Chapter 5:
Recurrence Calculations

Objectives

This chapter focuses on using spreadsheets to build models of processes that change with time – such as the calculation of a running total of some quantity, the compound growth of an investment, or the growth of a population, for example. Here is a partial list of what the chapter covers:

- Further practice of good design techniques
- Using recurrence formulas

Preparation

To be able to complete the exercises in this chapter in a reasonable amount of time it is essential that you are properly prepared, You should:

- Read the whole of this chapter very carefully

Exercise 1 - A Running Total Example

Suppose you have a row (or column) of numbers and you want to have another row that displays the cumulative (running) total of the first, i.e.

Row:	8	3	12	9	11	5
Cumulative total:	8	11	23	32	43	48

There are many applications of this – calculating the balance in a bank account as deposits are made for example.

Open the file Exercise 1 (ch5_Ex1) in Support Files (Chapter 5) on the course website. The model contains a Comments worksheet and a worksheet called CumulativeData (Figure 5.1), which contains just one long row of numbers at the moment. You want to add another row that will contain the running total.

Figure 5.1 – the Data worksheet of Exercise 1

	A	B	C	D	E	F	G	H	I	
1	**Data:**	124.93	120.43	193.25	177.00	179.01	254.79	42.80	242.22	
2										
3										

Add the row label first and then name the 2 rows. Select the first cell in the Total row. Since this is the first one it should simply contain the first value in the Data row, so construct a formula that will calculate this result for the first cell:

<p align="center">=Data</p>

The second cell in the total row should contain the sum of the first cell in that row and the second cell in the Data row. However, the formula

<p align="center">=Total + Data</p>

will cause a circular reference. Since the formula is in the Total range, it refers to itself. To avoid circularity we cannot use the name Total. The formula should be:

<p align="center">= B2 + Data</p>

Now select this second cell and fill it along the row. You should not be surprised to see that this formula produces the results we want. Examine a few of the formulas that were created when you filled along the row and make sure you understand why they produce the results we were aiming for.

You have just seen the two elements required when constructing a recurrence formula – initialising the first value and constructing the recurrence. You have also seen that this is the one time when it's not possible to use a named range in a formula.

Exercise 2 – **Bank Account Balance**

Open the file Exercise 2 (Lab5_Ex2) in Support files (Chapter 5) on the course website. The model contains a Comments worksheet and a worksheet called Account, which contains a balance brought forward (from a previous month perhaps) and lists of deposits and withdrawals. The aim is to add another column that shows the balance after each deposit or withdrawal transaction.

Enter and format a title for the column, define names for the existing data and start to create the formula for the first cell of the new column. This formula should be straightforward since you simply need to initialise the balance to that brought forward.

The formula for the second cell in the balance column must set up the recurrence calculation of the balance. It is in effect simply the implementation of:

<p align="center">Balance = previous balance + deposit - withdrawal</p>

Build this formula yourself using the spreadsheet syntax and fill it down the column. Since "balance" and "previous balance" are in the same range you need to use a cell reference for this part of the formula. Note that the formula will work because Excel treats blank cells as if they contain zeroes. If the deposit or withdrawal value is blank the formula will actually add or subtract zero.

Exercise 3 - Investment Growth

This next example calculates the growth of an investment which earns interest at a given interest rate and where the interest is added to the investment, i.e. a compound growth situation.

Open the file Exercise 3 (Lab5_Ex3) in Support files (Chapter 5) on the course website. The model contains a Comments worksheet, a worksheet called CapitalGrowth and one called Parameters. The CapitalGrowth worksheet is intended to show the compound growth of an investment as interest is added to the investment amount at the end of each period of the investment. There is a column for counting the periods and another for calculating the amount of the investment.

The Parameters worksheet contains the initial investment amount, the date of the investment, its maturity date, the interest rate (which is per annum), and the interest period. The interest period is the period of time after which interest is paid into the account. It might be yearly or every 6 months for example. The number of periods (between the investment date and the maturity date) is to be calculated and then used in the CapitalGrowth worksheet to make sure only that many periods are shown.

The Parameters Worksheet

The labels for the cells have been entered for you, so now enter some values – say $30,000 for the initial capital, 01/31/2003 for the investment date and 01/31/2013 for the maturity date. Enter an interest rate of 8% and an interest period of 0.5 (meaning ½ a year, or 6 months).

You would like to be able to change these values and still have the CapitalGrowth worksheet display sensible results. For example, if you changed the interest period to 1 the CapitalGrowth worksheet should only show 10 rows (since there are 10 years between the investment and maturity date); but with an interest period of 0.5 there would be 20 rows displayed, or with an interest period of 3 months there would be 40 rows displayed.

Create the formula to calculate the Term. In spreadsheets there are functions to manipulate dates. In this case we want to find the number of years between the investment date and the maturity date. We'll keep things simple and assume that the maturity date will always be the same day and month as the investment date.

The function YEAR will obtain the year from a date value, so the formula should implement the following description:

> *Obtain the year from the maturity date and subtract the*
> *year obtained from the investment date.*

First define appropriate names for the relevant data you are about to use in the formula and then implement this formula. Check that you have the correct answer.

Once the Term is computed you need a formula to calculate the number of periods by dividing the Term by the interest rate period

Commentary

Take a moment to experiment with entering dates and formatting numbers as dates. Try the following: select any empty cell, enter the integer 76 (press enter and then select the cell again). Choose the Home tab and then click on the little arrow icon at the lower right corner of the Number group. Choose Date from the list of possibilities in the Number tab of the Format Cells window (Figure 5.2). Choose any one of the Types. Does the result make sense to you? Try entering, in an empty cell, a different number – say 367 – and formatting it.

<u>Figure 5.2 – formatting a cell as date</u>
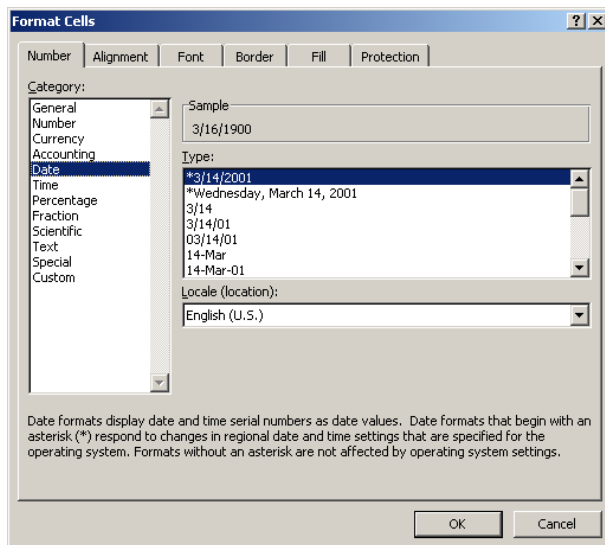
Also try different methods of entering a date value directly into a cell. You already entered 01/31/2003 but there are many other ways of entering a date. Format the date you entered as an integer by choosing General from the Format Cells window. Make sure you understand the results of this experimentation.

The CapitalGrowth Worksheet

The labels and headings have already been entered for you in this worksheet. You need to develop the formulas for calculating the values.

Focusing on the Period column you realise that you want a 0 to be in the first cell and that the column should then simply count up to the value set by the Number of Periods: in the Parameters worksheet, i.e. 20 using the parameter values suggested above.

The second cell should simply be the first cell + 1 – so enter this formula, and then fill it down the column for about 50 cells. You will find that the values do not stop at 20 – the actual number of periods - as you would like them to. The solution to this is to use an IF function in the formula which says (in English – you need to translate this to the spreadsheet syntax):

> *If the last value of* Period *is less than the number of periods specified in the* Parameters *worksheet, increment the* Period *value by 1, otherwise blank the cell (i.e. "").*

Change the formula in the second cell of the column to include the IF function and then fill it down the column again. Now there should be no values computed after Period=20.

Now consider the Capital column. The initial value should be $30,000 – the value of the investment amount from the Parameters worksheet. Enter the formula that will do this.

The second cell should be the initial amount plus the amount of interest earned in the first period, and the third cell should be the amount from the second cell plus the amount of interest earned in the (second) period.

The amount of interest will be calculated as the capital amount multiplied by the interest rate (which gives the interest for a whole year) multiplied by the interest period (which is given as a fraction of a year). This amount should be added to the capital in order to calculate the capital for the next period. Construct this formula yourself and fill it down the column. Figure 5.3 shows the first few values from the CapitalGrowth worksheet.

Figure 5.3 – part of the CapitalGrowth **worksheet.**

ch5_Ex3.xlsx				
	A	B	C	D
1	**Compound Capital Growth**			
2		**Period**	**Capital**	
3		0	$ 30,000.00	
4		1	$ 31,200.00	
5		2	$ 32,448.00	
6		3	$ 33,745.92	
7		4	$ 35,095.76	
8		5	$ 36,499.59	

You should now be able to change the interest period in the Parameters worksheet and see the CapitalGrowth worksheet correctly calculate the results. If you change the interest period to something small – one month for example – the number of periods required would be 120 with the investment and maturity dates used above. In this case you might find that you didn't fill the formulas in the CapitalGrowth worksheet down far enough to calculate so many values. You should correct this situation if necessary.

Exercise 4 – **Pepper Moth Population Simulation**

Introduction
In this part of the lab you will develop a model that simulates the population changes of a pepper moth colony. You will be deriving formulas from a description of the problem – not an easy task – so it is absolutely critical that you attempt to write these formulas down before sitting at the computer to use the spreadsheet.

The Population Model
This population simulation is intended to calculate the number of light and dark pepper moths. The population starts out being all light coloured and then each generation produces, by mutation, a few dark coloured moths. The dark coloured moths have a greater chance of survival because predators find them harder to see, so fewer of them are eaten. Thus each generation has more dark moths and fewer light coloured moths. The mutation works both ways, however, so a few of the moths born to dark coloured parents are actually light coloured because of mutation.

Given the population of light and dark moths we can calculate the number of light coloured moths in the next generation as follows:

- Predators eat some light moths, and the others survive to reproduce. The percentage of the current population of light moths that is eaten is the light predation rate.
- The surviving light moths reproduce at a certain rate. We'll pretend that light moths only mate with light moths to avoid having to model gene dominance etc. The reproduction rate might be 1.5 per moth in the current population.
- Some of the newly born moths from light coloured parents are in fact dark. This percentage is the light mutation rate.
- Some portion – perhaps 100% – of the old generation dies, and this is called the light coloured death rate.
- Some of the newly born moths from dark parents are in fact light coloured and hence must be added to the light coloured population. The percentage of offspring of dark parents that are light coloured is the dark mutation rate.

The dark moth population can be calculated in a similar fashion:

- A certain number are eaten by predators - the dark predation rate (different from the light predation rate).
- The survivors reproduce – the dark reproduction rate (same as the light reproduction rate).
- Some of the newly born moths are in fact light - the dark mutation rate.
- Some portion of the old generation dies - the dark coloured death rate (same as for light moths).
- Some of the newly born moths from light parents must be added to the dark coloured population - the light mutation rate.

So given some population of light and dark moths the population immediately after the birth of the next generation (and assuming the death of the previous generation occurs immediately after the birth of the next generation) is given by:

For light coloured moths:
> Next population = number of adult light survivors + number of light births – number of light births that were actually dark by mutation + number of dark births that were actually light by mutation – number of deaths

For dark coloured moths:
> Next population = number of adult dark survivors + number of dark births – number of dark births that were actually light by mutation + number of light births that were actually dark by mutation – number of deaths

Let's look more closely at the calculation for light coloured moths. The number of adult survivors is the population minus those that were eaten, which written as a formula is:

> light survivors = light population - light population * light predation rate

The predation rate might be 0.5 for example, meaning that half the population was eaten by predators.

The number of light births, written as a formula is:

> light births = light survivors * light birth rate

So if the light birth rate is 1.5 for example this means that for every adult in the light moth population 1.5 new moths are born. Remember that the light births include some that are actually dark by mutation, which must be subtracted from the light population:

> light mutations = light births * light mutation rate

If the light mutation rate is 0.01 it would mean that 1 in 100 light births is in fact dark.

The number of dark births that are in fact light and that must be added to the light population is:

dark births mutated to light = dark births * dark mutation rate

Finally the number of deaths – by natural causes such as old age – is given by:

light deaths = light survivors * light death rate

Analogous formulas can be written for the dark moth population.

The model is not meant to be scientifically valid. It is meant only as an example of the uses of a spreadsheet.

The Simulation Parameters

As you learned in the previous labs, an important element in designing a spreadsheet model is to place parameters in a separate worksheet, where they can be easily changed. Since formulas link calculations to these parameter values the effects of changing the values can be seen immediately.

In this case the model parameters are:
- The initial populations of light and dark moths
- The predation rates for light and dark moths
- The reproduction rates for light and dark moths (which should be the same)
- The mutation rates for light and dark moths (probably the same)
- The death rates for light and dark moths (probably the same)

Create a new Excel spreadsheet, name the first worksheet Comments (and write what comments you can now, but go back to complete the description after you have finished) and the second one Parameters. Enter titles and values for the parameters as shown in Figure 5.4. This worksheet will also eventually contain a graph of the population evolution.

Figure 5.4 – the Parameters worksheet for the Pepper Moth Simulation

	A	B	C
3		Light	Dark
4	Initial Population	2000000	0
5	Birth Rate	1.5	1.5
6	Death Rate	1	1
7	Predition Rate	0.35	0.35
8	Mutation Rate	0.01	0.01

The Population Evolution Worksheet

Name the third worksheet PopulationGrowth and enter the first three column headings as Light, Dark and Total with a heading of Moth Populations spanning these three columns in very top row of the worksheet.

Enter the formulas in the first cells of the Light and Dark columns to transfer the initial population values from the Parameters worksheet. The values in the Total column will simply be the sum of Light and Dark and you can enter that formula now too.

To calculate the other values in the moth populations columns it will be easiest if you set up other columns to contain each of the values described above – i.e. survivors, births, mutations, deaths for each of the light and dark moths. You can hide these columns later, if you wish, since they are intermediate results that you don't really want to see. The formulas for these quantities are described previously, although you will have to implement the formulas in the syntax of the spreadsheet.

The formula for the light population in the second row (i.e. the second generation) will then simply be

$$= \text{light survivors} + \text{light births} - \text{light mutations} - \text{deaths} + \text{dark mutations}$$

where these values come from the first row. Recall that in a recurrence formula you need to use explicit cell references. The formula for the dark population in the second row will be analogous to that for the light population. Fill the formulas down about 500 rows.

The first few rows row should contain the values shown in Figure 5.5 if you use the parameter values shown in Figure 5.4.

Figure 5.5 – the first few rows of the moth population evolution worksheet

	A	B	C	D	E	F	G	H	I	J	K
1	Moth Populations			Light Moths				Dark Moths			
2	Light	Dark	Total	Survivors	Births	Mutations	Deaths	Survivors	Births	Mutations	Deaths
3	2000000.0	0.0	2000000.0	1300000.0	1950000.0	19500.0	1300000.0	0.0	0.0	0.0	0.0
4	1930500.0	19500.0	1950000.0	1254825.0	1882237.5	18822.4	1254825.0	12675.0	19012.5	190.1	12675.0
5	1863605.3	37644.8	1901250.0	1211343.4	1817015.1	18170.2	1211343.4	24469.1	36703.6	367.0	24469.1

The columns containing the intermediate results (survivors, births etc.) need not be seen – they are somewhat distracting from the main results. Excel and most other spreadsheets provide a way to hide data that need not be seen. Select the columns you want to hide by dragging across the column letter headings at the top of the worksheet. It will be columns D through to K if you have constructed your worksheet as shown in Figure 5.5. You'll find the Column/Hide command on the Home tab, in the Cells group by clicking on Format. (Notice that you can also hide rows and whole worksheets; also observe that there is an Unhide command.)

Creating the Graph

Select the three columns Light, Dark and Total, and click on the Insert tab and select Line from the Charts group. If you include the headings in your selection you'll find that the headings are automatically used for the legend that labels the three different lines that will be in your graph.

Return to Chapter 2 of this manual to review building a multi-line graph.

Next you should "fine tune" the appearance of the graph to resemble Figure 5.6. In particular notice the vertical and horizontal axes in the figure. You should also not rely on colour to distinguish the different lines on the graph since your printer may not be colour.

Figure 5.6 – the Pepper Moth graph
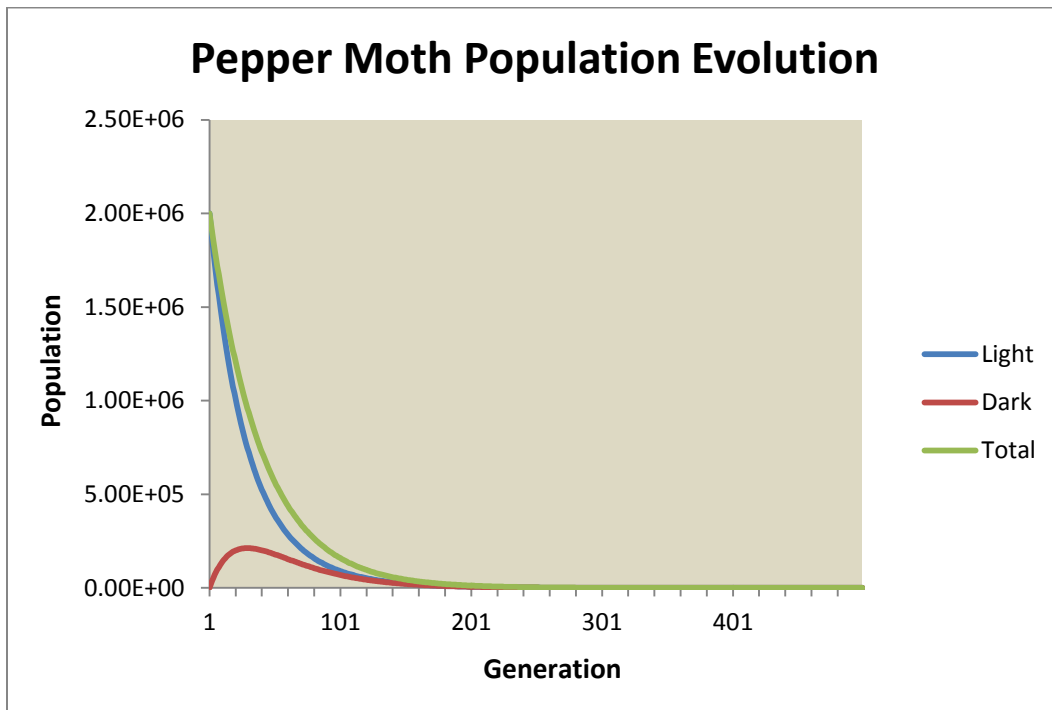
Experimenting with the Simulation

The predation rates for light and dark moths should not be the same, as given in Figure 5.4. The light predation rate should be larger than the dark predation rate, so you will have to change this.

Also, it is not sensible that the moths become extinct as implied by this graph. You should adjust the parameters so that the population reaches an approximate steady state.

Normally populations do not change rapidly – rather it takes many generations for evolution to occur. So the simulation is also unrealistic because drastic changes happen within about 100 generations. Examine which parameters affect the rapidity of the changes in the population.

A word of advice – do not change the parameters by large amounts, and also remember what a parameter was before you changed it so that you can change it back if something unusual happens.

You should aim to end up with a graph somewhat like Figure 5.7.

<u>Figure 5.7 – an almost steady state pepper moth population</u>

Note: the look of this graph is different from the one in fig. 5.6 because
it was not possible to generate a new version at this time.
The information in the graph is otherwise correct.

Summary

Notice that even though it's possible to create a sensible graph, there's no guarantee that our efforts actually model the real world. What we need to do is compare our projections to experimental data. This is, of course, what researchers do all the time. We'll look at an example in the next exercise.

Exercise 5a – **Importing Data**

As discovered at the end of the last exercise, it is desirable to be able to compare the data projected by a simulation with experimental data. In the best case, this data is stored in a machine-readable form so we can avoid the tedious and error-prone process of entering it manually. Since this is a common situation, Excel has a Wizard that facilitates the importing of data.

The file ExperimentalData.txt includes data collected in an experiment conducted by medical researchers in the 1960s. Measurements of 3 different chemicals in the blood were made every 9 minutes, over a period of 300 minutes.

Download the ExperimentalData.txt file from Support Files (Chapter 5) on the course website. Open a new workbook. Click on the File tab and then click on Open and navigate to the folder in which you saved the file. Since Excel normally looks for files that it has created (.xlsx), the ExperimentalData.txt file will not appear in the list. You need to specify that Excel is to look for .txt files. Change the setting of the Files of type: pick list at the bottom of the Open dialog box. You can use any of the settings that include .txt – perhaps All Files (*.*).

Select ExperimentalData.txt and click Open.

The window like that shown in Figure 5.8 will appear.

Figure 5.8 – Import Wizard (1 of 3)
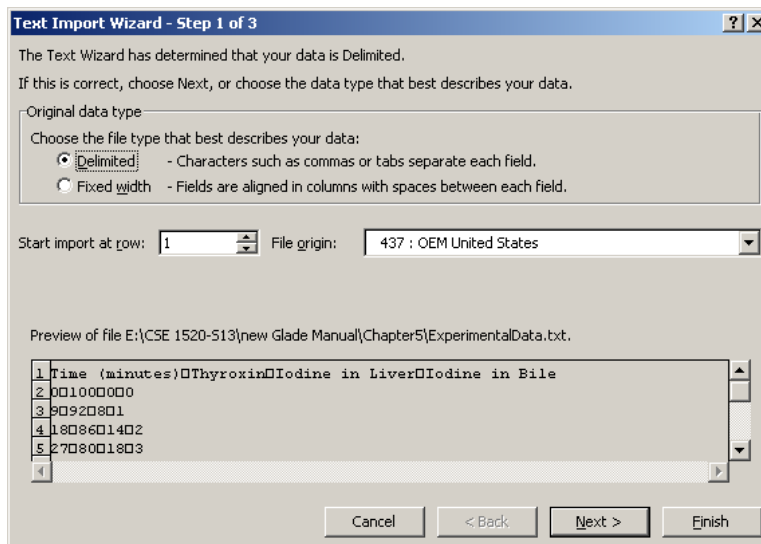
Excel has examined it and presented a Preview of the file. A brief look at the data shown might raise some concerns. There are supposed to be 3 data items for each of several time

values. But that's not what appears to be the case. However, Excel has also determined that the data is **delimited**. This means that a character has been placed between data items to separate them. Click Next to see how Excel organises the data.
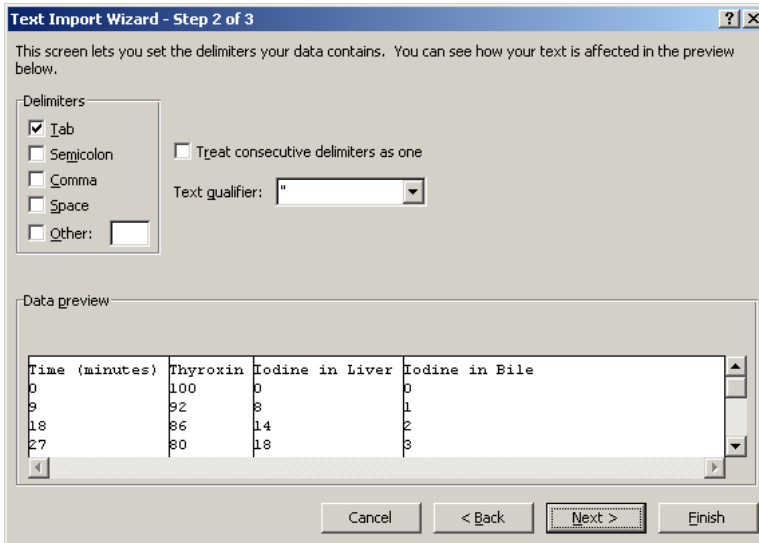
Figure 5.9 – Import Wizard (2 of 3)

Figure 5.9 indicates that Tab characters are present in the data. If these are the **delimiters** then the data will be imported as shown in the Data preview. You should select and deselect other options to explore this step in the Wizard. When you've completed this experimentation set the options back to the way they are in Figure 5.9 and click Next.

Figure 5.10 – Import Wizard (3 of 3)

The third step in importing data offers an opportunity to set formats for the each column of data, or to specify which columns not to import. We want all the columns and the General format is a good choice for all the data so just click Finish. Figure 5.11 shows the appearance of the data after it has been imported and the column headings have been formatted. Save your file. Be sure to save it as an Excel workbook and not a text file.

Figure 5.11 – The imported data

	A	B	C	D
1	Time (minutes)	Thyroxin	Iodine in Liver	Iodine in Bile
2	0	100	0	0
3	9	92	8	1
4	18	86	14	2
5	27	80	18	3
6	36	78	19	4
7	45	80	20	5
8	54	75	20	6
9	63	75	22	6
10	72	73	22	7

ExperimentalData

Exercise 5b - **Thyroxine Conversion in the Liver**

The Model

The object of this exercise is to simulate the functioning of the human liver as it cleans the blood of a chemical that has been injected. The chemical is called thyroxine. The liver can convert this into iodine, and from the liver the iodine is absorbed into the bile. In an experiment a subject is injected with thyroxine and then, over a period of time, measurements are made of the amount of thyroxine in the blood, the amount of iodine in the liver, and the amount of iodine in the bile.

You can see by inspecting the experimental data that as the conversion takes place the amount of thyroxine decreases, the amount of iodine in the liver increases and stays roughly constant, and the amount of iodine in the bile increases.

This process is modeled in the following way. The liver can process a portion of the thyroxine in the blood in a certain amount of time – let's say 3% (0.03) per minute. Thus every minute 3% of the thyroxine would be converted to iodine.

As a formula this is

    new amount of thyroxine = old amount of thyroxine
                                    - old amount of thyroxine* 0.03 * time interval

where new amount of thyroxine means the amount at the end of the time interval and old amount of thyroxine means the amount at the beginning of the time interval. For example, if at the beginning of a time interval the (old) amount of thyroxine was 100 then after a time interval of 2 minutes the (new) amount would be $100 – 100*0.03*2 = 94$. If we use this formula again in the next time interval (the "old" amount is now 94) we have the "new" amount as $94 – 94*0.03*2 = 88.36$. We started with the amount of thyroxine as 100 and after 2 minutes there was 94 and after 4 minutes there was 88.36 and so on.

This is only part of the model. As it converts thyroxine to iodine the liver tends to get saturated with iodine and begins to convert some back to thyroxine (at least we can think of it this way for our purposes) – let's say 7% (0.07) per minute of iodine is converted back to thyroxine. Thus although 3% of thyroxine was removed from the blood stream, 7% of that 3% was converted back to thyroxine, and added back into the blood stream. So as a formula this becomes:

    new amount of thyroxine = old amount of thyroxine
                                    - old amount of thyroxine* 0.03 * time interval
                                    + old amount of iodine in liver * 0.07 * time interval

This means that after a time interval of 2 minutes the new amount of thyroxine would be the old amount (i.e. the amount at the start of the 2 minute period) minus 2*3% of the old amount plus 2*7% of the old amount of iodine in the liver. Of course, to begin, the amount of iodine in the liver is zero (or very small) so the amount of iodine converted back to thyroxine is very small (zero for the first time interval in fact).

Next, let's consider the amount of iodine in the liver. It will increase as the liver converts thyroxine to iodine, but decrease as some is converted back to thyroxine and also as some is absorbed into the bile. By an argument analogous to the above this is represented by the formula:

    new amount of iodine in liver = old amount of iodine in liver
                                    + old amount of thyroxine * 0.03 * time interval
                                    - old amount of iodine in liver * 0.07 * time interval
                                    - old amount of iodine in liver * 0.02 * time interval

The last term subtracted in the formula represents 2% per minute of iodine being absorbed into the bile.

Finally let's consider the amount of iodine in the bile. It is given by the formula:

new amount of iodine in bile = old amount of iodine in bile
+ old amount of iodine in liver * 0.02 * time interval

These three equations represent the amount of thyroxine in the blood, iodine in the liver, and iodine in the bile as they change with time. For example, starting with 100 units of thyroxine, zero iodine in the liver, and zero iodine in the bile we would have:

After 2 minutes –

Amount of thyroxine =	100 – 100 * .03 * 2 + 0 * .07 * 2	= 94
Amount of iodine in liver =	0 + 100 * .03 * 2 – 0 * .07 * 2 – 0 * .02 * 2	= 6
Amount of iodine in bile =	0 + 0 * .02 * 2	= 0

After 4 minutes –

Amount of thyroxine =	94 – 94 * .03 * 2 + 6 * .07 * 2	= 89.2
Amount of iodine in liver =	6 + 94 * .03 * 2 - 6 * .07 * 2 - 6 * .02 * 2	= 10.56
Amount of iodine in bile =	0 + 6 * .02 * 2	= 0.24

After 6 minutes –

Amount of thyroxine =	89.2 – 89.2 * 0.03 * 2 + 10.8 * .07 * 2	= 85.36
Amount of iodine in liver =	10.56 + 89.2 * .03 * 2 - 10.56 * .07 * 2 - 10.56 * .02 * 2	= 14.01
Amount of iodine in bile =	0.24 + 10.8 * .02 * 2	= 0.672

Clearly we could continue this process – and indeed that is exactly what the recurrence formulas in the spreadsheet that you create will do.

The mathematically inclined might recognise these as the finite difference equations representing a system of three simple differential equations:

$$dx_1/dt = -k_{t,i} \, x_1 \; + \; k_{i,t} \, x_2$$
$$dx_2/dt = +k_{t,i} \, x_1 \; - \; k_{i,t} \, x_2 - k_{i,b} \, x_2$$
$$dx_3/dt = \; k_{i,b} \, x_2$$

where x_1, x_2, and x_3 are the amounts of thyroxine in the blood, iodine in the liver, and iodine in the bile respectively, and $k_{t,i}$, $k_{i,t}$, and $k_{i,b}$ are the amounts per unit time of thyroxine converted to iodine, iodine converted to thyroxine, and iodine absorbed into the bile respectively. It is not necessary that you understand these differential equations.

Creating the rest of the Worksheets

Much as you did for the Pepper Moth Population Simulation you will calculate values for the amounts of thyroxine in the blood, iodine in the liver, and iodine in the bile over a period of time using recurrence formulas.

Switch to a new sheet in the model you've begun and write a brief description of what you are about to do in a Comments worksheet.

Name another worksheet Parameters and enter names and values for the initial amounts of thyroxine, iodine in the liver, and iodine in the bile. Use initial values of 100, 0, and 0 respectively. Also enter the transfer coefficients – i.e.
- o   the percent of thyroxine converted to iodine per unit time,
- o   the percent of iodine converted back to thyroxine per unit time,
- o   and the percent of iodine absorbed into the bile per unit time.

The values can be 0.03, 0.07, and 0.02 respectively. Also enter a value for the time interval, which you should set to be 3 (minutes) rather than 2 minutes that we used in the discussion above.

Make sure to arrange things so that there is room in this worksheet for a graph.

Name another worksheet Simulation and set up column titles as shown in Figure 5.12. You should be able to construct recurrence formulas to calculate the Time column by adding the time interval to the time above, and the other three columns by implementing the formulas detailed above. Figure 5.12 shows the first few rows of the Simulation sheet.

Figure 5.12 – part of the Simulation worksheet in the Thyroxine Conversion model

	A	B	C	D
1	Time	Thyroxine	Iodine in Liver	Iodine in Bile
2	0	100	0	0
3	3	91	9	0
4	6	84.7	14.76	0.54
5	9	80.1766	18.3978	1.4256
6	12	76.824244	20.646288	2.529468
7	15	74.24578252	21.9859722	3.76824528
8	18	72.18071626	22.73188013	5.087403612
9	21	70.45814662	23.09053696	6.45131642
10	24	68.96592619	23.19732518	7.836748638
11	27	67.63043112	23.14098074	9.228588148

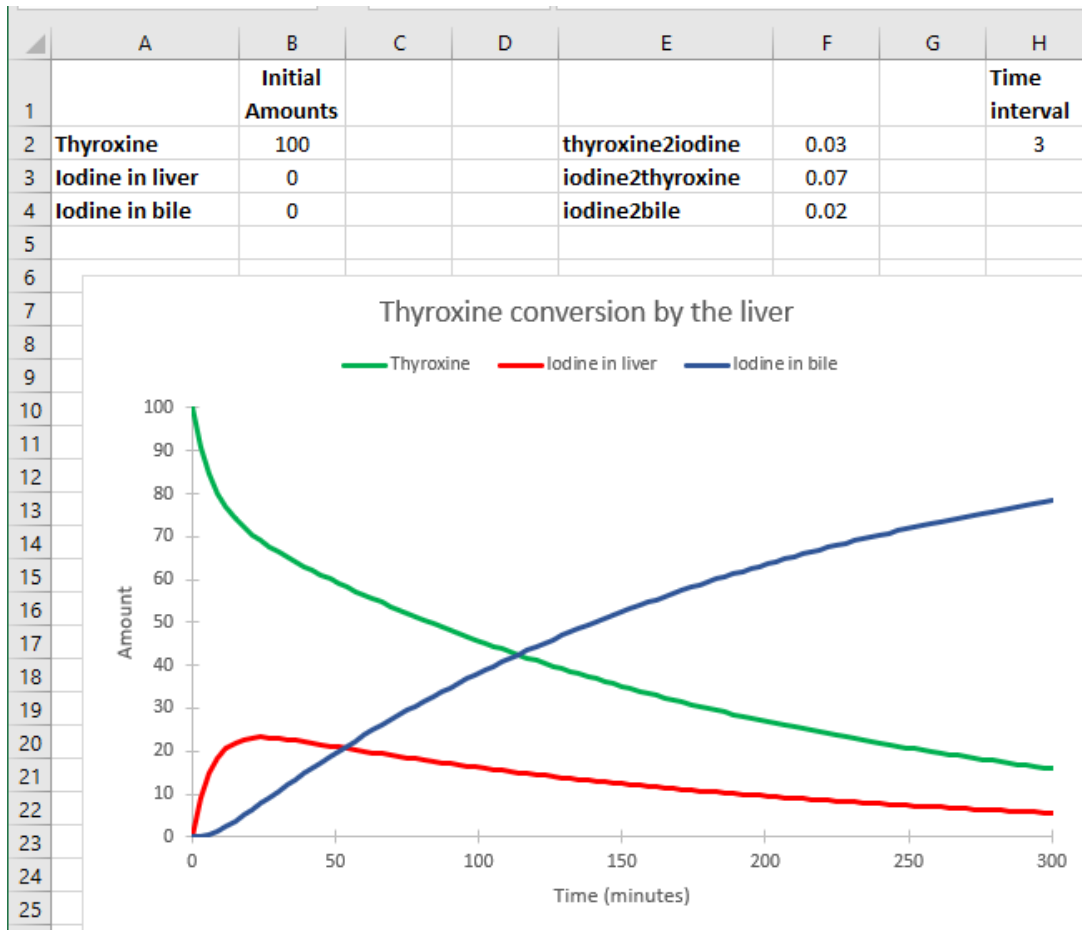Simulation   ExperimentalData

The Graph

The creation of this graph is a little more difficult than previous ones – there are deliberately few instructions here to force you to experiment with the chart building process in Excel and fine-tuning the appearance.

One hint is that you should choose the type as XY (scatter) and then later change the simulation series so that smooth lines join the points.

Figure 5.13 shows what you should aim to end up with.

Figure 5.13 – the graph from the thyroxine conversion simulation

Adjusting the Simulation to fit the Data

As you can see the simulation does not agree very well with the experimental data. This is not necessarily a sign that the simulation model is useless and that it can therefore tell us nothing about the actual physical process that was measured in the experiment. Before we can conclude that the model is invalid, we should try to adjust its parameter values to see if the calculated results can be made to more closely agree with experiment.

The graph in Figure 5.14 shows a better approximation – see if you can match it (or better it?)

So perhaps this model might tell us something after all. The values of the transfer coefficients that give this fit to experiment might mean something, and at least we know that the model itself is perhaps a reasonable way of thinking about how the liver cleans the blood.

Figure 5.14 – a reasonable fit of simulation to experiment

Note: the look of this graph is different from the one in fig. 5.13 because it was not possible to generate a new version at this time. The information in the graph is otherwise correct.