4 Why Use Statistics? Pros, Cons and Alternatives

Chapters 2 and 3 introduced some of the basic concepts of statistics – probability, sampling, methods of summarising patterns of data – and should have given you some of the flavour of statistical reasoning. There are, however, other ways of handling the uncertainties of life. This chapter starts by reviewing some of the key features of the statistical approach, and then looks at the role of some of these alternatives. If you like, you can skip this diversion: the main account of statistics continues in the next chapter.

▶ 4.1 Introduction

Smoking causes lung cancer. The married are happier, on average, than the unmarried.[48] According to recent press reports, the A889 in the Scottish Highlands is the most dangerous road in the UK. Happy workers are more productive than unhappy workers. Patients in different areas of the UK have different survival rates after treatment for cancer. What these all have in common is that they are statistical conclusions. In each case the research works by taking large samples of people, or car journeys, counting up the numbers in various categories (the happily married, the happily unmarried, the unhappily married, the unhappily unmarried and so on) and working out average measurements in these categories or something broadly similar.

We are living in the age of statistics. Every morning the news seems to include another result of statistical analysis, with clear implications for how we should live our lives: stop smoking, get married, avoid the A889 and so on. As large-scale data collection and analysis becomes easier, this trend is perhaps inevitable. The scale is important; in earlier times, informal statistics on, say, disease would have been noted on a local level, but there would have been far too little data to see anything but the crudest patterns.

As explained in Chapter 1, my approach in this book is to use the model of a bucket with some balls in it to describe many aspects of statistics. We can think of statistical methods as bucket and ball methods. For example, we can visualise the people in a survey to investigate the relation between marriage and happiness as balls with the relevant information (marital and happiness status) stamped on them. One person may be represented by a ball stamped with the words ‘married’ and ‘happy’; another’s ball might be stamped ‘unmarried’, ‘unhappy’. The statistician is then in a position to start counting balls to compare the various groups. There might be other bits of information stamped on the balls as well, age, for example. The bucket and ball model has the advantage of drawing attention to some of the characteristic features of the statistical approach. Let’s see what these are.

4.2 Characteristic features of the statistical approach

The first feature is that all balls have the same status. If we take a group of people – each represented by a ball – and work out their average age, or the proportion who are happy, each person in the group has the same input to the answer.

The second feature is that there is only a limited amount of information stamped on each ball (values of number and category variables) which allows us to compare balls and summarise the properties of a group. We might stamp on each ball ‘married’ or ‘unmarried’ and the age of the person. However, names are irrelevant, as are all the little details that would help us to get to know these people ‘as people’. Statistics treats people as if they were balls in a bucket. This is not, of course, to deny that people are people, but statistics deals with a much simplified view of reality. If, for example, we had two balls in a bucket, both stamped with ‘married’, ‘happy’, ‘36’, these two balls are *identical* from the point of view of a statistical analysis. There is no difference between them. Obviously they represent different people with different personalities and interests, but that’s another story.

The third feature is the fact that statistics then works by a mechanical process of counting the various balls in the bucket. There are rules for finding probabilities, averages or standard deviations which can be performed by a computer. Once we’ve decided which statistics we’re going to work out, we just follow the rules. Judgement is, of course, needed to decide on the starting point – which data to use and how to build our bucket and ball model – and interpret our results. But the statistical analysis itself is a mechanical, ball-crunching process.

What’s stamped on the balls is not always true. The data in the file `drink.xls` (Section 3.1) includes one female who apparently drank 140 units of alcohol over three days. This is difficult to believe, as anyone drinking this much would probably be seriously ill, if not dead. It is entirely possible, and in my view very likely, that the student in question just made these figures up. Or she may have been so drunk she couldn't remember how much she drank. Either way, we can't assume the data is true. The balls in the bucket are just a model of reality; this model may not correspond exactly with reality. Doing statistics with balls in a bucket, and not the people themselves, should help to remind you that statistics is a story which may not always correspond exactly with reality.

▶ 4.3 What sort of conclusions does statistics provide?

The result of this mechanical ball crunching is that we get to know about tendencies or probabilities – what is generally true, true on average or most probable. We find out that male students drink more than female students on average, or that the married are happier than the unmarried, on average. But, of course, there are some females who drink more than some males, and some unmarried people who are happier than some married people.

Care is needed when using statistical averages. Suppose the average (mean) number of children born to women in a country is 2.3. Is this a helpful piece of information for a study of population trends**?**

Yes, it obviously is. Would it be sensible to take it as indicating the typical family for a study of the experience of women in this country**?**

No, for equally obvious reasons: no woman has 2.3 children! Furthermore, the experience of women with no children is likely to be very different from the experience of women with (say) two children. And the experience of both will be very different from that of women with ten children. You need to be very careful about using statistical conclusions which 'average' over very diverse scenarios.

Sometimes we are just concerned about a single person or a single event. We then need to use statistical information about general tendencies or probabilities to come to the best conclusion we can. We might decide that the probability of fine weather on 13 June (see Section 2.1) is 68%, or that we are more likely to be happy if we get married. We can't be certain, of course. The weather and happiness are too complex to be certain. But the statistical approach still gives us useful information.

Usually, the precise causes behind these tendencies or probabilities is to some extent unknown. We may have an idea why smoking causes cancer, or why the married tend to be happier than the unmarried, but we probably do not understand the exact reasons. By taking large numbers of cases, statistics can demonstrate things whose detailed working remains mysterious, and which could never be recognised by studying individual cases. Sometimes this ability of statistics to see the general truth seems magic, surpassing what ought to be possible. We are told, on the basis of a public opinion poll, that 35% of the electorate will vote Labour, *despite the fact that we have not been asked*. How can the statistician tell without asking everyone eligible to vote? In fact there are likely to be errors but the size of these errors can be estimated with surprising accuracy (see Chapter 7).

Of course, if we have a sufficient understanding of a situation we don't usually need statistics. Jumping out of tall buildings is obviously likely to harm people. We know enough about gravity and human bodies for this to be entirely predictable, so there is no point in a statistical survey. But the health effects of eating meat or drinking red wine are less easily predicted, so a statistical analysis which enables us to spot the general tendency, despite all the other influences involved, is very helpful.

There is another side to statistics; it can be used to decide which of a number of possible hypotheses is true, or to assess the likely error in a numerical measurement. This side of statistics is introduced in Chapter 6, where we will use as an example the problem of deciding whether a volunteer in a telepathy experiment, Annette, is actually telepathic or not (see Section 6.2). There are two hypotheses here: either Annette is telepathic or she isn't. Analysing this statistically is a little more complicated (see Chapter 6). In one approach to the problem (Section 6.4), we envisage a bucket with balls representing *possible worlds*; we might conclude that balls representing worlds in which Annette is telepathic are much more numerous than those representing worlds in which she is not telepathic, so she probably is telepathic. Another approach is explained in Section 6.6. In both approaches, the conclusions depend on probabilities: we can never be 100% sure which hypothesis is true.

Statistics has a reputation, in some quarters, for being boring. There are a number of possible reasons for this. It's usually viewed as a mathematical subject and mathematics is sometimes seen as boring; I hope my approach in this book goes some way towards dealing with this. However, the perspective here may also not have immediate appeal. A simplified view of reality in which people (or whatever) are reduced to a few crude pieces of information stamped on some balls in a bucket may strike you as a rather boring attitude to the world. The advantage, of course, is that this simplification enables you to see through the fog of detail and see that, for example, smoking does cause cancer.

There is also the problem that the statistical results are sometimes not quite real. We are told that in the 1930s 'urban air pollution probably killed at least 64,000 extra people each year in the UK',[49] but we can never be sure who these people were. Some people who died of respiratory disease would have died even if the air had been less polluted, whereas others would not, but we cannot be sure which category a particular person is in. They are balls in a bucket, not real people. This makes the problem more difficult to identify with than, say, the people killed by the events of 11 September 2001. In this case, journalists can get our attention and sympathy by telling us stories of individual people, which is far more difficult with air pollution. But in the long run, air pollution is more important.

▶ 4.4 What are the alternatives to statistics?

Attitudes to statistics tend to be polarised. There are those who think that statistics is the answer to all problems and those who think it is never any help. The most helpful attitude falls somewhere in the middle. Statistics is very useful, but there are other ways in which we can deal with the uncertainties of life. I'll mention some of these possibilities briefly in this section, and then discuss three of them at slightly greater length in the next three sections.

Science – proper science in the sense of physics, chemistry or biology – tries to derive laws which explain exactly how things work. If it's successful, of course, then we don't need statistics. If we know enough about happiness to be able to deduce how happy people will be from their circumstances, then we have no need for a statistical analysis. Statistics is just the last resort when we don't understand things properly. So the first alternative to statistics is to try and understand things in the manner of science. If you know anything about research in management, education or the social sciences, you will probably be surprised by the idea of science as an alternative to statistics. Statistics tends to be seen as the scientific approach. It is true that natural science uses statistics to analyse measurement errors, and some theories are formulated in terms of statistics. But there is nothing statistical about $E = mc^2$, it is supposed to be exactly true everywhere, with no statistical qualifications.[50] Statistics should *not* be identified with the scientific approach: some approaches to science may be statistical but many are not.

There are several other approaches to analysing uncertainty in numerical terms besides probability (which leads on to statistical methods). These include 'surprise theory', 'certainty factors' and several others.[51] 'Fuzzy logic' is the only one of these to have made much of an impact and is the subject of Section 4.6. 'Chaos' is another term with a high profile. This refers to random behaviour in systems which are obeying explicit, deterministic (non-probabilistic) laws. You might think this is a contradiction in terms – if it's obeying laws, it's not random – but, as we'll see in Section 4.7, this is not always quite true. Finally, and perhaps most obviously, we might study particular cases, or individual people, as an alternative to statistics. Journalists do this when they want to bring, say, a famine, or the attacks on 11 September, to life for their audience. As well as reporting the statistics of the famine, they will also follow the stories of individuals to show what the famine really means. But do such anecdotes have a place in the serious business of research?

4.5 Anecdotes about specific cases

The problem with taking too much notice of special cases is that they may not be typical. They may not reflect the broader picture. It is always tempting to generalise from a few personal experiences. If you happen to know a couple of female students who drink too much, and a male student who is a teetotaller, you may well not believe that male students drink more than female students. Studies of the way people assess probabilities intuitively have found that it is very common for people to be misled into giving too much weight to personal experiences and high-profile examples.[52] This is not entirely surprising: the antidote is, of course, to collect and analyse the data systematically.

Special cases, then, should not be given too much weight, statistically speaking. But that's not to say they are of no value. For example, there are people who have a natural immunity to the HIV virus which causes AIDS.[53] From a broad statistical point of view they may be a negligible minority, but it would obviously be interesting to know why they are immune, because this may lead to possible ways of helping other patients. Similarly, total quality management is an approach to managing a business which is said by the gurus to bring tremendous benefits. The trouble is that there is widespread scepticism about whether it actually works. Now imagine that an organisation were to be found where total quality management was, indisputably, working well and bringing the anticipated benefits. Obviously, this organisation would be worth studying in depth to see exactly how total quality management had worked its wonders, in the hope that something would be learned which could be helpful to other organisations.

In both examples, the study of the particular case is worthwhile because it illustrates *possibilities* which may be of wider interest.[54] There is no guarantee of this, and certainly no statistical conclusions are possible. But the study of these particular cases may help us to learn about possibilities of which we may otherwise have been unaware. Obviously the cases should be chosen carefully; they may be of special interest like the two examples above, or they may be chosen because they seem typical in some sense. This principle is also accepted as an adjunct to standard statistical practice. You should always check data for outliers that do not fit the pattern, and then check to see if you can see what is going on. For example, the girl who said she drank 140 units in three days (Section 3.2) was probably lying, but if not, the reasons for her exceptional tolerance might be medically interesting.

Statistical process control (SPC) is a set of statistical methods used for monitoring business and industrial processes. A large part of the purpose of SPC is to identify unusual situations so that they can be investigated more fully. For example, many industrial processes are monitored statistically so that if something goes wrong the problem will be picked up by the statistics because it does not fit the normal pattern. Then the people in charge can investigate to see what happened and take appropriate action. The statistics here are used to distinguish the unusual cases so that they, and only they, can be explored further.

▶ 4.6 Fuzzy logic

Suppose you were asked to find out whether it is true that students drink too much on Saturday night. We've already looked at some data on this question in Chapter 3. Looking at the data in Table 3.1, would you say that students do drink too much on Saturday night? Yes or no**?**

You might say that you can't answer this because the question is too vague, in several different senses. The first type of vagueness is due to the fact that some students do undoubtedly drink too much (26 units is 13 pints of beer which cannot be healthy), whereas others drink nothing, which is obviously not too much. We could plot a histogram (Figure 3.1) to show this variation. This is what statistics is about: we can analyse this using concepts like probability, averages and measures of spread.

There are, however, other types of vagueness. The most important is *how much* does a student have to drink for it to count as too much? We could take the UK government guideline of 3 units a night, which means that we would draw a very sharp distinction between 'too much' (4 or even 3.5 units) and 'OK' (3 units). On the other hand, we might feel it makes more sense to use a continuous scale from 2 units – definitely OK – to 10 units – definitely too much. This is the idea behind fuzzy logic. Instead of treating statements as true or false (a so-called 'two valued logic'), we use a continuous scale from 0 – definitely false – to 1 – definitely true. What 'truth value' would you give to the assertion that a student who had 6 units had drunk too much**?**

The obvious answer is 0.5, on the grounds that 6 units is midway between 2 units (truth value of 0) and 10 units (truth value of 1). The relationship I am assuming between units drunk and truth value is shown in Figure 4.1. However, I can't really say that 0.4 or 0.6 is wrong. And why did we choose

Figure 4.1 Truth values for 'drinks too much'

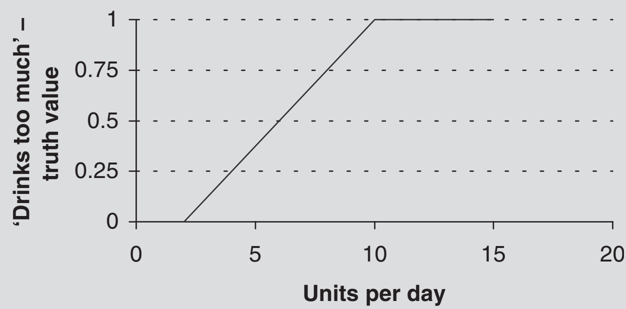

2 units for the zero of the truth value and 10 units for the top of the scale? Inevitably it's a bit arbitrary. It's based on my judgement and different people might have different views. Fuzzy logic is itself a bit fuzzy.

Kosko's book *Fuzzy thinking*,[55] puts the case for avoiding sharp true/false or black/white distinctions, and using a fuzzy scale wherever possible. And in many contexts he certainly has a point. I was listening to a pundit being asked on the radio whether men were now able to cook and look after themselves at home. The expected answer was yes or no. The real answer would acknowledge the fact that men vary, and that household skills need to be assessed on a fuzzy or sliding scale. It's not just a matter of competent or incompetent, but rather there are varying degrees of competence. And, of course, we need statistics to cope with the fact that men are not all the same.

In the black and white world of true and false, we might have a rule which says that *if* the student has drunk too much *and* the route home is dangerous, *then* the student should be restrained from going home. The difficulty here is that 'drunk too much' and 'dangerous' are fuzzy quantities. How much too much has the student drunk? How dangerous is the route home? Is it a moderately busy road, or does it involve a path along the top of a precipice? Obviously we need to take this into account, and fuzzy logic is one way of doing this. Fuzzy logicians would then set up a rule to relate the two truth values to the recommendations about whether we should stop the student from going home. It is an antidote to the black and white reasoning of true and false.

In practice, only logicians really need this antidote. When doing a survey, it is common practice to ask people, for example, how much they enjoy their work on a five-point scale ranging from 1 representing 'not at all enjoyable' to 5 representing 'very enjoyable'. Then a score of 3 would represent some sort of moderate enjoyment. Similarly, I asked my students (see Chapter 3) how much they drank, rather than simply whether they drank. Both are more useful than asking yes/no questions. And instead of simply asking whether it will rain or not on the day of a fete (Section 2.1), it might be more useful to ask the fuzzier question of how good the weather will be.

Fuzzy scales like these should not be confused with probability. Saying somebody's enjoyment of work is 3 on a 0 to 5 scale – or 60% – is obviously different from saying that if you were to encounter a random worker, the probability of their enjoying their work is 60%. It's important not to confuse fuzzy logic with probability. Often we have statistical questions about fuzzy variables. Unfortunately, some statisticians like to answer these questions in black and white, true or false terms. The conclusion may be that the hypothesis is false, instead of a fuzzier, but more useful conclusion about how much truth there is in it (see Chapter 8, especially Sections 8.4 and 8.6.3).

▶ 4.7 Chaos

The words 'chaos' and 'chaotic' have rather different meanings in mathematics and everyday life. In ordinary language if we talk of chaos we mean confusion, things being arranged with no apparent pattern. This is similar to the idea of randomness: a random sample is one chosen without any reference to the patterns in the data. In mathematics, chaos has a more restricted meaning: 'stochastic behaviour occurring in a deterministic system'.[56] Let's see what this means.

A 'deterministic' system is one which is determined by exact rules which can be used to predict what will happen next. For example, the movement of the earth, moon and sun are determined by mechanical forces which are described (with almost complete accuracy) by the laws of motion devised by Isaac Newton. These laws can be used to predict where the earth, moon and sun are at any given time. They were used to predict the onset of a total eclipse of the sun in Penzance, Cornwall at 37 seconds past 11.11 am on 11 August 1999. This was predicted years in advance with almost complete accuracy. This is a deterministic system; if we know where the sun, earth and moon are now, and how fast they are going and in which direction, we can predict eclipses for hundreds of years.

'Stochastic' behaviour, on the other hand, is random behaviour, or behaviour which can only be described in terms of probabilities. At first sight, there would seem to be no place for randomness in deterministic systems, because we ought to be able to use the laws of the system to work out exactly what will happen. There is, however, often a practical problem which means that this is not true.

Imagine that you have been given a rather unusual type of bank account. Instead of working out the balance in the account every day by adding up what has been paid in and what has been taken out, the balance is £1250 on 1 January, and thereafter it is worked out by a simple rule. Every day, the bank applies the rule and works out the new balance. This is a deterministic system: everything is decided by the rule, with nothing being left to chance.

Table 4.1 shows some daily balances for four of these accounts a few months after they were opened. Three of the four are based on a deterministic rule, but the bank has cheated for the fourth which is based on a rule using random numbers. Can you tell which account is based on random numbers**?**

It should be obvious that the first is not random and the second has an obvious cycle. This leaves the last two, both of which look fairly random. In fact, the third is deterministic, and the fourth is random. Despite the fact that

Table 4.1 Balances in four bank accounts opened on 1 January

Date	*Account 1*	*Account 2*	*Account 3*	*Account 4*
20 May	£268	£970	£1746	£1546
21 May	£268	£2	£1112	£412
22 May	£268	£1744	£25	£230
23 May	£268	£970	£1901	£325
24 May	£268	£2	£1623	£1596
25 May	£268	£1744	£775	£1286
26 May	£268	£970	£101	£38
27 May	£268	£2	£1616	£470
28 May	£268	£1744	£758	£67
29 May	£268	£970	£117	£274
30 May	£268	£32	£1559	£1885
31 May	£268	£1744	£624	£1026
01 Jun	£268	£970	£282	£1776
02 Jun	£268	£2	£1030	£1274
03 Jun	£268	£1744	£2	£1610
04 Jun	£268	£970	£1993	£24
05 Jun	£268	£2	£1972	£1850
06 Jun	£268	£1744	£1889	£60
07 Jun	£268	£970	£1581	£946
08 Jun	£268	£2	£676	£1007
09 Jun	£268	£1744	£210	£1423

All balances rounded off to the nearest whole number in the table, although the next day's balance is worked out from the *unrounded* number.

it is deterministic, it is difficult to see a pattern in Account 3: no values are repeated and there is no obvious cycle. There is some sort of weak pattern: every value of £1000 or more is followed by a lower number, which would probably not be true if the sequence were genuinely random. However, this pattern does not enable you to predict, say, five days in advance.

Imagine now that you have the option of taking the money from one of these accounts, but you have to name the day five days in advance. How would you do this for each of the accounts? Can you predict the likely balance in each account on 14 June**?**

This is easy for the first, the balance will be £268 whenever you choose to take it. In the second the predictable cycle makes it almost as easy; my prediction would be £2. The last two, however, are more or less unpredictable. According to the spreadsheet on which Table 4.1 is based, the balance in Account 3 on 14 June is £356, and in Account 4, £412. But I'm sure you did not manage to predict this! The rule the bank uses for Accounts 1–3 is actually very simple:

1. Divide yesterday's balance by 1000.
2. Subtract 1.
3. Multiply the answer by itself, remembering that if you multiply a negative number by itself the answer is positive.
4. Multiply the answer by 500 for Account 1, by 1750 for Account 2, and by 2000 for Account 3. This is the new balance.

To take an example, for Account 3 on 2 January, step 1 gives 1.25 (the balance on 1 January was £1250), step 2 gives 0.25, step 3 gives 0.0625 and step 4 gives the new balance of £125. Applying this process to this balance (0.125, –0.875, 0.765625, 1531) gives the balance for 3 January. And so on. There is no randomness involved, just explicit, arithmetical rules. It is easy to set this up on a spreadsheet, as I did to generate Table 4.1.[57]

And yet the result for Account 3 looks pretty random. A seemingly simple rule has resulted in chaos. But two slightly different rules, differing only in step 4, result in the more predictable patterns you might expect from an arithmetical rule. The clue to what is going on here is to look at what happens if we change the opening balance on 1 January slightly. Let's imagine the bank puts in an extra pound to start us off (£1251 instead of £1250). This shouldn't make too much difference, or will it? For Account 1 the change makes no difference after the first few days. The account seems destined to settle down to £268. For Account 2, the balance on 9 June (day 160) changes from £1744 to £2. There is, however, much the same cycle of three values, the balance the next day becomes £1745. The changed balance has just shifted the cycle by a day.

For Account 3, however, the balance on 9 June changes from £210 to £419. Even smaller changes have a massive impact by the time we get to 9 June. For example, a balance on 1 January of £1250.01 leads to £1975 on 9 June, £1250.02 leads to £751, and £1250.03 leads to £640. Tiny changes in the starting conditions get magnified to enormous changes later on. This leads to some oddities. For example, on 25 January the balance is £1982, on the 26th it is £1928, and it has returned to £1982 by 30 January. Does this mean that it is now starting a five-day cycle and must return to £1982 after another five days on 4 February**?**

No! On 31 January the balance is £1930, slightly different from the £1928 on the 26th, and by 4 February the difference has got bigger, the balance being £1961. The point is that the balance on 25 January is not *exactly* the same as on 30 January, and the difference gets magnified as the rule is applied every day.

If you have tried to reproduce the balances in Account 3 on your own computer, you will probably have obtained different answers. The spreadsheet on which Table 4.1 is based was not Excel; when I transferred the worksheet to Excel, the answers changed yet again. The balance for Account 3 on 9 June is £1917 instead of £210. Why do you think this is**?**

Most fractional numbers cannot be expressed exactly as decimals; they have to be rounded off to a particular number of decimal places. The difference is that each spreadsheet stores a different number of decimal places (for example 0.333 versus 0.333333); this means that the numbers stored are slightly different and these differences get magnified in the way we have just seen. Whenever the laws of a system mean that the final result is very sensitive to the initial conditions, this sort of *chaotic* behaviour may result.

One of the standard examples of chaotic behaviour is the weather. The atmosphere obeys the laws of physics so, if we were clever enough, we should be able to use information about what the atmosphere is doing now to predict what it will be doing in a second's time, and use this to work out what it will be doing in the next second, and so on, for ever. This is roughly what weather forecasters use their computer for. But it doesn't work for more than a few days, perhaps up to week. After that the inevitable inaccuracies in the data on which the computer works means that the forecasts get far less reliable. The problem is that small differences in the initial conditions may get magnified. The problem is far, far worse when predicting weather than when predicting the bank balances above, of course. With the bank balance there is just one initial condition, the initial balance. With weather forecasting, there will be millions of initial conditions, corresponding to the condition of the atmosphere all over the world. Needless to say, millions of initial conditions are likely to produce more of a headache than one.

The problem has been christened the butterfly effect.[58] A butterfly's wing produces a tiny change in the atmosphere. The atmosphere is now different from the way it would have been without the butterfly. These differences may get magnified – just like the penny extra on the balance above – and mean that the difference between the world with the butterfly and the world without it may be the difference between a hurricane on the other side of the world and no hurricane. Perhaps the hurricane which devastated southern England in 1987 would not have happened if a particular butterfly somewhere had not flapped its wings. Nobody knows if this is so; the weather is chaotic and unpredictable.

So where does this leave us? The main point to bear in mind is that understanding that something is chaotic does not solve many problems. Probabilities and statistics are still useful for weather forecasting. Chaotic events are random,[59] so statistics will probably be helpful. The mathematical theory of chaos[60] may also help. It may tell us which sort of values to expect, whether we will get chaos (like Account 3), or more predictable patterns like Accounts 1 and 2. And if we have got a chaotic situation, understanding the deterministic laws may help to some extent. For example, there is some sort of pattern in Account 3: the balance seems to go down in small steps, then suddenly jump again (but never to more than £2000), and then decline in another series of steps. This pattern could be helpful for predicting tomorrow's balance, but not next year's.

Which brings me back to that eclipse. I went to Cornwall to try to see it. The forecast of the time of the eclipse was, of course, correct. The weather, despite being a deterministic system, was chaotic and could only be forecast probabilistically. The forecast the day before was that the sky would probably be overcast, in which case the eclipse would not be visible. In the event the sky was overcast!

▶ 4.8 Exercise

4.8.1 Experimenting with chaos

Use a spreadsheet to calculate the balance in the bank accounts described in Section 4.7. Do you get the same results as in Table 4.1? Can you see any way of forecasting the balance in Account 3? Now try experimenting with different numbers in step 4 of the rule. Which numbers give you chaotic results?

▶ 4.9 Summary of main points

- Statistical methods enable you to see general patterns and tendencies, which provide an invaluable perspective in many walks of life. Statistics

often allow you to see through a fog of detail to truths which may not otherwise be apparent.
- It is often tempting to focus on particular cases, perhaps those which are familiar, accessible or have a high profile. This has its place for demonstrating what is possible, but may lead to misjudging the overall picture.
- There are many other approaches to handling uncertainty, two of which, fuzzy logic and the idea of chaos, are briefly reviewed in this chapter.