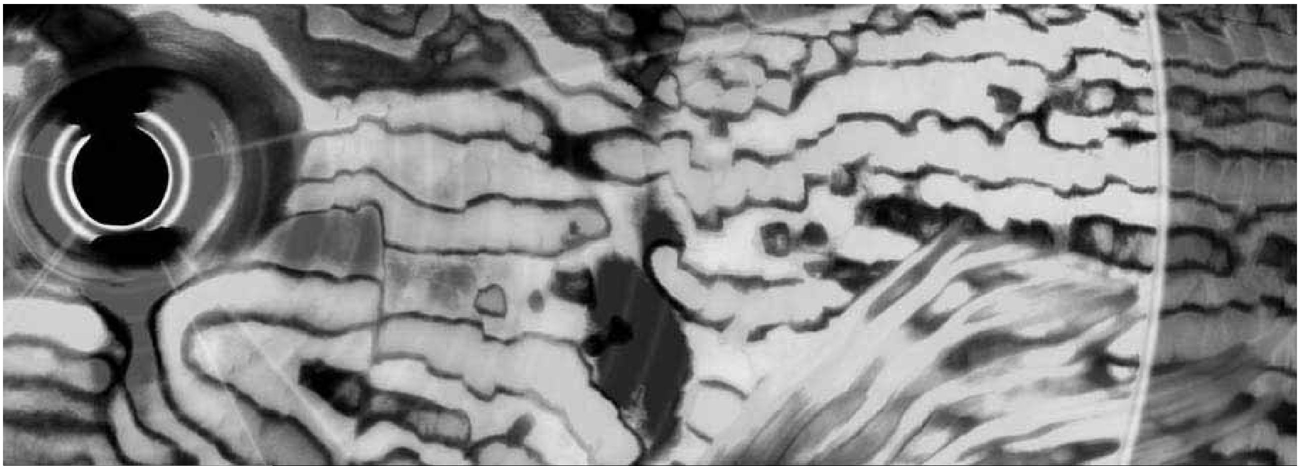

1 Doing science

Where do questions come from?

You're out for a walk one autumn afternoon when you notice a squirrel picking up acorns under some trees. Several things strike you about the squirrel's behaviour. For one thing it doesn't seem to pick up all the acorns it comes across; a sizeable proportion is ignored. Of those it does pick up, only some are eaten. Others are carried up into a tree, where the squirrel disappears from view for a few minutes before returning to the supply for more. Something else strikes you: the squirrel doesn't carry its acorns up the nearest tree but instead runs to one several metres away. You begin to wonder why the squirrel behaves in this way. Several possibilities occur to you. Although the acorns on the ground all look very similar to you, you speculate that some might contain more food than others, or perhaps they are easier to crack. By selecting these, the squirrel might obtain food more quickly than by taking indiscriminately any acorn it encountered. Similarly, the fact that it appears to carry acorns into a particular tree suggests this tree might provide a more secure site for storing them.

While all these might be purely casual reflections, they are revealing of the way we analyse and interpret events around us. The speculations about the squirrel's behaviour may seem clutched out of the air on a whim, but they are in fact structured around some clearly identifiable assumptions, for instance that achieving a high rate of food intake matters in some way to the squirrel and influences its preferences, and that using the most secure storage site is more important to it than using the most convenient site. If you wanted to pursue your curiosity further, these assumptions would be critical to the questions you asked and the investigations you undertook. If all this sounds very familiar to you as a science student, it should, because, whether you intended it or not, your speculations are essentially scientific. Science is simply formalised speculation backed up (or otherwise) by equally formalised observation and experimentation. In its broadest sense most of us 'do science' all the time.

1.1 Science as asking questions

Science is often regarded by those outside it as an open-ended quest for objective understanding of the universe and all that is in it. But this is so only in a rather trivial sense. The issue of objectivity is a thorny one and, happily, well beyond the scope of this book. Nevertheless, the very real constraints that limit human objectivity mean that use of the term must at least be hedged about with serious qualifications. The issue of open-endedness is really the one that concerns us here. Science is open-ended only in that its directions of enquiry are, in principle, limitless. Along each path of enquiry, however, science is far from open-ended. Each step on the way is, or should be, the result of refined question-asking, a narrowing down of questions and methods of answering them to provide the clearest possible distinction between alternative explanations for the phenomenon in hand. This is a skill, or series of skills really, that has to be acquired, and acquiring it is one of the chief objectives of any scientific training.

While few scientists would disagree with this, identifying the different skills and understanding how training techniques develop them are a lot less straightforward. With increasing pressure on science courses in universities and colleges to teach more material to more people and to draw on an expanding and increasingly sophisticated body of knowledge, it is more important than ever to understand how to marshal information and direct enquiry. This book is the result of our experiences in teaching investigative skills to university undergraduates in the life sciences. It deals with all aspects of scientific investigation, from thinking up ideas and making initial exploratory observations, through developing and testing hypotheses, to interpreting results and preparing written reports. It is not an introduction to data-handling techniques or statistics, although it includes a substantial element of both; it simply introduces these as tools to aid investigation. The theory and mechanics of statistical analysis can be dealt with more appropriately elsewhere.

The principles covered in the book are extraordinarily simple, yet, paradoxically, students find them very difficult to put into practice when taught in a piecemeal way across a number of different courses. The book has evolved out of our attempts to get over this problem by using open-ended, self-driven practical exercises in which the stages of enquiry develop logically through the desire of students to satisfy their own curiosity. However, the skills it emphasises are just as appropriate to more limited set-piece practicals. Perhaps a distinction – admittedly over-generalised – that could be made here, and which to some extent underpins our preference for a self-driven approach, is that with many set-piece practicals it is obvious *what* one is supposed to do but often not *why* one is supposed to do it. Almost the opposite is true of the self-driven approach; here it is clear why any investigation needs to be undertaken, but usually less clear what should be done to see it through successfully. In our experience, developing the ‘what’ in the context of a clear ‘why’ is considerably more instructive than attempting to reconstruct the ‘why’ from the ‘what’ or, worse, ignoring it altogether.

1.2 Basic considerations

Scientific enquiry is not just a matter of asking questions; it is a matter of asking the *right questions* in the *right way*. This is more demanding than it sounds. For a start, it requires that something is known about the system or material in which an investigator is interested. A study of mating behaviour in guppies, for instance, demands that you can at least tell males from females and recognise courtship and copulation. Similarly, it is difficult to make a constructive assessment of parasitic worm burdens in host organisms if you are ignorant of likely sites of infection and can't tell worm eggs from faecal material.

Of course, there are several ways in which such knowledge can be acquired: e.g. the Internet/World Wide Web, textbooks, specialist academic journals (many of which are now available electronically through licensed subscribers like universities and colleges), asking an expert, or simply finding out for yourself through observation and exploration.

These days, the first choice for browsing information is often the Internet/World Wide Web. The advantages of such 'online' searching in terms of speed and convenience hardly need detailing here, but there *are* dangers, as we indicate later. A good way of accessing reliable scientific information like this is to use one or more of the professional Web-based literature databases, such as the Web of Knowledge (**wok.mimas.ac.uk**), PubMed (**www.ncbi.nlm.nih.gov/pmc/**) Google Scholar (**scholar.google.co.uk**) or BIOSIS (**science.thomsonreuters.com/training/biosis/**). These search the peer-reviewed (and therefore quality-controlled) academic journals for articles containing information relevant to your request. Each of these provides tips on how best to use them, but a handful of basic ones is given in Box 1.1.

Whichever mode of acquiring information is preferred, however, a certain amount of background preparation is usually essential, even for the simplest investigations. In practical classes, some background is usually provided for you in the form of handouts or accompanying lectures, but the very variability of biological material means that generalised and often highly stylised summaries are poor substitutes for hard personal experience. Nevertheless, given the inevitable constraints of time, materials and available expertise, they are usually a necessary second best. There is also a second, more important, reason why there is really no substitute for personal experience: the information you require may not exist or, if it does exist, it may not be correct. The Internet/World Wide Web is a particular hazard here because of the vast amount of unregulated information it makes available, often dressed up to appear professional and authoritative. *Such material should always be treated with caution and verified before being trusted*. Where academic information is concerned, a first step might be to check the host site to see whether it is a recognised institution, like a university or an academic publisher; another might be to look for other research cited in the information, for instance in the form of journal citations (*see* section 4.3.1), which can be cross-checked. Entering the author's name into the search field of one of the web-based professional literature databases (Box 1.1) to see whether this person has a published research track record can be another approach.

Box 1.1 Searching online literature databases

Searchable online literature databases, like the Web of Knowledge, BIOSIS or PubMed, allow you to search for articles by particular authors, or on particular topics, or according to some other category, such as a journal title or research organisation. An example of the kind of search fields on offer, in this case for the Web of Knowledge, is shown in Fig. (i).

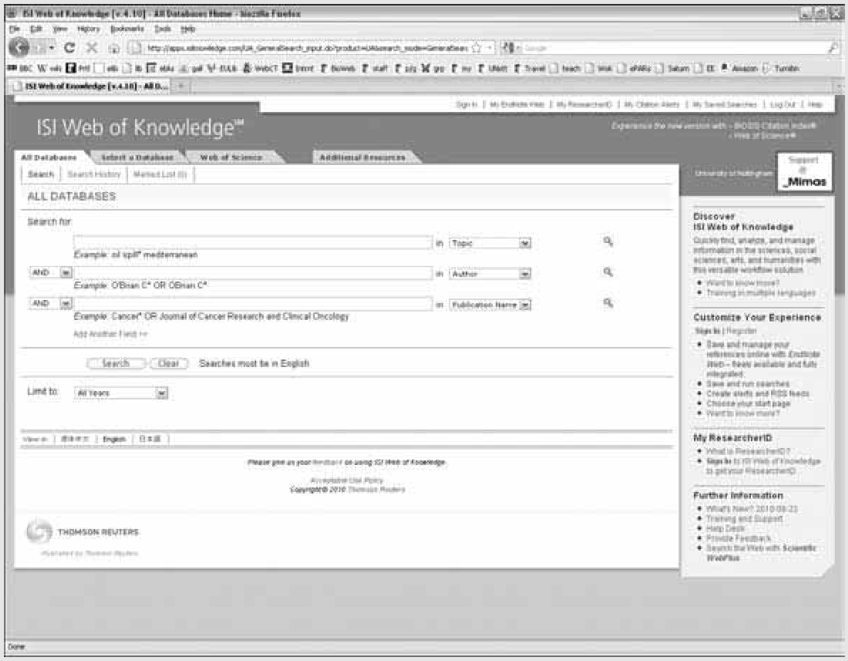

Figure (i) A screen capture from the Web of Knowledge as it appeared in 2010. Like other similar sites, it is regularly updated, so the exact appearance of the search field screen changes.

The key to using the search fields effectively lies in the precision with which you specify your terms: too general and you will be swamped with articles that are of little or no interest; too narrow and you will wind up with only one or two and miss many important ones. To help with this, the search fields provide various means of linking terms so that searches can be focused (the AND, OR, NOT options – called 'operators' – in Fig. (i)). However, the process inevitably involves some compromises.

For example, suppose you were interested in steroid hormone secretion as a cause of immune depression in laboratory mice. You might start, seemingly reasonably, by typing '*steroid hormone* AND *immune depression* AND *laboratory mice*' into the 'Topic' search field in Fig. (i) and hitting the 'Search' button. Disappointingly, and rather to your surprise, this yields nothing at all – apparently nobody has published anything on steroid hormones and immune depression in mice. At the other extreme, a search for '*immune* AND *mice*' yields over 45,000 articles, a wholly unmanageable number, of which many can be seen at a glance to be irrelevant to your needs. Clearly, something between the two is what is required.

The reason the first search turned up nothing is not, of course, because nobody has published anything on the topic, but because the search term was restricted to a very specific combination of phrases. It could well be that people have published on the effects of steroid hormones on immune depression in mice but didn't use the precise phrases selected. For instance, they may have reported 'depressed immune responsiveness' or 'depressed immunity', rather than 'immune depression', and referred to specific hormones, such as testosterone or cortisol, rather than the generic term 'steroid'. There are various ways of catering for this. In the Web of Knowledge, the form '*immun** SAME *depress** AND *mice*' in the 'Topic' search field allows the system to search for any term beginning with 'immun' or 'depress', such as 'immune', 'immunity', 'immunocompetence', 'depression', 'depressed' and so on, thus picking up all the variants. The term 'SAME' ensures similar combinations of phrase are recognised, in this case, say, 'immune depression', 'depressed immunity' or 'depressed immune response'. Running the search again in this form yields around 450 articles, much better than zero or 45,000, but with quite a lot of them still redundant. If the search is specified a little more tightly as '*immun** SAME *depress** AND *mice* AND *hormone*', however, it turns up around 40 articles, and all much more on target.

All the searchable databases use these kinds of approaches for refining searches, some very intuitive, some less so. One thing you will quickly notice, though, is that exactly the same search can turn up a different number and selection of articles depending on which database you are using – BIOSIS, for example, manages to find something under the initial over-specific search that drew a blank on the Web of Knowledge. For this reason, it is good practice to run searches on a selection of databases.

Using general-purpose search engines, like Yahoo! or Google, can often turn up information from the professional literature too, but just as often you're likely to get information from unregulated personal websites, or other sources of uncertain provenance. Taking received wisdom at face value can be a dangerous business – something even seasoned researchers can continue to discover, the famous geneticist and biostatistician R. A. Fisher among them.

In the early 1960s, Fisher and other leading authorities at the time were greatly impressed by an apparent relationship between duodenal ulcer and certain rhesus and MN blood groups. Much intellectual energy was expended trying to explain the relationship. A sceptic, however, mentioned the debate to one of his blood-group technicians. The technician, for years at the sharp end of blood-group analysis, resolved the issue on the spot. The relationship was an artefact of blood transfusion! Patients with ulcers had received transfusions because of haemorrhage. As a result, they had temporarily picked up rhesus and MN antigens from their donors. When patients who had not been given transfusions were tested, the relationship disappeared (Clarke, 1990).

Where at all feasible, therefore, testing assumptions yourself and making up your own mind about the facts available to you is a good idea. Indeed, science is often characterised as systematic scepticism – a demand for evidence for every assertion. It is impossible to draw up a definitive list of what it is an investigator needs to know as essential background; biology is too diverse a subject, and every investigation is to some extent unique in its factual requirements. Nevertheless, it is useful to indicate the kinds of information that are likely to be important. Some examples might be as follows:

Question

Can the material of interest be studied usefully under laboratory conditions or will unavoidable constraints or manipulations so affect it that any conclusions will have only dubious relevance to its normal state or functions?

For instance, can mating preferences in guppies usefully be studied in a small plastic aquarium, or will the inevitable restriction on movement and the impoverished environment compromise normal courtship activity?

Or, if nutrient transfer within a plant can be monitored only with the aid of a vital dye, will normal function be maintained in the dyed state or will the dye interfere subtly with the processes of interest?

Question *Is the material at the appropriate stage of life history or development for the desired investigation?*

There would, for instance, be little point in carrying out vaginal smears on female mice to establish stages of the oestrous cycle if some females were less than 28 days of age. Such mice may well not have begun cycling.

Likewise, it would be fruitless to monitor the faeces of infected mice for the eggs of a nematode worm until a sufficient number of days have passed after infection for the worms to have matured.

Question *Will the act of recording from the material affect its performance?*

For example, removing a spermatophore (package of sperm donated by the male) from a recently mated female cricket in order to assay its sperm content may adversely affect the female's response to males in the future.

Or, the introduction of an intracellular probe might disrupt the aspect of cell physiology it was intended to record.

Question *Has the material been prepared properly?*

If the problem to be investigated involves a foraging task (e.g. learning to find cryptic prey), has the subject been trained to perform in the apparatus and has it been deprived of food for a short while to make it hungry?

Similarly, if a mouse of strain X is to be infected with a particular blood parasite so that the course of infection can be monitored, has the parasite been passaged in the strain long enough to ensure its establishment and survival in the experiment?

Question *Does the investigation make demands on the material that it is not capable of meeting?*

Testing for the effects of acclimation on some measure of coping in a new environment might be compromised if conditions in the new environment are beyond those the organism's physiology or behaviour have evolved to meet.

Likewise, testing a compound from an animal's environment for carcinogenic properties in order to assess risk might not mean much if the compound is administered in concentrations or via routes that the animal could never experience naturally.

Question *Are assumptions about the material justified?*

In an investigation of mating behaviour in dragonflies, we might consider using the length of time a male and female remain coupled as an index of the amount of sperm transferred by the male. Before accepting this, however, it would be wise to conduct some pilot studies to make sure it was valid; it might be, for instance, that some of the time spent coupled reflected mate-guarding rather than insemination.

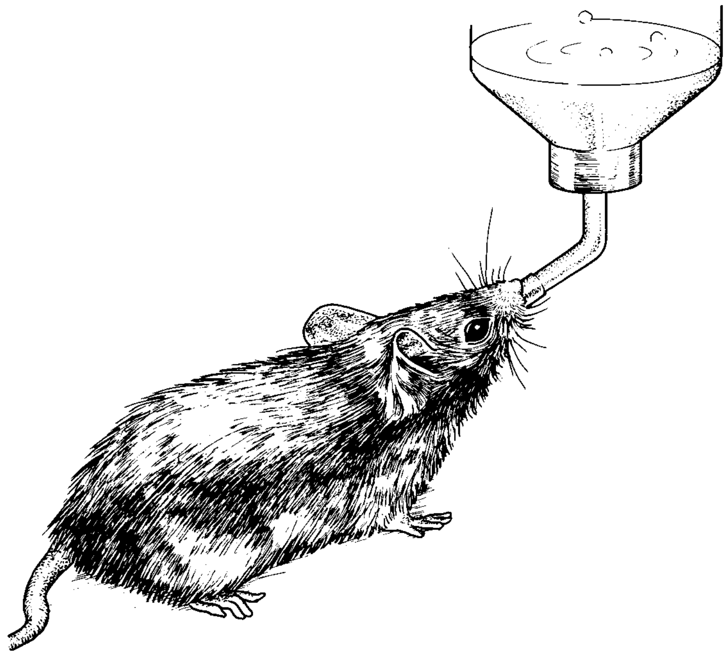

By the same token, assumptions about the relationship between the staining characteristics of cells in histological sections and their physiological properties might need verifying before concluding anything about the distribution of physiological processes within an organ.

The list could go on for a long time, but these examples are basic questions of practicality. They are not very interesting in themselves but they, and others like them, need to be addressed before interesting questions can be asked. Failure to consider them will almost inevitably result in wasted time and materials.

Of course, even at this level, the investigator will usually have the questions ultimately to be addressed – the whole point of the investigation – in mind, and these will naturally influence initial considerations. Before we develop this further, however, there is one further, and increasingly prominent, issue we must address, and that is the *ethics* of working with biological material.

1.2.1 Ethical considerations

Because biological material is either living, or was once living, or is derived from something that is or was living, we are sensitive to the possibility that another living organism may be harmed in some way as a result of what we are doing. Of particular concern is the possibility that our activity might cause such an organism to suffer, physically or psychologically. We try very hard to avoid suffering ourselves because, by definition, it is extremely unpleasant, so the question arises as to whether we should risk inflicting it on another living being simply because we are interested in finding something out about it. This is not an easy question to answer, not least because of the difficulty of knowing whether species very different from ourselves, such as invertebrates, are capable of experiencing anything that might reasonably be called suffering in the first place. However, good science is mindful of the possibility, and works to various guidelines and codes of practice, some enforced by law, to give organisms the benefit of the doubt. While minimising the risk of suffering is important in itself, there is also a straightforward practical reason why we should take care of the organisms we use, whatever they may be, since any results we obtain from them could be affected if the organism is damaged or in some way below par.

Suffering may not be the only potential ethical concern. If material is coming from the field, for example, there could be conservation issues. Is the species concerned endangered? Is the habitat it occupies fragile? Are there unwelcome consequences for populations or habitats of removing material and/or returning it afterwards? Questions like this can lead to acute dilemmas. For instance, the fact that a species is becoming endangered may mean there is a desperate need for more information about it, but the very means of acquiring the information risks further harm.

As awareness of these issues increases, ethical considerations are beginning to play a more explicit role in the way biologists approach their work, not just in terms of taking greater care of the organisms they use, and being better informed about their needs, but at the level of how investigations are designed in the first place. Take sample size, for instance. Deciding on a suitable sample size is a basic problem in any quantitative study. It might involve an informal judgement on the basis of past experience or the outcome of other studies, or it might depend on power tests (*see* Box 3.14) to calculate a sample size statistically. Where there are ethical concerns, a power test would arguably be better than 'guesstimation' because it would provide an objective means of maximising the likelihood of a meaningful result while minimising the amount of material needed (a smaller sample would risk the outcome being swamped by random noise, while a larger one would use more material than necessary). But, of course, the ideal sample size indicated by the power test might demand more material than can be sustained by the source, or involve a very large number of animals in a traumatic experimental procedure. The value of proceeding then has to be judged against the likely cost from an ethical perspective, a task with considerable room for debate. Detailed discussion of these issues is beyond the scope of this book, but a good idea of what is involved can be found in Bateson (1986, 2005), who provides a digestible introduction to trading off scientific value and ethical concerns, and the extensive ethical guidelines for teachers and researchers in animal behaviour published by the Association for the Study of Animal Behaviour (ASAB) and its North American partner, the Animal Behavior Society (ABS) (see **asab.nottingham.ac.uk** or **www.animalbehaviorsociety.org** or each January issue of the academic journal *Animal Behaviour*). It is also well worth looking at the website of the UK National Centre for the 3Rs (**www.nc3rs.org**; the three Rs stand for the Replacement, Refinement and Reduction in the use of animals in research), a government-funded organisation dedicated to progressing ethical approaches to the use of animals in biology. For discussion of more philosophical issues, see, for example, Dawkins (1980, 1993) and Barnard & Hurst (1996). It is important to stress that, tricky as these kinds of decision can be, ethical considerations should *always* be part of the picture when you are working with biological material.

1.3 The skill of asking questions

1.3.1 Testing hypotheses

Charles Darwin once remarked that without a hypothesis a geologist might as well go into a gravel pit and count the stones. He meant, of course, that simply gathering facts for their own sake was likely to be a waste of time. A geologist is unlikely to profit much from knowing the number of stones in a gravel pit. This seems self-evident, but such *undirected* fact-gathering (not to be confused with the often essential descriptive phase of hypothesis development) is a common problem among students in practical and project work. There can't be many science teachers who have not been confronted by a puzzled student with the plea: 'I've collected all these data, now what do I do with them?' The answer, obviously, is that the investigator should know what is to be done with the data before they are collected. As Darwin well knew, what gives data collection direction is a working *hypothesis*. Theories and hypotheses are absolutely vital to science, otherwise 'we shall all be washed out to sea in an immense tide of unrelated information' (Watt, 1971). With them, 'the enormous ballast of factual information, so far from being about to sink us, is used to reveal patterns and processes so that we need no longer to record the fall of every apple' (Dixon, 2000).

The word 'hypothesis' sounds rather formal and, indeed, in some cases hypotheses may be set out in a tightly constructed, formal way. In more general usage, however, its meaning is a good deal looser. Verma & Beard (1981), for example, define it as simply:

> a tentative proposition which is subject to verification through subsequent investigation. . . . In many cases hypotheses are hunches that the researcher has about the existence of relationships between variables.

A hypothesis, then, can be little more than an intuitive feeling about how something works, or how changes in one factor will relate to changes in another, or about any aspect of the material of interest. However vague it may be, though, it is formative in the purpose and design of investigations because these set out to test it. If at the end of the day the results of the investigation are at odds with the hypothesis, the hypothesis may be rejected and a new one put in its place. As we shall see later, *hypotheses are never proven, merely rejected if data from tests so dictate, or retained for the time being for possible rejection after further tests*.

1.3.2 How is a hypothesis tested?

If a hypothesis is correct, certain things will follow. Thus if our hypothesis is that a particular visual display by a male chaffinch is sexual in motivation, we might expect the male to be more likely to perform the display when a female is present. Hypotheses thus generate *predictions*, the testing of which increases or decreases our faith in them. If our male chaffinch turned out to display mainly

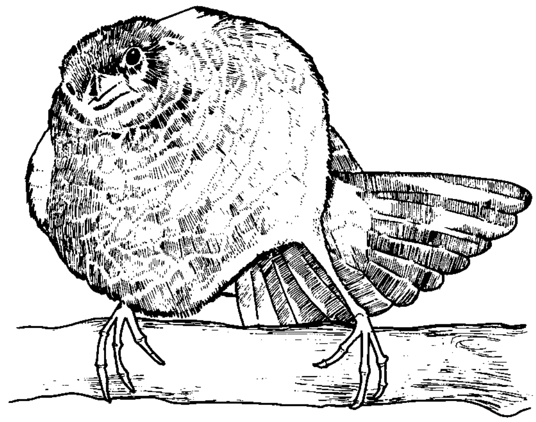

when other males were around and almost never with females, we might want to think again about our sexual motivation hypothesis. However, we should be wrong to dismiss it solely on these grounds. It could be that such displays are important in defending a good-quality breeding territory that eventually will attract a female. The context of the display could thus still be sexual, but in a less direct sense than we had first considered. In this way, hypotheses can produce tiers of more and more refined predictions before they are rejected or tentatively accepted. Making such predictions is a skilled business because each must be phrased so that testing it allows the investigator to discriminate in favour of or against the hypothesis. While it is best to phrase predictions as just that (thus: *males will perform more of display* y *in the presence of females*), they sometimes take the form of questions (*do males perform more of display* y *when females are present?*). The danger with the question format, however, is that it can easily become too woolly and vague to provide a rigorous test of the hypothesis (e.g. *do males behave differently when females are present?*). Having to phrase a precise prediction helps counteract the temptation to drift into vagueness.

Hypotheses, too, can be so broad or imprecise that they are difficult to reject. In general the more specific, mutually exclusive hypotheses that can be formulated to account for an observation the better. In our chaffinch example, the first hypothesis was that the display was sexual. Another might be that it reflected aggressive defence of food. Yet another that it was an anti-predator display. These three hypotheses give rise to very different predictions about the behaviour and it is thus, in principle, easy to distinguish between them. As we have already seen, however, distinguishing between the 'sexual' and 'aggressive' hypotheses may need more careful consideration than we first expect. *Straw man hypotheses* are another common problem. Unless some effort has gone into understanding the material, there is a risk of setting up hypotheses that are completely inappropriate. Thus, suggesting that our displaying chaffinch was demonstrating its freedom from avian malaria would make little sense in an area where malaria was not endemic. We shall look at the development of hypotheses and their predictions in more detail later on.

1.4 Where do questions come from?

As we have already intimated, questions do not spring out of a vacuum. They are triggered by something. They may arise from a number of sources.

1.4.1 Curiosity

Questions arise naturally when thinking about any kind of problem. Simple curiosity about how something works or why one group of organisms differs in some way from another group can give rise to useful questions from which testable hypotheses and their predictions can be derived. There is nothing wrong with 'armchair theorising' and 'thought experiments' as long as, where possible, they are put to the test. Sitting in the bath and wondering about how migratory birds manage to navigate, for example, could suggest roles for various environmental cues like the sun, stars and topographical features. This in turn could lead to hypotheses about how they are used and predictions about the effects of removing or altering them. By the time the water was cold, some useful preliminary experiments might even have been devised.

1.4.2 Casual observation

Instead of dreaming in the bath, you might be watching a tank full of fish, or sifting through some histological preparations under a microscope. Various things might strike you. Some fish in the tank might seem very aggressive, especially towards others of their own species, but this aggressiveness might occur only around certain objects in the tank, perhaps an overturned flowerpot or a clump of weed. Similarly, certain cells in the histological preparations may show unexpected differences in staining or structure. Even though these aspects of fish behaviour and cell appearance were not the original reason for watching the fish or looking at the slides, they might suggest interesting hypotheses for testing later. A plausible hypothesis to account for the behaviour of the fish, for instance, is that the localised aggression reflects territorial defence. Two predictions might then be: (a) *on average, territory defenders will be bigger than intruders* (because bigger fish are more likely to win in disputes and thus obtain a territory in the first place) and (b) *removing defendable resources like upturned flowerpots will lead to a reduction in aggressive interactions*. Similarly, a hypothesis for differences in cell staining and structure is that they are due to differences in the age and development of the cells in question. A prediction might then be: *younger tissue will contain a greater proportion of* (what are conjectured to be the) *immature cell types*.

1.4.3 Exploratory observations

It may be that you already have a hypothesis in mind, say that a particular species of fish will be territorial when placed in an appropriate aquarium environment.

What is needed is to decide what an appropriate aquarium environment might be so that suitable predictions can be made to test the hypothesis. Obvious things to do would be to play around with the size and number of shelters, the position and quality of feeding sites, the number and sex ratio of fish introduced into the tank, and so on. While the effects of these and other factors on territorial aggressiveness among the fish might not have been guessed at beforehand, such manipulations are likely to suggest relationships with aggressiveness that can then be used to predict the outcome of further, *independent* investigations. Thus if exploratory results suggested aggressiveness among defending fish was greater when there were ten fish in the tank compared with when there were five, it would be reasonable to predict that aggressiveness would increase as the number of fish increased, *all other things being equal*. An experiment could then be designed in which shelters and feeding sites were kept constant but different numbers of fish, say 2, 4, 6, 8, 10 or 15, were placed in the tank. Measuring the amount of aggression by a defender with each number of fish would provide a test of the prediction.

1.4.4 Previous studies

One of the richest sources of questions is, of course, past and ongoing research. This might be encountered either as published literature (*see* Box 1.1) or 'live' as research talks at conferences or seminars. A careful reading of most published papers, articles or books will turn up ideas for further work, whether at the level of alternative hypotheses to explain the problem in hand or at the level of further or more discriminating predictions to test the current hypothesis. Indeed, this is the way most of the scientific literature develops. Some papers, often in the form of mathematical models or speculative reviews, are specifically intended to generate hypotheses and predictions and may make no attempt to test them themselves. At times, certain research areas can become overburdened with hypotheses and predictions, generating more than people are able or have the inclination to test. If this happens, it can have a paralysing effect on the development of research. It is thus important that hypotheses, predictions and tests proceed as nearly as is feasible hand in hand.

1.5 What this book is about

We've said a little about how science works and how the kind of question-asking on which it is based can arise. We now need to look at each part of the process in detail, because while each may seem straightforward in principle, some knotty problems can arise when science is put into practice. In what follows, we shall see how to:

- frame hypotheses and predictions from preliminary source material,
- design experiments and observations to test predictions,

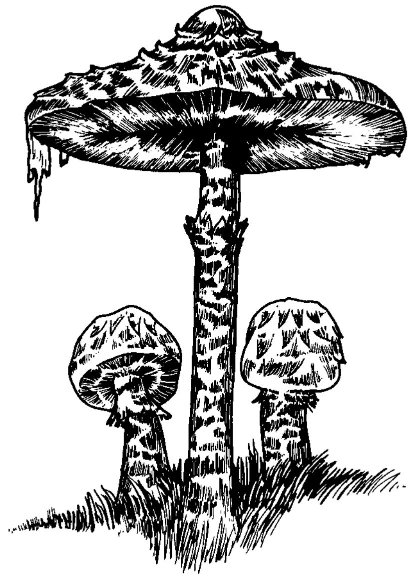

- analyse the results of tests to see whether they are consistent with our original hypothesis, and
- present the results and conclusions of tests so that they are clear and informative.

The discussion deals with these aspects in order so that the book can be read straight through or dipped into for particular points. A summary at the end of each chapter highlights the important take-home messages, and the self-test questions at the end show what you should be able to tackle after reading the book.

Remember, the book is about asking and answering questions in biology – it is not a biology textbook or a statistics manual, and none of the points it makes are restricted to the examples that illustrate them. At every stage you should be asking yourself how what it says might apply in other biological contexts, especially if you have an interest in investigating them!

We suggest you get hold of a statistical package in order to follow the analyses. The two we use in detail in the book are AQB and R. AQB is a set of Excel sheets that can be downloaded free from **www.pearsoned.co.uk/barnard**. You can download the installation file of R free from **www.r-project.org**, and follow the simple instructions to install it on your computer. This book will enable you to run analyses in R; a simple and very clear guide to learning more about R is available in Zuur *et al.* (2009).

References

Barnard, C. J. & Hurst, J. L. (1996) Welfare by design: the natural selection of welfare criteria. *Animal Welfare* **5**, 405–433.

Bateson, P. (1986) When to experiment on animals. *New Scientist* **109**, 30–32.

Bateson, P. (2005) Ethics and behavioral biology. *Advances in the Study of Behavior* **35**, 211–233.

Clarke, C. (1990) Professor Sir Ronald Fisher FRS. *British Medical Journal* **301**, 1446–1448.

Dawkins, M. (1980) *Animal suffering: The science of animal welfare*. Chapman and Hall, London.

Dawkins, M. (1993) *Through our eyes only?: The search for animal consciousness*. Oxford University Press, Oxford.

Dixon, A. F. G. (2000) *Insect predator–prey dynamics: Ladybird beetles and biological control*. Cambridge University Press, Cambridge.

Watt, K. E. F. (1971) Dynamics of populations: a synthesis. In P. J. den Boer & G. R. Gradwell (eds) *Dynamics of populations*. PUDOC, Wageningen, pp. 568–580.

Verma, G. K. and Beard, R. M. (1981) *What is educational research? Perspectives on techniques of research*. Gower, Aldershot.

Zuur, A. F., Ieno, E. N. & Meesters, E. H. W. G. (2009) *A beginner's guide to R*. Springer, London.