the set that drink gin and tonic will be referred to as G&T.

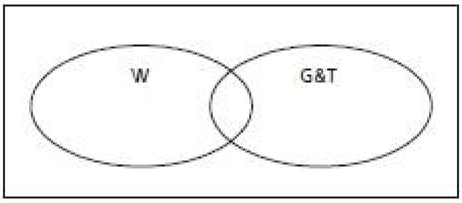

Figure 9.10 Two sets of women who drink wine, and gin and tonic

The next step in the process is to insert the probabilities in the Venn diagram and to start the calculations. Let x be the probability that a woman drinks both wine and gin and tonic. This is represented by the intersection of the two ellipses. This assumption now allows us to say that the area of the set W, which represents the women who only drink wine and not gin and tonic is 0.4–x. In the same way we can say that the area of the set G&T, which represents women who only drink G&T and not wine is 0.7–x. We also know from our data that the probability of a woman not drinking wine and not drinking gin and tonic is 0.1. These are all represented in Figure 9.11.

Uncertainty and risk

This section gives some background to the concepts of uncertainty and risk, and in particular, how probability can be used as a measure of risk. Although uncertainty and risk are connected concepts, they are sufficiently different to require separate explanations. Whilst a detailed understanding of uncertainty and risk is not required for your exam, it is useful to have an appreciation of these terms, and also the practical role that probability plays when assessing risks.

The terms uncertainty and risk are used in different ways by different commentators, researchers and professional practitioners in the business world, and it is useful to start with a discussion of what these words or terms mean.

Starting with risks, every business is said to face a series of risks, and as a result of these risks it is difficult to be able to say how the business will actually perform. These risks have a number of different origins or drivers and can include risks relating to marketing, production, technology, personnel, information systems and financial. (Note that this is not intended to be an exhaustive list of the sources of business risk.) For example, the marketing risks which a business faces may include the chance that the product which they are selling going out of fashion. It may include the chance that a large and powerful competitor may enter the market and take traditional customers away.

Another risk might be that due to a recession and the corresponding reduction in buying power, clients, and as a result sales decline. There may also be the risk of new government legislation being aimed at market controls which make the product too expensive for its traditional customers. Thus, from these examples, it may be seen that a risk can be a threat o the business.

If we considered production risks, there would be a similar list of issues, which may, for example, include the risk of a disruption in the supply of raw materials, or the risk of substantial price increases in essential services such as electricity and gas costs. Another risk could be that new legislation might require a much more costly regime of waste disposal in order to protect the environment.

It is not difficult to list similar types of risks for technology, personnel and information aspects of the business and the reader may do that for himself/herself. The category of risks referred to as financial risks are sometimes perceived to be different. Financial risks include the availability of credit, the cost of borrowing, the value (price) of the business' shares if they are quoted on the stock market. The issues which are discussed under this heading are often thought to be more volatile than those under the other more general business risks. But financial risks pose essentially the same sort of problems as other risks.

It is important for every organisation to be aware of the risks which it faces and there are no organisations which do not have a set of business and financial risks which directly affect them. Businesses need to assess the risks which they face and to take appropriate action. The assessment of risk is not a trivial matter and it requires considerable skill. The first step is to list all the possible risks, preferable by the major activities and functions of the business. Once this list is complete then an assessment needs to be made to estimate the likelihood of the risk occurring, as well as the amount of damage which the business could sustain if the risk materialises. The likelihood of the risk occurring is normally expressed on a scale from 0 (zero) to 1 where 0 means the risk will not actually occur and 1 means the risk will certainly occur1. You will recognise this as the description of probability, which was introduced in Section 9.2. Thus, a score of 0.50 suggest that there is a 50% chance that the risk will occur. However, there is no rigorous way of assessing the likelihood of a risk. It is simply a question of management judgement.

With regard to the question of how much damage a risk could do to the business, it is possible to make a more detailed and objective assessment of this and the use of financial estimates play a large role in this activity.

Once these numbers have been estimated then they may be used to calculate the expected value of a given risk. The expected value (EV) of a risk is the product of the chance (probability) of it happening multiplied by the size of the damage it will do to the organisation if the risk occurs. The expected value combines the probability and the damage of the result of the risk to give a figure which represents the relative importance of the risk to the business.

The calculation of expected value is for each item in the list of marketing risks is shown in Figure. Note that the expected value of the individual risks may be summed to give a total expected value of the risks emanating out of the marketing activity of the business.

The zero and 1 positions are theoretically limits on a spectrum. If the possible threat is rated at 0 then it need not be included in risk analysis as it will not occur. If the possible threat is rated at 1 then it need not be included in risk analysis as it has either occurred or will certainly occur and therefore there is not the element of chance which is inherent in the definition of risk.

Marketing risks

Type of risk	Probability of occurrence	Estimate of financial damage	Expected value
Product out of fashion	0.30	1,200,000	360,000
Entry of big competitor	0.25	1,500,000	375,000
Recession	0.10	2,000,000	200,000
Legislation changes	0.05	3,000,000	150,000
Total Marketing risks			1,085,000

From Figure it may be seen that the most damaging risk which the business faces is the possibility of the entry of big competitor. If this occurs then the loss to the business is expected to be £375,000. The second greatest risk is the possibility that the product could go out of fashion. If this occurs then the loss to the business is expected to be £360,000. The size of the other risks may be read from Figure 9.13.

In a similar way, analyses may be undertaken for other risks; production, technology, personnel, information systems and financial. There are two courses of action which management may take in the face of these risks. The first is to initiate risk avoidance measures, and the second is to establish a programme which will mitigate the impact of the risk if it should occur. However, a detailed discussion

of this is beyond the scope of this book. The technique described above whereby the expected values are calculated, may also be applied to other business calculations in which it is appropriate to include risk assessments. There are two major approaches to this. Both approaches call for the use of a range of estimates of the projected values of cost and benefits for the production of budgets. These techniques are frequently used in capital investment appraisal or assessment. By using the maximum estimated and the minimum estimated values, a range of possible outcomes for the investment are calculated. The results of these calculations which will be a range of values themselves will show the result of the investment if the impact of the risks are minimal i.e. few of the threats materialise, and also the result of the investment if the impact of the risks are large i.e. most of the threats materialise. Management judgement is then required to decide which of these scenarios is the most plausible. There are sophisticated variations of this approach which use a technique known as Monte Carlo simulation, although this is beyond the scope of this book.

Before concluding this section it is appropriate to more comprehensively define risk in broad terms. The risk of a project is the inherent propensity of the estimates concerning the cost, time or benefits for the project not to be achieved in practice, due to foreseeable and unforeseeable circumstances.

Although risk is often spoken of in a negative context i.e. the project will cost more than budgeted for, or take longer than originally believed, it is obviously the case that FUNDAMENTALS OF BUSINESS MATHEMATICS **375** PROBABILITY sometimes projects are completed below budget and before their deadlines. Thus risk may enhance the potential of a project as well as detract from it. It is clear that risk is based on the fact that the future is always unknowable or uncertain in the sense that we are unable to be sure of anything before it happens. With regard to the concept of uncertainty when it is not possible to make any estimate of the probability or the impact of a future event or threat we do refer to its risk – rather we refer to uncertainty. For example, we cannot state with any degree of confidence about the risk that any large banks will become insolvent. This is because we have neither a way of estimating the chance nor a probability of that happening (nor the ability of estimating the impact that such an event would have on our society). However, we can safely say is that at present, the future of banks is uncertain.

Thus uncertainty may be thought of as a sort of risk about which nothing may be estimated. While risk is a concept which is used extensively by business and management practitioners, the concept of uncertainty is employed by economists when they are referring more generally about business affairs in the economy.

REVIEW QUESTIONS

1. Define probability. Describe Addition rules of probability?
2. Describe the process of probability of opposites. What is the multiplication rule of probability?
3. Describe discrete probability and distributions. What are the limitations of this approach?
4. What are the Characteristics of the distribution normal?
5. Describe Venn diagrams. Discuss the use of Venn diagrams to assist with probability.
6. Write a short note on marketing risks.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

UNIT-10 SPREADSHEET SKILLS USING EXCEL

CONTENTS

INTRODUCTION

Introduction

A spreadsheet is a multipurpose piece of software which may be used for calculations, drawing graphs or handling data in a way similar to a database program. All these functions are available in most spreadsheets at both an elementary level and a highly sophisticated level. In a spreadsheet like Excel, complex problems may be handled by using macros. A spreadsheet may also be described as a computer program that allows data to be entered and formulae to be created in a tabular format. It was designed to mimic a large paper-based worksheet with many rows and columns. Spreadsheet information is stored in cells and the power of this technology lies in the way each cells can store numerical or alphabetical data or a formula for operating on other cells. A cell can also hold references to other spreadsheets or objects (such as graphics).

The first spreadsheet program was called VisiCalc (standing for Visible Calculator). The program was created by two American Computer Science researchers called Dan Bricklin and Bob Frankston and was first released through Software Arts in 1979. It was in 1982 that the next improvement in spreadsheet technology came about with the release of Lotus 1-2-3. A much faster running program and increased functionality helped Lotus 1-2-3 to become the market leader – a place it kept until 1989 when Excel arrived as part of Microsoft's Office suite of programs to run on its Windows operating system. Although Lotus produced a Windows version of 1-2-3, it did not stand up against the popularity of the Office suite and thus Excel has become the de facto spreadsheet in use today.

For any readers using an alternative, the basic functionality of most spreadsheets is pretty similar and the exercises and examples used in this book can still largely be followed. In your exam, it is important that you input your answer exactly how you would enter it into Excel, for example, including the leading = sign. It is, of course, possible to enter alternative, but equivalent correct formulae; the assessment software will handle this.

Spreadsheet terminology

It is worth clarifying the different application areas within the spreadsheet.

Workbooks and Worksheets

An Excel file is referred to as a workbook. A workbook can consist of a single worksheet or can be a combination of multiple worksheets, charts, databases etc. An Excel file is saved on disk with a .xls file extension.

Cells

A worksheet is described by column letters and row numbers and each row/column coordinate is referred to as a cell. In Excel there are 256 columns labelled A through IV and 65,536 rows. This in theory provides 16,777,216 cells into which information can be placed! In actual fact the number of cells that can really be used is restricted by the specification of the computer and the complexity of the data and formulae being worked on.

Note on macros and application development

It is possible in Excel to record a series of keystrokes and/or mouse clicks which can be stored in a macro. The macro can then be run whenever that series of keystrokes and/or mouse clicks is required. For example, to print a specific area of a spreadsheet, or to save a file with a particular name. In addition to recordable macros, Excel has a powerful computer language called Visual Basic for Applications (VBA). With VBA it is possible to program the spreadsheet to perform in very individualistic ways.

The development and use of macros requires a substantial understanding of the spreadsheet and is thus beyond the scope of this book.

Getting started with Excel

When the Excel program is launched a blank spreadsheet is displayed. Figure shows what this looks like and highlights some of the main features of the system.

Note: This figure assumes that Excel has not been customised in any way. If your system has been installed with customised toolbars the screen may not look the same. It would be preferable to re-install Excel without customisation for the purposes of working with this book.

Workbooks of files

When Excel is loaded a new, blank workbook is displayed as shown in Figure. This workbook is called BOOK1 and consists of a number of blank worksheets. Each worksheet is labelled on a tab at the bottom of the workbook. You will be able to customise the name of these worksheets as you use them, but at this stage they are labelled Sheet 1, Sheet 2 etc. Sheets can be moved or copied between workbooks, and you can reorganise sheets within a workbook. In addition you can have several workbooks open at the same time, each in its own window.

Worksheets

Most of the work you do will use worksheets. As you can see in Figure 10.1 a worksheet is a grid of rows and columns, forming a series of cells. Each cell has a unique address. For example, the cell where column C and row 8 intersect is referred to as cell c8. You use cell references when you create formulae or reference cells in command instructions. The active cell is the one into which data will be placed when you start typing. You can determine the active cell by the bold border it has around it. When you open a new workbook this will be cell a1 on Sheet1. To change the active cell you can either use the directional arrow keys on the keyboard to move one cell at a time to the left, right, up or down, or you can use the mouse to move the pointer into the required cell and then click once on the left mouse button.

Scroll bars

To the right and the bottom right of the screen there are scroll bars which allow you to scroll up and down and left and right around the active window. Click on the down arrow in the vertical scroll bar which will scroll the worksheet down by one row.

Status bar

At the bottom of the screen are the horizontal scroll bar and the status bar which display information about the current document or the task you are working on. The exact information displayed will vary according to what you are doing. When you open a new workbook there are indicators to the right of the status bar that are highlighted if the caps lock key, num lock key or scroll lock key is activated.

Toolbars

As in all Windows applications the toolbars allow quick access to commonly used commands. On starting Excel the Standard and Formatting toolbars are displayed. Move the pointer over one of the toolbar buttons and notice the name is displayed in a small box below the selected button. This is called a ToolTip. A brief description of what the button does is displayed.

Good spreadsheet design

Whether a spreadsheet is being developed for specific business mathematical calculations or as a forecasting plan, a profit and loss account or a marketing plan it is essential that duecare and attention be given to its design and structure. Establishing some rules as to how all the spreadsheets in a department or organisation are developed enables different people to look at different plans and feel familiar with the layout, style, reports, charts, etc. This is in much the same way as users feel familiar with software applications that have a similar interface such as those in the Microsoft Office suite of products.

The objectives of good design in spreadsheet terms are exactly the same as those required for any other software development:

1. to ensure that the spreadsheet is as error free as possible;
2. to ensure that the spreadsheet can be used without much training or control;
3. to minimise the work required to enhance or change the spreadsheet.

If care is taken to ensure sound structure and good design a spreadsheet will be straightforward to develop, easy to read, simple to use, not difficult to change and will produce the required results.

The plan developed over a number of stages in this chapter illustrates a variety of aspects of the principles of spreadsheet design and development. The series begins with a plan that has had little or no thought put into its design and layout and as the chapter proceeds ways of improving and enhancing the plan are identified and explained.

Getting started

The spreadsheet in Figure is a simple profit projection that may be of use to the author, but is unlikely to be helpful to anyone else. This is clearly a quick one-off plan which has been prepared with very little care and which may well not even be saved on the disk.

Problems with this spreadsheet

The immediately obvious problems with this spreadsheet are that it has no title, it is not clear what the columns represent, i.e. are they different periods or perhaps different products, and the author is unknown.

With regards the data itself, the figures are hard to read as there are varying numbers of decimal places. Whilst perhaps there has been a growth in sales and price, the percentage has not been indicated. The costs line could be misleading as no indication of where the costs have been derived is supplied.

Positive aspects of this spreadsheet

If the author of the spreadsheet required a quick profit estimation based on known data and growth rates for sales units, price and costs then the spreadsheet has supplied that information quickly and in a more concise form than would have been achievable using a calculator and recording the results on paper.

Ownership and version

In Figure the three major shortfalls of the first spreadsheet have been remedied. The plan has also been given a title and author details have been included. It is important that every business plan have a clear owner who is responsible for overseeing the accuracy and maintenance of the system. A name plus some form of contact details should always be included.

Problems with this spreadsheet

The construction of the data and results is still unclear and the lack of formatting makes the figures hard to read. The costs remain grouped together.

Positive aspects of this spreadsheet

In addition to the owner details having been added to the plan, the date when the plan was written is a useful feature. The date becomes particularly important when the question of spreadsheet versions arise. Note that the date has been entered here as text. If a date function had been used it would be continually updated each time the file is retrieved, whereas here it is the date of the last update that is required. The ruling lines above and below specific sections of the spreadsheet are also quite helpful. This can be quickly achieved using the automatic formatting features. These are accessed via the Format Auto format command.

Formatting

In Figure the data for the four quarters is totalled and reported as an annual figure. The values in the plan have also been formatted with the majority of figures being formatted to zero decimal places and the price line to two decimal places. One of the automatic formatting options has been selected to shade and outline the plan.

Problems with this spreadsheet

By looking at the plan in Figure it can be seen that the sales and the costs both increase over time. However it is not clear by how much because the sales growth factor and the increase in costs have been incorporated into the formulae as absolute references. The inclusion of absolute values in formulae is not recommended and can lead to GIGO. To change the sales growth factor in Figure two processes are required. First, cell c5 is accessed and the edit key pressed. The growth factor is changed and enter is pressed. This has changed the formula in this one cell, but only once the formula has been extrapolated across into cells d5 and d6 is the amendment complete. It is not difficult to see that there is room for error here in a number of different ways.

Positive aspects of this spreadsheet

Having a current date and time indicator displayed on the spreadsheet ensures that a hard copy report will reflect the date, and perhaps more importantly the time it was printed. This is achieved through the now( ) function, which can be formatted with a range of different display options. Because it is likely that a spreadsheet will be recalculated, even if it is set to manual calculation, before printing, the date and time will always be up to date. It is of course possible to include the date and time in headers and footers, but during the development phase of a system the page layout is less relevant than printing the section being worked on and so thought should be given to the positioning of the now function. The cells in this plan have now been formatted, which makes the data easier to read. When formatting a spreadsheet it is important to consider the entire plan and not just the cells that are currently being worked on. The entire spreadsheet should be formatted to the degree of accuracy required for the majority of the plan then those cells that need to be different, such as percentages, can be reformatted accordingly. This is quickly achieved by right clicking on the top left corner of the spreadsheet at the intersection between the column letters and row numbers and then select format cells. Whatever formatting is now applied will affect the entire worksheet.

It is important to understand that formatting cells only changes the display and does not affect the results of calculations that are still performed to the full degree of accuracy, which is usually 16 significant decimal places. This is why a cell containing the sum of a range of cells might display an answer that does not agree with the result of visually adding the values in the range.

The only safe way to ensure that the results of a calculation are actually rounded to a given number of decimal places the round function is required. Figure shows two tables representing the same extract from a profit and loss account. In both cases all the cells have been formatted to zero decimal places, but in Table B the round function has been incorporated in the formulae for cells f15 through f20. The formula entered into cell f15, which can then be copied for the other line items is: = round(sum(b15:e15),0) The effect of the round function can be seen in cell f20. By visually adding up the numbers in the range f15 through f19 the result is 78111 whereas the formatting of these cells without the use of the round function in Table A returns a value of 78112 in cell f20. Having applied the round function to a cell any future reference made to that cell will use the rounded value.

Excel does offer an alternative to the round function in the Precision as displayed option within tools: options: calculation. This command assumes that calculations will be performed to the level of accuracy currently displayed. The danger of using this command is that when data is changed or added to the spreadsheet the command is no longer valid and it is then necessary to repeat the command to update the spreadsheet – this is another invitation to gigo.

Documentation

Spreadsheet developers are notoriously bad at supplying documentation and other supporting information about the plan. There are a number of features offered by Excel to assist in the documenting of plans including the insert comment command. Figure 10.6 shows a comment being entered onto a plan – notice how the user name of the comment author is included. This is useful when a team of people are working on a system. The presence of a comment is indicated by a small red triangle on the cell – to read the comment move the cursor over the cell and it will automatically be displayed. A word of caution concerning the use of comment boxes – they take up a considerable amount of space and if used widely they can make a noticeable difference to the size of a file. To clear all the comments use the edit clear comments command.

The provision of a hard copy report showing the logic used to create a plan is also helpful as this is the ultimate reference point if a formula has been overwritten and needs to be reconstructed.

In Excel there is a shortcut key to display the formulae which is ctrl = (accent grave). Alternatively this can be achieved through the tools options view command and then check the Formulas box.

In addition to providing documentation for a spreadsheet system, looking at the contents of the cells as opposed to the results can also be a helpful auditing tool. For example, Figure 10.7 highlights the fact that there are still values embedded in formulae which is not good practice and is addressed in the next version of the plan. A third form of documentation which can be particularly useful for large systems is the ‘sentence at the end of the row’ technique. Requiring less file space than comment boxes, and always on view it can be useful to have a brief description of the activity taking place in each row of a plan.

Minimising absolute values

One of the reasons that spreadsheets have become such an integral part of the way we do business is the fact that they facilitate quick, easy and inexpensive what-if analysis. What-if analysis may be defined as the process of investigating the effect of changes to assumptions on the objective function of a business plan. Performing what-if analysis on the opening sales assumption or the opening price assumption is quite straightforward, involving placing the cursor on the figure and entering the new value. On pressing enter the spreadsheet is re-evaluated and all cells which refer to the changed values, either directly or indirectly are updated. The success of performing even the simplest what-if analysis is dependent on the spreadsheet having been developed with the correct series of relationships. For example, changing the opening sales value in Figure 10.8 would automatically cause the other quarter sales values to recalculate, as well as the revenue, costs and profit lines, because they relate, through the cell references in the formulae, either directly or indirectly to the sales value in cell b5.

However, as already mentioned this plan incorporates absolute values in the formulae for sales and costs growth. Furthermore, the price is a fixed value and has been entered once into cell b6 and the value has then been copied into the other periods. This presents problems when what-if analysis is required on any of these factors.

Problems with this spreadsheet

Because no growth in the price is required the opening value of 12.55 has been copied for the four quarters. Whilst this is fine all the time a price of 12.55 is required, it presents a problem when the price requires changing. With this spreadsheet it would be necessary to overwrite the price in the first quarter and then copy the new value for the remaining three quarters. The same applies if the sales growth or the cost factors required changing. To prevent these problems arising, a different approach to the development of the plan needs to be taken.

In the first instance all growth and cost factors should be represented in a separate area of the spreadsheet – even on a different sheet altogether in the case of a large system with a lot of input. The factors can then be referenced from within the plan as and when they are required. Figure 10.9 shows the adapted layout for this plan after extracting the sales growth and costs factors.

Having the growth and cost factors in separate cells means that the formulae need to be changed to pick up this information. Figure 10.10 shows the amended formulae for this plan. Note that the references to cells d15,d16 and d17 are fixed references. This is achieved by placing the $ symbol before the column letter and row number, i.e. $ d $ 15, and means that when the formula is copied the reference to cell d15 remains fixed. A shortcut key to add the $ symbols to a cell reference is f4.

In this plan an option in the growth factors has been included for the price, despite the fact that in this plan the price does not change. It is important to always think ahead when developing any plan and although the price does not currently change, it might be necessary to include a percentage increase in the future. Having the facility for change built-in to the plan could save time later – and for the time being the growth factor is simply set to zero.

Removing the growth and cost factors from the main body of a business plan is the first step in developing a data input form which will ultimately separate all the input data from the actual logic of the spreadsheet. This separation of the data allows the logic cells to be protected from accidental damage. This is discussed further in the Template section of this chapter.

Control checks for auditing

As already mentioned, spreadsheet users are not inherently good at auditing plans as thoroughly
as perhaps they should, and therefore an important aspect of spreadsheet design is to build into the system checks on the arithmetical accuracy that will raise the alarm if things begin to go wrong. This might include validating input data through the use of an if function, or performing a cross-check on a calculation.

When creating calculation checks the first step is to select a number of key items from the model, whose result can be calculated using a different arithmetic reference. A separate sheet can be allocated for data validation and arithmetic checks. For example, in Figure 10.11 below the Year End Gross Profit has been calculated by referencing the individual total values in column f and then by totalling the values in the Gross Profit row. An if function is then applied to compare the two results and if they are not the same the word 'error' is displayed in cell d8.

The formulae required in cells d6 and d7, which calculate the year-end gross profit from the plan illustrated in Figure

Charts

It is useful to support the information supplied in business plans with charts. In the profit and loss account used in this chapter various charts might be useful, for example to show the relative impact of price and sales volume figures. Although charts can be placed on the same worksheet as the plan, it is usually preferable to keep graphs on separate chart sheets. The exception might be if it is appropriate to view changes on a chart at the same time data in the plan is changed, or if a spreadsheet is to be copied into a management report being created in Word. Figure 10.12 is an example of the type of chart that might be produced from the plan used in this chapter.

Tips for larger plans

The plan used in this chapter has been a simple quarterly plan, but in many cases business plans will be larger and more complex. Figure 10.13 is an extract from a five-year quarterly plan. Although it is not obvious by looking at Figure , each year in this report has been formatted with a different colour font. This is a useful technique when working with large models because it enables the user to quickly know which part of the plan is being viewed or worked on, without having to scroll around the spreadsheet to see the titles.

This colour coding can then be carried over to summary reports, and other reports pertaining to the different parts of the plan.

From a design point of view it is preferable to place different reports associated with a plan on separate worksheets. The report in Figure 10.14 , which has been placed on a separate sheet called Summary is created by referencing the cells from the yearly totals in the main plan.

Templates

A business plan that requires time and effort to design and implement is likely to be in regular use for some time. In addition, the data in the plan will almost certainly change, either as situations within the business change, or on a periodic basis. In such circumstances it is advisable to convert the developed plan into a template, into which different data can be entered whenever necessary.

A template is a plan that contains the logic required, i.e the formulae, but from which the data has been removed. When new data is entered so the formulae will be calculated. Figure shows the simplest approach to creating a template. Taking the one-year quarterly plan used in this chapter the input data and growth factors have been removed and these cells have been highlighted by shading the cells. When the input cells are set to zero, all other cells that are directly or indirectly related to those cells should also display zero. The only exception to this is if there are division formulae in which case a division by zero error will be displayed. The act of removing the data is in itself a useful auditing tool, because if values are found in any cells this indicates that there is an error in the way that the plan was developed which can be rectified. When the template is complete the spreadsheet should be protected and then only the input cells unprotected in order that the user can only enter data into the designated cells. This is a two-step process. First, the cells into which data can be entered are unprotected using the format cells protection command and removing the tick on the Locked box. The second step is to then enable protection by selecting Tools Protection Protect Sheet.

It is also important to save the file now as a Template file as opposed to a Worksheet file. This is achieved by selecting file save as template (.XLT) in the file type box. The location of the template file defaults to the directory where other Microsoft Office template files are located. To use the template file new is selected which accesses the Template directory and when a file is selected a copy of it is opened, leaving the original template unchanged on the disk.

Data input forms

A further enhancement that makes working with templates easier to control is to remove all the data from the main plan and place it on one or more data input forms which will normally be located on separate worksheets. Figure is a data input form for the quarterly plan, and Figure shows the amended formulae in the plan which picks up the data from the input form.

There are many benefits to be derived from using data input forms including the fact that the data can be checked more easily. Sometimes it might be possible to design an input form that is compatible with a forecasting or accounting system so that the data can be electronically picked up from the other system without having to type it in again. Even if this is not possible, the order of items in the data input form does not have to be the same as the order in which they are referenced in the logic, which means that the data input form can be created to be as compatible with the source of the input data as possible. Furthermore, the worksheet containing the logic for the plan can be protected, and if necessary made read-only in order to maintain the integrity of the system. It is not a trivial task to change existing systems to be templates with data input forms, and it will also take a little longer to develop a new system in this way, as opposed to incorporating the data with the logic. However, the ease of data input and ongoing maintenance should make the additional effort worthwhile.

The use of spreadsheets by management accountants

There are a number of different ways in which the management accountant can use a spreadsheet in his or her work. In the first place, spreadsheets are especially useful in the performance of calculations. In addition to the basic mathematical operators discussed such as addition, subtraction, division, multiplication etc., there are many other functions which will be of direct use. These include NPV, IRR, PV to mention only three. There are infact more than 350 built-in functions in Excel. When it comes to repetitive calculation the management accountant can set up templates that can be used again and again. There are many different aspects to the way that Excel can be used for planning and those who are interested in more detail should consult the Elsevier Cima Publication, Financial Planning Using Excel – Forecasting, Planning and Budgeting Techniques, by Sue Nugus, 2005. In addition to the calculation side of the spreadsheet the management accountant will find useful the ease with which graphs and charts can quickly be created in Excel.

REVIEW QUESTIONS

1. Discuss about spreadsheet terminology. What are the Positive aspects of this spreadsheet?
2. What are the Problems with spreadsheet?
3. Describe Control checks for auditing in spreadsheet.
4. What are Charts and Templates?
5. Give the detail of Data input forms.

FURTHER READINGS

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap

BIBLIOGRAPHY

1. Business Mathematics and Statistics- Andy Francis
2. Agarwal B.M.
3. Introduction to Business Mathematics- R. S. Soni
4. Business Mathematics : Theory & Applications- Jk. Sharma
5. Business Mathematics- Trivedi Kashyap
6. Business Mathematics and Statistics Quantitative Techniques- Calman Goozner, Thomas P. Walsh
7. Fundamentals of Business Maths -Cima, Graham Eaton