


# ASKING QUESTIONS IN BIOLOGY

3rd Edition

A guide to Hypothesis-testing,  
Experimental Design and  
Presentation in Practical  
Work and Research Projects



Chris Barnard  
Francis Gilbert  
Peter McGregor



# Asking Questions in Biology



We work with leading authors to develop the strongest educational materials in biology, bringing cutting-edge thinking and best learning practice to a global market.

Under a range of well-known imprints, including Prentice Hall, we craft high-quality print and electronic publications that help readers to understand and apply their content, whether studying or at work.

To find out more about the complete range of our publishing please visit us on the World Wide Web at:  
[www.pearsoned.co.uk](http://www.pearsoned.co.uk)



Third Edition

# Asking Questions in Biology

A Guide to Hypothesis-testing, Analysis and  
Presentation in Practical Work and Research

Chris Barnard, Francis Gilbert

School of Biology, University of Nottingham

and Peter McGregor

Duchy College, Cornwall

PEARSON  
Prentice  
Hall

Harlow, England • London • New York • Boston • San Francisco • Toronto  
Sydney • Tokyo • Singapore • Hong Kong • Seoul • Taipei • New Delhi  
Cape Town • Madrid • Mexico City • Amsterdam • Munich • Paris • Milan

**Pearson Education Limited**

Edinburgh Gate  
Harlow  
Essex CM20 2JE  
England

and Associated Companies throughout the world

*Visit us on the World Wide Web at:*  
[www.pearsoned.co.uk](http://www.pearsoned.co.uk)

---

First published under the Longman imprint 1993  
Second edition 2001  
**Third edition 2007**

© Pearson Education Limited 1993, 2001, 2007

The rights of Chris Barnard, Francis Gilbert and Peter McGregor to be identified as authors of this work have been asserted by them in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without either the prior written permission of the publisher or a licence permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd., Saffron House, 6–10 Kirby Street, London EC1N 8TS.


ISBN 978 0 13 222435 2

**British Library Cataloguing-in-Publication Data**

A catalogue record for this book is available from the British Library

10 9 8 7 6 5 4 3 2 1  
10 09 08 07


Typeset in 9.5/12.5pt Concorde BE by 35  
Printed by Ashford Colour Press Ltd., Gosport



# Contents

<b>Preface</b>	vii
<b>1 Doing science</b>	
<b>Where do questions come from?</b>	<b>1</b>
1.1 Science as asking questions	2
1.2 Basic considerations	3
1.3 The skill of asking questions	9
1.4 Where do questions come from?	10
1.5 What this book is about	12
References	13
<b>2 Asking questions</b>	
<b>The art of framing hypotheses and predictions</b>	<b>14</b>
2.1 Observation	14
2.2 Exploratory analysis	19
2.3 Forming hypotheses	39
2.4 Summary	48
Reference	48
<b>3 Answering questions</b>	
<b>What do the results say?</b>	<b>49</b>
3.1 Confirmatory analysis	49
3.2 What is statistical significance?	51
3.3 Significance tests	54
3.4 Testing hypotheses	136
3.5 Testing predictions	142
3.6 Refining hypotheses and predictions	148
3.7 Summary	149
References	150

<b>4 Presenting information</b>	
<b>How to communicate outcomes and conclusions</b>	<b>152</b>
4.1 Presenting figures and tables	152
4.2 Presenting results in the text	165
4.3 Writing reports	166
4.4 Writing for a more general readership	180
4.5 Presenting in person: spoken papers and poster presentations	182
4.6 Plagiarism	195
4.7 Summary	196
Reference	197
<b>Test finder and help guide</b>	<b>198</b>
<b>Some self-test questions</b>	<b>201</b>
<b>Appendix I: Table of confidence limits to the median</b>	<b>207</b>
<b>Appendix II: How to calculate some simple significance tests</b>	<b>208</b>
<b>Appendix III: Significance tables</b>	<b>223</b>
<b>Answers to self-test questions</b>	<b>229</b>
<b>Index</b>	<b>236</b>
<b>Quick test finders</b>	<b>244</b>



## Preface

Science is a process of asking questions, in most cases precise, quantitative questions that allow distinctions to be drawn between alternative explanations of events. Asking the right questions in the right way is a fundamental skill in scientific enquiry, yet in itself it receives surprisingly little explicit attention in scientific training. Students being trained in scientific subjects, for instance in sixth forms, colleges and universities, learn the factual science and some of the tools of enquiry such as laboratory techniques, mathematics, statistics and computing, but they are taught little about the process of question-asking itself.

The first edition of this book had its origins in a first-year undergraduate practical course that we, and others since, have run at the University of Nottingham for several years. The approach adopted there now also forms the basis for a more advanced second-year course. The aim of the courses is to introduce students in the biological sciences to the skills of observation and enquiry, but focusing on the process of enquiry – how to formulate hypotheses and predictions from raw information, how to design critical observations and experiments, and how to choose appropriate analyses – rather than on laboratory, field and analytical techniques *per se*. This focus is maintained in the third edition. However, as in the second edition, we have responded to a number of positive suggestions from people who have used the book, either as teachers or as students, which we think enhance further its usefulness in teaching practical biology generally.

The most major change has been with respect to the presentation of statistical tests. In the previous editions, we presented these as boxes of formulae for hand calculation, supported by worked examples of each test in the Appendices. Partly because of this, we also limited the tests to fairly basic ones. Feedback from users and reviewers, however, as well as our own experience, strongly indicates a preference for a computer-based presentation; that is, presenting tests as they are encountered in computer statistical packages. We have thus entirely replaced the hand-calculation boxes in the main text with boxes based on procedures and output in two widely used commercial packages: Microsoft Excel and Statistical Packages for the Social Sciences (SPSS), along with a package (AQB) written specifically to accompany the book that allows additional important tests not covered by the two (or any other, as far as we are aware) commercial packages. However, we have included the hand-calculation boxes in the Appendices because, as previously, we consider it important that the underlying arithmetic of the (basic at least) tests is understood. As well as adopting a computer package approach, we have greatly expanded the range of tests covered to



include, among other things, repeated-measures designs, analysis of covariance, multiple regression and principal components analysis. This therefore caters for the wider spectrum of eventualities likely to be encountered in extended project and research contexts. The Test Finder, Quick Test Finder and Help sections have been expanded accordingly.

With increasing emphasis on the wider communication and public understanding of science, we have introduced entirely new sections on presenting information as a talk or poster paper at meetings, and writing for a broader, non-specialist readership, such as that of a newspaper or magazine. We have also added sections on using on-line literature databases, such as the Web of Science and Biosis, and the ethical implications of working with biological material, the latter a now essential consideration in any study as the legal regulation of procedures, including those used in teaching, gathers pace.

The book looks at the process of enquiry during its various stages, starting with unstructured observations and working through to the production of a complete written report. In each section, different skills are emphasized and a series of main examples runs through the book to illustrate their application at each stage.

The book begins with a look at scientific question-asking in general. How do we arrive at the right questions to ask? What do we have to know before we can ask sensible questions? How should questions be formulated to be answered most usefully? Chapter 1 addresses these points by looking at the development of testable hypotheses and predictions and the sources from which they might arise.

Chapter 2 looks at how hypotheses and predictions can be derived from unstructured observational notes. Exploratory analysis is an important first step in deriving hypotheses from raw data and the chapter introduces plots and summary statistics as useful ways of identifying interesting patterns on which to base hypotheses. The chapter concludes by pointing out that although hypotheses and their predictions are naturally specific to the investigation in hand, testable predictions in general fall into two distinct groups: those dealing with some kind of *difference* between groups of data and those dealing with a *trend* in the quantitative relationship between groups of data.

The distinction between difference and trend predictions is developed further in Chapter 3, which discusses the use of confirmatory analyses. The concept of statistical significance is introduced as an arbitrary but generally agreed yardstick as to whether observed differences or trends are interesting, and a number of broadly applicable significance tests are explained. Throughout, however, the emphasis is on the use of such tests as tools of enquiry rather than on the statistical theory underlying them. Having introduced significance tests and some potential pitfalls in their use, the book uses the main worked examples to show how some of their predictions can be tested and hypotheses refined in the light of testing.

In Chapter 4, the book considers the presentation of information. Once hypotheses have been tested, how should the outcome be conveyed for greatest effect? The chapter discusses the use of tables, figures and other modes of presentation and shows how a written report should be structured. The chapter then moves on to consider the presentation of material in spoken and poster

paper formats, and how to recast written reports of results for a general, rather than specialist, readership.

At the end of the book is a number of appendices. These provide expanded self-test questions and answers sections based on the material in the previous chapters, formulae for calculating the simpler significance tests and some statistical tables for use in significance testing.

We said that the book had its inception in our introductory practical course. This course was developed in response to an increasingly voiced need on the part of students to be taught how to formulate hypotheses and predictions clearly and thus design properly discriminating experiments and observations. As we descend into ever more neurotically prescribed teaching and assessment procedures, the need for students to be given clear guidance on such aspects of their work becomes correspondingly greater. As always, both our practical teaching and the book have continued to benefit enormously from our ongoing and enjoyable interaction with our undergraduates. Their insights and enquiries continue to hone the way we teach and have been the guiding force behind all the discussions in the book.

Finally, we should like to thank all the people who have commented on the book since its first appearance, and encouraged us to think about the further amendments we have made in this present edition. In particular, we thank Tom Reader for the generous amount of time he has spent discussing the book with us, commenting on drafts of many of the amendments and giving freely of his experience in and enthusiasm for the business of communicating science. We also thank six anonymous reviewers of the second edition for their helpful suggestions as to what a third edition might comprise; we have followed these almost to the letter and the final product has undoubtedly been enhanced by them.

*Chris J. Barnard*  
*Francis S. Gilbert*  
*Peter K. McGregor*  
November 2006

## Acknowledgements


We are grateful to Blackwell Science Ltd. for permission to reproduce copyright material: Figs 2.6 and 3.3 modified from *Choosing and using statistics*, published by Blackwell Science Ltd (Dytham, C. 1999).

## Companion Web Site

Material is available on the Internet to support students and lecturers using this book, at the web address [www.pearsoned.co.uk/barnard](http://www.pearsoned.co.uk/barnard). This comprises downloadable Microsoft Excel® sheets that enable the student or lecturer to carry out in Excel all of the statistical tests described and included within the book. The tests covered include  $\chi^2$ , correlation and regression, and the specific and general hypotheses tested by analyses of variance, both parametric and non-parametric. The package forms a useful supporting resource to the text.

DEDICATED TO THE MEMORY  
OF CHRIS BARNARD

Chris Barnard succeeded in completing this edition before his untimely death in June 2007 at the age of 55. We dedicate this to his memory: his talents as a scientist, educator and communicator of science, raconteur, poet, artist and musician will be greatly missed.



# 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 Science (<http://wos.mimas.ac.uk>), PubMed ([www.pubmedcentral.nih.gov](http://www.pubmedcentral.nih.gov)) or BIOSIS ([www.biosis.org](http://www.biosis.org)). 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. Using general-purpose search engines, like Yahoo! or Google, can often turn up

## Box 1.1 Searching online literature databases

Searchable online literature databases, like the Web of Science, 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 Science, is shown in Fig. (i).

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



**Figure (i)** A screen capture from the Web of Science as it appeared in 2006. Like other similar sites, it is regularly updated, so the exact appearance of the search field screen changes.

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 Science, 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 Science. For this reason, it is good practice to run searches on a selection of databases.

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. 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* [www.asab.org](http://www.asab.org) or [www.animalbehavior.org](http://www.animalbehavior.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](http://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 and 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*.


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 and 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 defensible resources like overturned 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 two, four, six, eight, ten 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:

1. frame hypotheses and predictions from preliminary source material,
2. design experiments and observations to test predictions,
3. analyse the results of tests to see whether they are consistent with our original hypothesis, and
4. 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!

## References

- Barnard, C. J. and 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.
- Dawkins, M. (1993) *Through our eyes only?: The search for animal consciousness*. Oxford University Press, Oxford.
- Verma, G. K. and Beard, R. M. (1981) *What is educational research? Perspectives on techniques of research*. Gower, Aldershot.





# 2

## Asking questions

### The art of framing hypotheses and predictions

So far, we've discussed asking questions in a very general way. Simply being told that science proceeds like this, however, is not particularly helpful unless it is clear how the principles can be applied to the situation in hand. The idea of this chapter is thus to look at the development of the procedure in the context of various investigations that you might undertake in practical and project work. We shall assume for the moment that the material of interest is derived from your own observations. We shall start, therefore, with the problem of making observations and directing them in order to produce useful information.

## 2.1 Observation

### 2.1.1 Observational notes and measurements

When first confronted with an unfamiliar system, it is often difficult to discern anything of interest straight away. This seems to be true regardless of the complexity of the system. For instance, a common cause of early despair among students watching animals in a tank or arena for the first time is the *mêlée* of ceaselessly changing activities, many of which seem directionless and without an obvious goal. An equally common complaint is that the animals seem to be doing nothing at all worth mentioning. It is not unusual for *both* extremes to be generated by the same animals.

In both of the earlier mentioned cases, the problem almost always turns out to be not what the animals are or are not doing, but the ability of the observer to observe. This is because observation involves more than just staring passively at material on the assumption that if anything about it is interesting then it will


also be obvious. To be revealing, observations may need to be very systematic, perhaps involving manipulations of the material to see what happens. They may involve measurements of some kind since some things may be apparent quantitatively rather than qualitatively. In themselves, therefore, observations are likely to involve a certain amount of question-asking. Their ultimate purpose, however, is to provide the wherewithal to frame testable hypotheses and the discriminating predictions that will distinguish between the hypotheses.

To see how the process works, we shall first give some examples of observational notes and then look at the way these can be used to derive hypotheses and predictions. These examples, therefore, develop through the book from initial observational notes, through framing and testing hypotheses, to producing a finished written report. The examples are based on the kinds of preliminary notes made during practical, field course and research project exercises by our own students at various stages of their undergraduate training. They come from four different fields of study, but their common feature is that they provide scope for open-ended investigation and hypothesis testing. Of course, the fact that we happen to have selected these particular examples to illustrate the process is irrelevant to the aim of the book, as the range of other examples running through it amply demonstrates. What emerges from the examples applies with equal weight in all branches of biology from molecular genetics and cell biology to psychology and comparative anatomy.

### Example 1

**Material:** Samples of leaves collected in the field from early-successional (dandelion, *Taraxacum officinale*; plantain, *Plantago lanceolata*; poppy, *Papaver rhoeas*), mid-successional (clover, *Trifolium repens*; ox-eye daisy, *Leucanthemum vulgare*) and late-successional (dogrose, *Rosa canina*; ragwort, *Senecio jacobaea*; blackthorn, *Prunus spinosa*; goldenrod, *Solidago canadensis*) plant species; graph paper; binocular microscope.

**Notes:** Collected leaf samples show a lot of variation in damage. Some leaves have several semi-circles eaten in from the edges; some have numerous holes through the tissue, many brown round the edges; others show extensive damage with most of the leaf missing in many cases and damage to the stems and twigs they are on. There is considerable size variation in the leaves. Look at the size range in undamaged leaves and divide into three size classes. Measure size by drawing round whole leaves on graph paper and counting the squares within the boundary. Count number of leaves with different kinds of damage in each size class: 'small' leaves – 41 with small hole damage, four with marginal damage, 17 with severe damage; 'medium-sized' leaves – 12 with small holes, 29 with marginal damage, 34 with severe damage; 'large' leaves – two with small holes, 22 with marginal damage, 40 with severe damage. One obvious possibility is that size reflects height off the ground and the effects of different kinds of herbivore, perhaps mainly slugs and snails on the smaller leaves, caterpillars and other insects on the medium-sized (shrubs, bushes?) leaves and maybe cattle or deer on the largest. It also looks as if the smaller (low-growing?) leaves have much more damage than the larger ones (84 per cent showing some damage,



compared with 47 and 23 per cent in the other two categories). Maybe this is just because bigger plants have more leaves, so more escape damage. While examining the samples, notice several other things. Some leaves have tough 'skins', often with hairs on the surface; these are mainly from the medium-sized category (per cent with hairs visible to naked eye: small, 0; medium, 21; large, 9). Some of the bigger leaves smell strongly or have sticky or latexy sap when squeezed, but the last is also the case with what look like dandelion leaves. The stems of some of the medium-sized and bigger ones have thorns or sticky hairs. There's also more colour variation in these two categories, some leaves being reddish rather than green. It looks like the tougher, more strongly smelling samples generally have less damage than the others.

### Example 2

**Material:** Stained blood smears, vials of preserved ectoparasites, gut nematodes and faecal samples from live-trapped bank voles (*Clethrionomys glareolus*), microscope with eyepiece graticule, clean microscope slides and coverslips, pipettes.


**Notes:** Looking at blood smears under a microscope, notice range of red and white blood cells. Some red cells in some of the smears have small stained bodies in them. These turn out to be a stage in the life cycle of a protozoan parasite, *Babesia microti*, which infects voles. Some slides seem to have much higher densities of red cells than others. The number of fleas and ticks in each vial varies a lot. Several voles didn't seem to have any, while some had a large number. Divide samples by age and sex, and do some counts of infected red cells; scan a roughly standard-width field along the graticule scale for each vole and count number of infected cells (adult males: 21, 45, 3, 0, 64; adult females: 16, 1, 13, 0, 0; juvenile males: 0, 5, 34, 0, 0; juvenile females: 0, 0, 0, 16, 0). Count ticks recovered from same groups (adult males: 0, 8, 3, 0, 7; adult females: 0, 1, 0, 0, 2; juvenile males: 2, 2, 0, 5, 3; juvenile females: 0, 0, 0, 2, 0). Smear some faecal samples onto slides and inspect under microscope. Lots of fragments of plant material and detritus. Some samples have clear oval objects which, on asking, turn out to be nematode eggs. Sometimes there are very large numbers of eggs, sometimes none. Too difficult to count all the eggs in each sample, so designate a rank score from 0 (no eggs) to 5 (more than 100 eggs). Scores for adult males: 3, 5, 5, 0, 3; for adult females: 2, 5, 0, 2, 2; for juvenile males: 2, 2, 4, 2, 1; for juvenile females: 0, 1, 0, 2, 1. On looking at the tubes of preserved worms from the same animals, notice that those with high egg scores do not always have more worms, but some have a greater range of worm sizes. Measure small samples of worms against scale on graticule (ranges for five males: 12–25 units, 9–18 units, 13–28 units, 17–22 units, 11–14 units; ranges for five females: 14–23, 13–19, 12–20, 13–30, 15–29 units).



### Example 3

**Material:** Vials of water containing suspensions of soil-dwelling nematodes from three sites differing in heavy metal and organophosphate pollution, microscope with eyepiece graticule, clean microscope slides and coverslips, pipettes, diagnostic key to common species morphologies.

**Notes:** Pipette  $2 \times 0.2$ -ml droplets of each sample onto a clean slide and examine under microscope. Count adult worms of identifiable species on each slide. Sample 1 (heavy metal polluted site) six apparent species, call A–F for the moment: numbers – A 27, B 5, C 17, D 3, E 32, F 2; Sample 2 (unpolluted) – B 43, C 4, D 15, F 18, plus four different species G 20, H 31, I 4, J 12; Sample 3 (organophosphate polluted site) – A 5, C 48, E 19, H 11. Juvenile worms also present but not readily identifiable to adult species. Nevertheless, numbers in each sample are 23 (Sample 1), 31 (Sample 2), 0 (Sample 3). Can also detect adult females with eggs. Number of females with eggs of each species in samples: Sample 1 – A 6, B 0, C 4, D 0, E 5, F 0; Sample 2 – B 13, C 0, D 2, F 6, G 5, H 11, I 0, J 1; Sample 3 – A 0, C 12, E 7, H 5. Samples contain quite a lot of detritus, some of which can be identified as decomposing nematode material but not related to species. Take ten arbitrarily chosen fields per slide and see how many contain at least one piece of decomposing material: Sample 1 – 5, Sample 2 – 3, Sample 3 – 7. Repeat for a further two droplets per sample. Number of adults in second set: Sample 1 – A 15, C 21, D 9, F 4; Sample 2 – B 31, C 6, D 21, F 11, G 16, H 24, J 3; Sample 3 – A 8, C 35, D 3, E 14, H 6. Juveniles per sample – 17




(Sample 1), 19 (Sample 2), 5 (Sample 3). Females with eggs: Sample 1 – A 3, C 4, D 0, F 0; Sample 2 – B 10, C 0, D 8, F 4, G 4, H 7, J 0; Sample 3 – A 0, C 11, D 0, E 5, H 0. Number of fields per slide with decomposing material: Sample 1 – 3, Sample 2 – 3, Sample 3 – 9.

#### Example 4

**Material:** Stock cage of virgin female and stock cage of virgin male field crickets (*Gryllus bimaculatus*), two or three 30 × 30-cm glass/Perspex arenas with silver sand substrate, dish of water-soaked cotton wool and rodent pellets, empty egg boxes, assorted colours of enamel paint, fine paintbrush, paint thinners, rule, bench lamps.

**Notes:** Females are distinguished from males by possession of long, thin ovipositor at the back. Put four males into an arena. After rushing about, males move more slowly around the arena. When they meet, various kinds of interaction occur. Interactions involve variety of behaviours: loud chirping, tapping each other with antennae, wrestling and biting. Interactions tend to start with chirping and antenna tapping, and only later progress to fighting. Count number of encounters that result in fighting (15 out of 21). Put in three more males so seven in total and count fights again (8 out of 25). Take out males and choose another five. Take various measurements from each male (length from jaws to tip of abdomen, width of thorax, weight) and mark each one with a small, different-coloured dot of paint. Introduce individually marked males into arena. Count number of fights initiated by each male per encounter (red, thorax width 6.5 mm – 4/10; blue, width 7.5 mm – 8/11; yellow, width 7.0 mm – 8/10; silver, 6.0 mm – 3/12; green, 6.5 mm – 3/9). Continue observations and count number of encounters won by each male (win decided if opponent backs off) (red, 3 wins/6 encounters; blue, 4/5; yellow, 7/7; silver, 1/4; green, 4/8).



Introduce three sections of egg box to arena and leave males for 10 minutes. After 10 minutes, some males (blue, yellow, red) hiding under egg box shelters or sitting close to them. Males in or near boxes chirping frequently. Count number of encounters resulting in fight (12 out of 16). Count fights/encounter at different distances from burrows for each of the four males in turn – within 5 cm: 4 attempts in 6 encounters for yellow, 2/2 for red, 0/3 for silver, 4/5 for blue; between 5 and 10 cm: 5/11 (yellow), 2/3 (red), 1/5 (silver), 2/2 (blue);

10–15 cm: 3/7 (yellow), 4/8 (red), 0/2 (silver), 3/6 (blue); 15–20 cm: 2/9 (yellow), 1/6 (red), 1/5 (silver), 2/4 (blue); 20–25 cm: 1/8 (yellow), 1/12 (red), 0/2 (silver), 2/8 (blue). Introduce two females. Females move around the arena but end up mainly around the box sections. Show interest in males and occasionally mount. Other males sometimes interfere when female mounting particular male. Number of approaches to, and number of mounts with, different males: red, two approaches, no mounts; blue, two approaches, two mounts; yellow, three approaches, two mounts; silver, no approaches; green, one approach, one mount.

## 2.2 Exploratory analysis

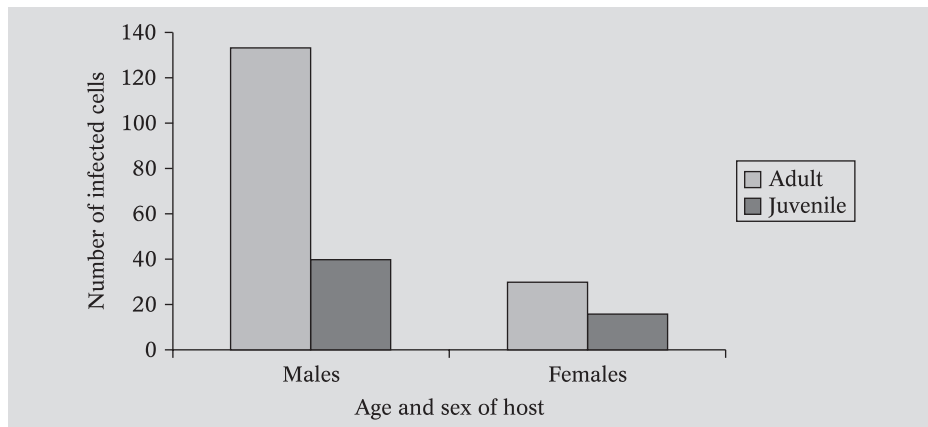
Observational notes are, in most cases, an essential first step in attempting to investigate material for the first time. However, as the examples amply demonstrate, they are a tedious read and, as they stand, do not make it easy to formulate hypotheses and design more informative investigations. What we need is some way of distilling the useful information so that points of interest become more apparent. If we have some numbers to play with – and this underlines the usefulness of making a few measurements at the outset – we can perform some exploratory analyses.

Exploratory analysis may involve drawing some simple diagrams or plotting a few numbers on a scattergram or it might involve calculating some *summary* or *descriptive statistics*. We shall look at both approaches shortly using information from the various sets of observational notes. These sorts of analyses almost always repay the small effort demanded but, like much basic good practice in any field, they are often the first casualty of impatience or prejudice of what is interesting or to be expected. It is always difficult to discern pattern simply by ‘eyeballing’ raw numbers and the more numbers we have the more difficult it becomes. A simple visual representation like a scattergram or a bar chart, however, can turn the obscure into the obvious.

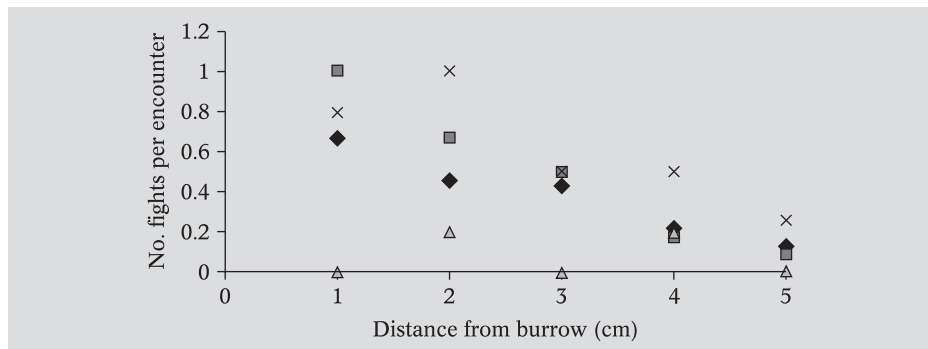
### 2.2.1 Visual exploratory analysis

There are several instances in the examples of observational notes where similar measurements were made from different kinds of material or under different conditions. For instance, infected blood cell counts were taken from adult and juvenile voles of both sexes, while fighting in male crickets was observed at different distances from artificial burrows. In both cases there seem to be some differences in the numbers recorded from different kinds of material (age and sex of host) or under different conditions (distance from a burrow). What do the differences suggest?

Eyeballing the blood cell data suggests some differences both between males and females and between adults and juveniles. A simple way to visualise this might be to total up the number of infected cells scored for each category of animal and present them in a bar chart (Fig. 2.1). If we do this, it looks as though



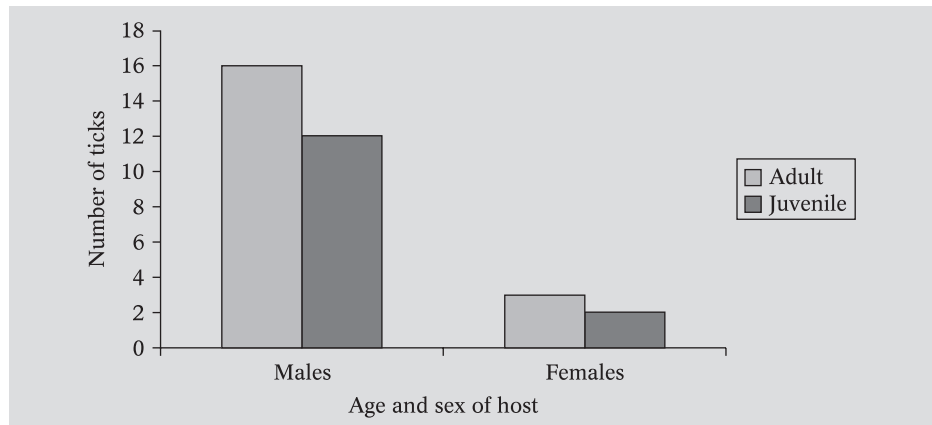
**Figure 2.1** The total number of infected red blood cells in voles of different age and sex (see Example 2, Notes).



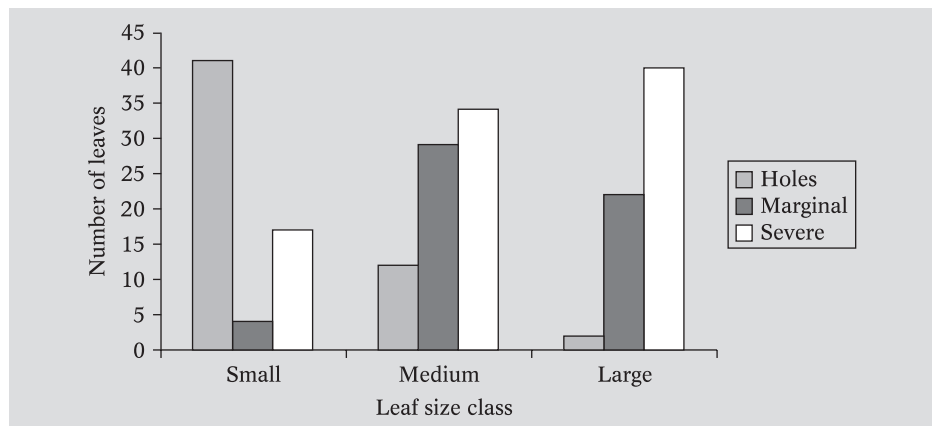
**Figure 2.2** The number of fights per encounter between crickets at different distances from the nearest burrow; symbols represent individuals (see Example 4, Notes).

males generally have higher infection levels than females, regardless of age, and that adults have more infection than juveniles, regardless of sex. From this, we might be tempted to suggest that adult males are particularly prone to infection compared with other classes of individual. As we shall see later, however, we might want to be cautious in our speculation.

In the cricket example, we can see that the number of fights per encounter seems to vary with distance away from an artificial burrow. An obvious way to see how is to plot a scattergram of number of fights per encounter against distance. The result (Fig. 2.2) suggests that, while there is a fair spread of values at each distance, there is a tendency for more encounters to result in a fight when crickets are close to a burrow. It is important to bear in mind that, as each cricket was observed in turn, results for the different animals are independent of each other and the trend is not a trivial outcome of a fight scored for one cricket also counting as a fight for his opponent. As we shall see later, this and other kinds of non-independence can be troublesome in drawing inferences from data.



**Figure 2.3** The number of ticks recovered from voles of different age and sex (see Example 2, Notes).



**Figure 2.4** The number of small, medium and large leaves showing different kinds of damage by herbivores (see Example 1, Notes).

These are just two examples. We could do similar things with various other measurements in the observational notes. Figure 2.3, for instance, suggests that the number of *Babesia*-infected cells in different classes of host might have something to do with tick burden since both show the same broad association with class (see Fig. 2.1). Figure 2.4 hints at an association between the size of a leaf and the kind of damage it sustains. The Notes provide scope for other exploratory plots; try some for yourself.

Casting data in the form of figures like this is helpful not just because visual images are generally easier for most people to assimilate than raw numbers, but because they can expose subtleties in the data that are less apparent in numerical form. The plot of number of fights per encounter against distance from a burrow in Fig. 2.2, for instance, suggests that while the likelihood of fighting decreases further away from a burrow, there is considerable individual variation (the different symbols) within the trend. We shall see later that this variation



leads to some interesting insights into the role of burrows in the behaviour of both male and female crickets.

The sorts of plots we have used so far are helpful in seeing at a glance whether something interesting might be going on. However, the data could be presented in a different way to make the figures more informative. It is clear from the scattergram in Fig. 2.2 and the raw data in the notes relating to Figs 2.1, 2.3 and 2.4 that there is a lot of variation in the numbers recorded in each case. Adult male voles, for example, did not all have high *Babesia* burdens; indeed one male didn't have any infected cells at all in the sample examined. This variability has at least two important consequences as far as exploratory plots are concerned. First it suggests that simply plotting totals in Figs 2.1 and 2.3 is likely to be misleading, because the totals are made up from a wide range of numbers. A large total could be due to a single large result, with all the rest actually being smaller than the results contributing to the other, lesser, totals, in which case our interest in the apparent differences between the bars in the figures might diminish somewhat. Second, variability in the data might obscure some potentially interesting tendencies in scattergrams. What we need, therefore, is a way of summarising data so that: (a) the interesting features are still made clear, but (b) the all-important variability is also presented, though in a way that clarifies rather than obscures patterns in the data. In short, we need some *summary statistics*.

## 2.2.2 Summary statistics

The usual way of summarising a set of data so as to achieve (a) and (b) above is to calculate a *mean* (average) or *median* value and then to provide as a measure of the variability the associated *standard error* (for a mean) or *confidence limits* (for a median).

### Means and standard errors

The mean (often represented by  $\bar{x}$  ('*x*-bar')) is simply the sum of all the individual values in the data set divided by the number of values (usually referred to as  $n$ , the sample size). Formally, the mean is expressed as:

$$\bar{x} = (1/n) \sum_{i=1}^{i=n} x_i$$

The expression  $\sum_{i=1}^{i=n} x_i$  indicates that the first ( $i = 1$ ) to the  $n$ th ( $i = n$ ) data values ( $x$ ) are summed ( $\Sigma$  is the summation sign) and can be expressed more simply as  $\Sigma x$ . This summed value is then multiplied by  $1/n$  (equivalent to dividing by the sample size  $n$ ).

Since the mean is calculated from a number of values, we need to know how much confidence we can have in it. By 'confidence' we mean the reliability with which we could take any such set of values from the material and still end up with the same mean. A statistician would phrase this in terms of our sample mean ( $\bar{x}$ , the one we have calculated) reflecting the true mean (usually denoted  $\mu$ ) of the population from which the data values were taken. Suppose, for instance,

we measured the body lengths of ten locusts caught in each of two different geographical areas and obtained the following results:

Length (cm)	
Area 1	Area 2
6.3	6.0
7.1	6.4
6.2	6.3
6.5	6.0
7.0	5.9
6.7	6.5
6.5	6.1
7.0	6.2
6.8	6.2
7.1	6.4
$\bar{x}$ 6.7	6.2

Then suppose that we had obtained the following instead:

Length (cm)	
Area 1	Area 2
8.3	8.0
5.1	5.4
7.2	5.3
5.5	6.5
8.0	8.4
5.7	5.5
5.5	7.1
8.0	5.2
8.8	5.2
5.1	5.4
$\bar{x}$ 6.7	6.2

In both cases, the mean body lengths from each area are the same and we might want to infer that there is some difference in body size between areas. In the first case, the range of values from which each mean is derived is fairly narrow but different between areas. We might thus be reasonably happy with our inference. In the second case, however, the values vary widely within areas and there is considerable overlap between them. Now we might want to be more cautious about accepting the means as representative of the different areas. We can see that this is the case from the columns of numbers but we need some way of summarising it without having to present raw numbers all the time. We can do this in several ways. The most usual is to calculate the standard error, a quantitative estimate of the confidence that can be placed in the mean and which can be presented with it as a single number. The calculation is simple and is shown in Box 2.1.

### Box 2.1 Standard deviations and standard errors

The *standard deviation* (usually abbreviated to s.d. or SD) measures the spread of actual data values around the mean, on the assumption that these data follow the *normal distribution* (see p. 55); it is a measure of the confidence you have that any particular data value will fall within a particular range (the mean + 1 s.d. and the mean – 1 s.d., hence the mean  $\pm$  s.d.). The standard error of the mean, usually just called the *standard error* (abbreviated to s.e. or SE) measures the spread of multiple *sample means* around the true population mean. Normally you will only be taking a single sample, and hence the SE is an expression of the confidence you have that your sample mean falls within a particular range (mean  $\pm$  s.e.) of the true population mean. Since sample means are almost always normally distributed, almost whatever the distribution of the raw data, it is always OK to cite a standard error with your sample mean. The calculations are as follows:

#### The standard deviation:

1. Calculate the sum of all the data values in your group ( $\Sigma x$ ).
2. Square the individual data values and sum them, giving ( $\Sigma x^2$ ).
3. Calculate  $\Sigma x^2 - (\Sigma x)^2/n$  (remember that  $n$  is the sample size, the number of values in your set of data). This is actually a quick way of calculating the sum of squared deviations from the mean:  $\Sigma(x - \bar{x})^2$ . The deviations are squared so that positive and negative values do not simply cancel each other out.
4. Dividing by  $n - 1$  gives the *variance* of the sample, an important intermediate quantity in many statistical tests.
5. Taking the square root of the variance gives the *standard deviation*.

Most scientific calculators will give you the mean of a set of numbers and most will also give you the *standard deviation* (Box 2.1), usually represented as  $\sigma$  or  $s$ . If your calculator has both  $\sigma$  and  $\sigma_{n-1}$  buttons, it is the  $\sigma_{n-1}$  one that you want. The standard deviation will become important later, but for the moment we can simply use it to obtain the standard error. All we need to do is call up the standard deviation, square it, divide it by  $n$  and take the square root.

Whichever way you calculate the standard error (by hand or by calculator), it should be presented with the mean as follows:

$$\bar{x} \pm \text{s.e.}$$

The  $\pm$  sign indicates that the standard error extends to its value on either side of the mean. The bigger the standard error, therefore, the more chance there is that the true mean is actually greater or smaller than the mean we've calculated. We can see how this works with our locust data. Let's look at the two sets of values for Area 1, first calculating the  $x^2$  value of the first example:

Locust	Body length ( $x$ )	$x^2$
1	6.3	39.69
2	7.1	50.41
3	6.2	38.44
4	6.5	42.25
5	7.0	49.00
6	6.7	44.89
7	6.5	42.25
8	7.0	49.00
9	6.8	46.24
10	7.1	50.41
$n = 10$	$\Sigma x = 67.2$	$\Sigma x^2 = 452.58$

The steps are then:

1.  $\Sigma x = 67.2$
2.  $\Sigma x^2 = 452.58$
3.  $\Sigma x^2 - (\Sigma x)^2/n = 452.58 - (67.20)^2/10 = 452.58 - 451.58 = 1$
4. Divide by  $n - 1 = 1/9 = 0.11$
5. Divide by  $n = 0.11/10 = 0.01$
6. Take the square root =  $\sqrt{0.01} = 0.11$

Thus the mean length of locusts in the first example for Area 1 is:

$$6.72 \pm 0.11 \text{ cm}$$

If we repeat the exercise for the second example, however:

1.  $\Sigma x = 67.2$
2.  $\Sigma x^2 = 471.18$
3.  $\Sigma x^2 - (\Sigma x)^2/n = 471.18 - 451.58 = 19.6$
4. Divide by  $n - 1 = 19.6/9 = 2.2$
5. Divide by  $n = 2.2/10 = 0.22$
6. Square root =  $\sqrt{0.22} = 0.47$

The mean is now expressed as:

$$6.72 \pm 0.47 \text{ cm}$$

We could leave the mean and standard error expressed numerically like this, or we could present them visually in a bar chart. If we opt for the bar chart, then the mean can be plotted as a bar and the standard error as a line through the top centre of the bar extending the appropriate distance (the value of the standard

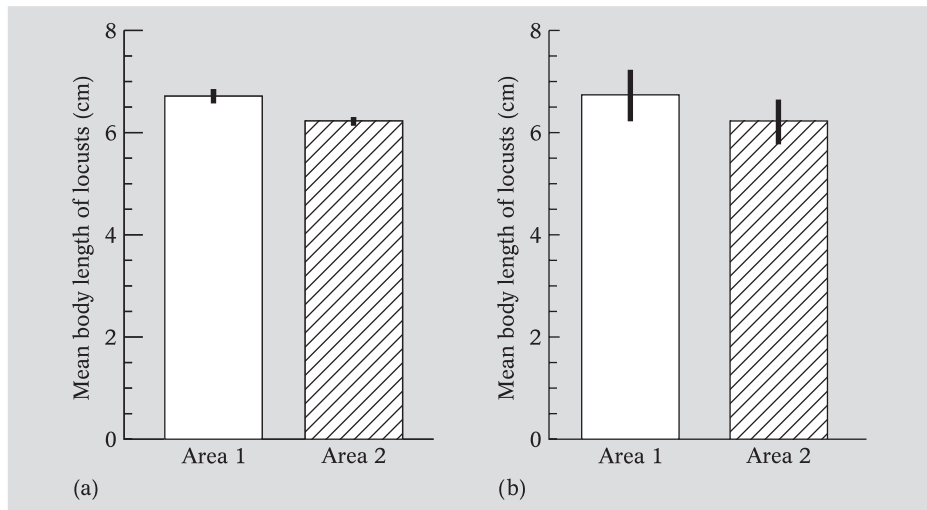


Figure 2.5

error) above and below the mean. Figure 2.5a, b shows such a plot for the two sets of example locust data.

### Medians and confidence limits

An alternative summary statistic we could have used is the median. There are good statistical reasons, to do with the distribution of data values (*see* Chapter 3), why we may need to be cautious about using means and standard deviations (but not standard errors – see Box 2.1 – unless sample sizes are very low). The use of standard deviations in particular makes important assumptions about the distribution of the data from which they are calculated that may not hold in many cases. Using medians and confidence limits avoids these assumptions. Later, we shall see that statistical tests of significance can also avoid them. Finding the median is simple. All we do is look for the central value in our data. Thus, if our data comprised the following values:

5 7 11 21 8 12 14

we first rank them in order of increasing size:

5 7 8 11 12 14 21

and take the value that ends up in the middle, in this case 11. If we have an even number of values so there isn't a single central value, we take the halfway point between the two central values, thus in the following:


2 6 8 20 23 38 40 85

the median is 21.5 (halfway between 20 and 23). Note that the median may yield a value close to or very different from the mean. In the first sample, the mean is 11.1 and thus similar to the median. In the second sample, however, it is much greater at 27.58.

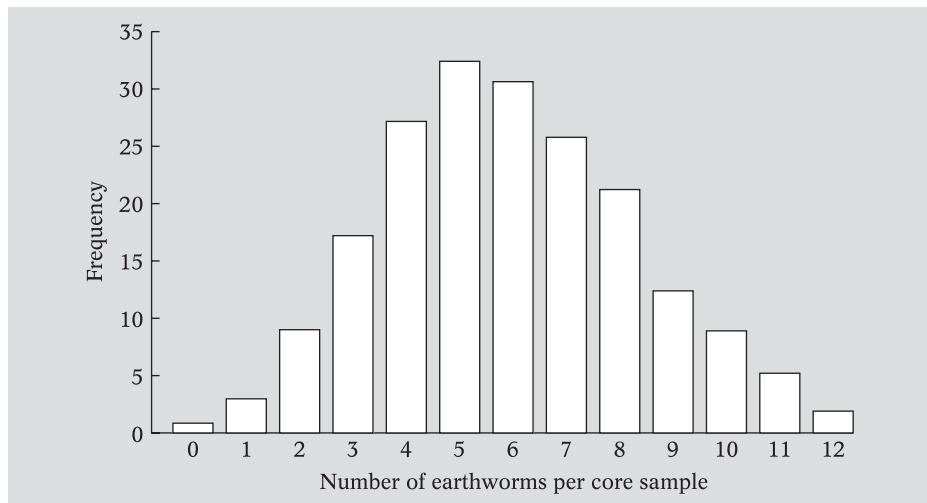
Again, we want some way of indicating how much confidence to place in the median. By far the simplest way is to find the confidence limits to the median using a standard table, part of which is shown in Appendix I. All we need to do is rank order our data values as before, count the number of values in the sample ( $n$ ), then use  $n$  to read off a value  $r$  from the table. Normally, we would be interested in the  $r$ -value appropriate to confidence limits of approximately 95 per cent ('approximately' because, of course, the limits always have to be two of the values in the data set, one above the median and one below – if  $n$  is less than 6, 95 per cent confidence limits cannot be found). This  $r$ -value then dictates the number of values in from the two extremes of the data set that denote the confidence limits. Thus in our first sample data set, there are seven values. Reference to Appendix I shows that for  $n = 7$ ,  $r = 1$ ; the confidence limits to the median of 11 are therefore 5 and 21. However, if we had a sample of nine values (say 7, 11, 15, 22, 46, 67, 71, 82, 100)  $r$  for approximately 95 per cent confidence limits is 2, so for a median of 46, the limits would be 11 and 82.

As with the mean and standard error, we can represent medians and their associated confidence limits visually in a bar chart.

Another common way of expressing the distribution of values about the median is the interquartile range, which is the difference in values between the data point one-quarter of the way down the rank order of values in the sample and the point three-quarters of the way down. This is often presented in the form of a so-called 'box and whisker' plot (Fig. 2.6). This shows the median value as a bold bar within a box that represents the interquartile range. The 'whiskers' are



**Figure 2.6** A box-and-whisker plot summarising the effect of a drug treatment and control on patient response. See text for details. Note the outlier indicated for the eleventh data point in Treatment 1.

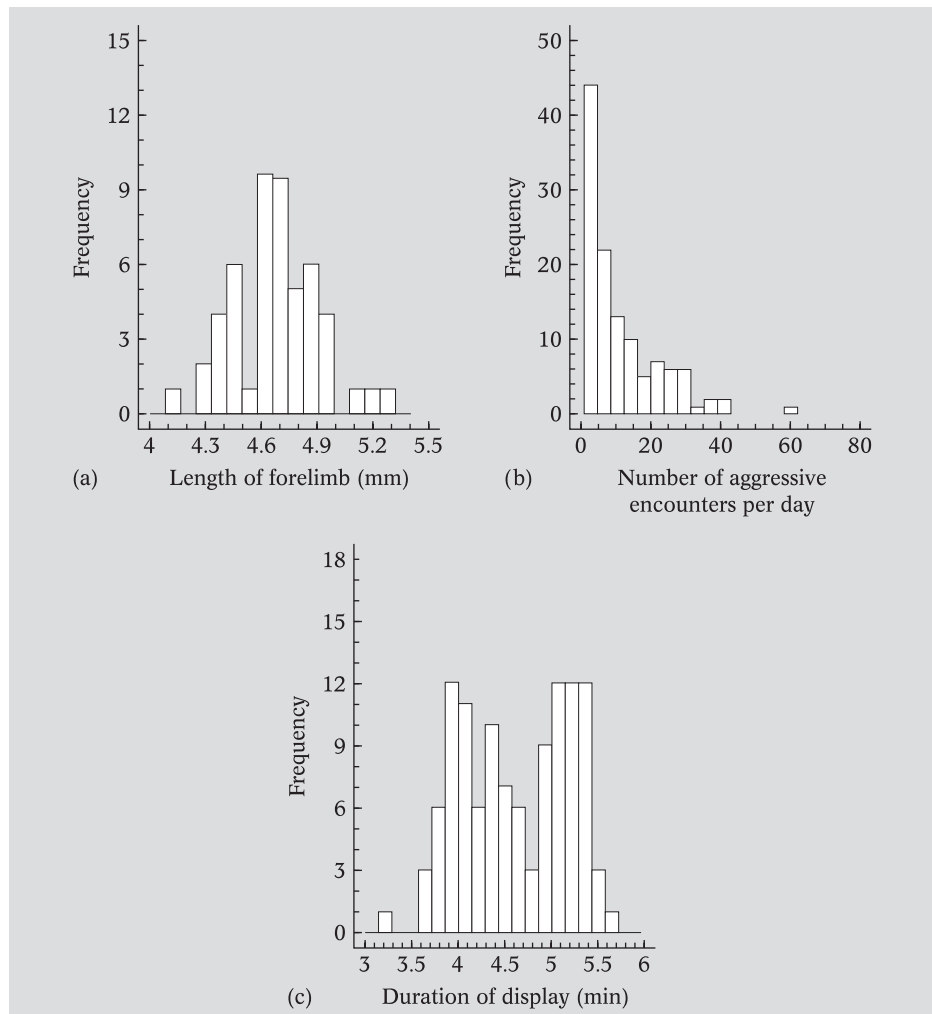


**Figure 2.7** A bar-chart frequency distribution of different numbers of earthworms recovered from core samples of soil (modified from Dytham, 2003).

not error bars in this case, but extend to the largest value in the sample that is within 1.5 interquartile ranges of the top or bottom of the box. Some computer statistical packages also present ‘outlier’ values as additional points on box and whisker plots; these can provide useful alerts to typographical errors in data entry.

### Frequency distributions

Means and standard errors, medians and confidence limits, then, are two conventional ways of summarising central tendency and spread of values within data. As we have seen, they are particularly useful in making quick ‘eyeball’ comparisons between two or more data sets. Such ‘eyeball’ comparisons, however, are only one reason why visual summaries of data sets can be worth plotting. Another is to allow the features of a single set of data to be explored fully, perhaps to examine the distribution of values and decide on an appropriate method of confirmatory analysis (*see later*). The usual way of doing this is to plot a *frequency distribution* of the values (the number of times each occurs in the data set), either as a bar chart, like the means in Fig. 2.7, or as a histogram. Bar charts, in which the bars in the figure are separated by a small gap along the horizontal ( $x$ ) axis, are used to plot the distribution of discrete values where there are sufficiently few of these to make such a plot feasible. Histograms, in which the bars are contiguous, are used to plot the distribution of classes of values (e.g. 1–10, 11–20, 21–30, etc.), usually where there is a wide and continuous spread of individual values. Figure 2.8 shows a range of frequency distribution histograms. It is obvious almost at a glance why we urged caution in using the mean as a general measure of central tendency and why frequency distributions can be a crucial first step in deciding how to analyse and present data. In Fig. 2.8a the distribution is more or less symmetrical with a peak close to the centre. Formally the peak is referred to as the *mode* and, in symmetrical, unimodal distributions like Fig. 2.8a, the mode and the arithmetic mean amount to the



**Figure 2.8** Histogram frequency distributions of: (a) the length of the right forelimb of field cricket (*Gryllus bimaculatus*) nymphs, (b) the number of aggressive encounters per day between male mice (*Mus musculus*) in an enclosure and (c) the duration of the aggressive displays of male Siamese fighting fish (*Betta splendens*) in pairwise encounters.

same thing. In these cases, therefore, the mean is an adequate measure of central tendency. In Fig. 2.8b, c, however, it is clear that the mean and mode are far from the same thing (in fact Fig. 2.8c has two modes and the distribution is said to be *bimodal*). Here, it makes little quantitative sense to use the arithmetic mean as a measure of central tendency. The shape of these distributions becomes extremely important when deciding on further analysis of the data, as we shall see shortly.

Of course, when plotting such distributions, one has first to decide on the number of categories into which to cast the data along the  $x$ -axis. Perhaps not surprisingly, there is no hard and fast rule about this. Dytham (2003), in his excellent book, advocates 12–20 categories as a rough rule of thumb, or, as an



alternative,  $\sqrt{n}$ , where  $n$  is the number of data points in the sample. But, in the end, it's a matter of judgement and common sense.

Exploratory analysis, then, is a way of summarising data, visually or numerically, to make it easier to pick out interesting features. The calculations and plots we've suggested here can, of course, be produced by various computer packages rather than by hand. All the main statistical packages, like SPSS, Minitab, Statistica, GenStat, Statgraphics, GLIM, S-Plus and R, can do this, but many people prefer to enter their data into a general spreadsheet for organising and checking before transferring them into a statistical package. Microsoft's Excel is by far the most widely used of these spreadsheets, partly because it allows a range of exploratory analyses via simple programmable formulae, and easily produces data tables and plots (though see Box 2.3) that can be 'cut and pasted' into other documents or PowerPoint slides. Data in Excel can also be read directly into several different statistical packages, so provide a convenient basis for further, more sophisticated analysis. We shall use Excel as our example spreadsheet package throughout the book; our example output uses the Microsoft Office 2003 version of the package. As our example commercial statistical package, we have chosen SPSS (*Statistical Packages for the Social Sciences*), again because it is used widely in higher education and research and, with an important exception we shall introduce later, offers the range of tests required. The version we present in the screen captures is v13. Box 2.3 summarises how to do the main kinds of exploratory analyses in Excel and SPSS, but before looking at that, it is worth considering some general 'dos' and 'don'ts' in preparing data for analysis using computer packages. These are summarised in Box 2.2.

### Box 2.2 Preparing data for analysis using computer packages

Although the modest data sets acquired during practical classes and projects can often be explored and analysed by hand without difficulty, it is now almost unheard of to do so. Instead, data are usually entered into one of the array of computer spreadsheet, database or statistical packages available commercially or via the Internet/World Wide Web, and analysed using these. Needless to say, these provide a quick and convenient means of dealing with data and are the virtually universal instruments of choice these days. However, as anyone who has ever used computer programs will know, it is important to maintain a healthy level of wariness when using them. The all-too-human potential for programmer and/or user error can lay serious traps for the unwary. On top of that, the immediacy of the results – click a button and there they are – does away with any need to understand what is

actually going on. *When it comes to statistics, therefore, computer packages should be seen as a convenient adjunct to a proper formal training in the subject, not as a quick-fix substitute for it.* It is thus impossible to over-emphasise the need to *check* and *double-check* that a package has done what you think it has when you carry out an operation – utterly meaningless rubbish can look perfectly plausible on superficial inspection.

While data can be entered into any program of your choice, we strongly recommend using Microsoft's spreadsheet package Excel, rather than a statistical package *per se*, for entering and organising your data ahead of analysis. Excel is a powerful program that allows data to be explored, prepared and analysed in many different ways. It comes as part of the standard installation of packages on many computers when they are purchased, so is widely used and supported. Data in Excel

can also be imported easily into most statistical packages, so Excel is a convenient ‘universal donor’ when it comes to switching between different packages. The same is not true for many of the statistical packages themselves, where data often have to be converted into other formats (often Excel) before being exportable elsewhere, sometimes risking corruption on the way. Excel’s power and flexibility, however, is a two-edged sword. On the one hand it allows large data sets to be organised and explored easily; on the other, it provides a wealth of opportunities for error and confusion. Two of the commonest user problems are mistakes in using formulae to calculate new variables from existing ones (*see below*), and mistakes in using menu commands to reorganise data within a worksheet (most often inadvertently selecting only data visible on-screen, instead of the whole data set, when using the ‘Sort’ command, thus shifting some data out of line with the rest). The key to using Excel to prepare data for use in statistical packages is ‘*keep it simple*’.

### Formatting data worksheets

When you boot up Excel, you are presented with the empty rows and columns seen in Fig. (i). This simple layout is all you need. Excel allows a vast range of row/column formats within its worksheets, and the tendency among students is to make liberal use of this capability, with header rows spanning several columns, worksheets subdivided into sections, textual notes and keys interspersed with data,

data split across different worksheets and so on (e.g. Fig. (ii)). While all this is fine if the file is simply being used as a spreadsheet in its own right, it is completely unworkable if you want to import the data into a statistical package. As a general rule, statistical packages expect data in simple column format, with each variable occupying a separate column and the first cell of the column containing the name of the variable. Each row then represents one study sample to which the data in each column cell along that

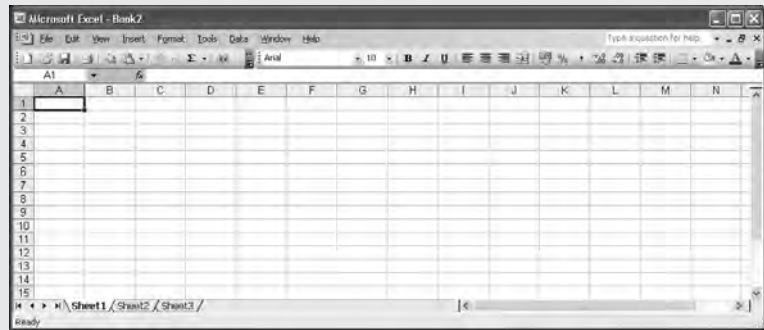


Figure (i) A newly opened Excel screen awaiting data entry.

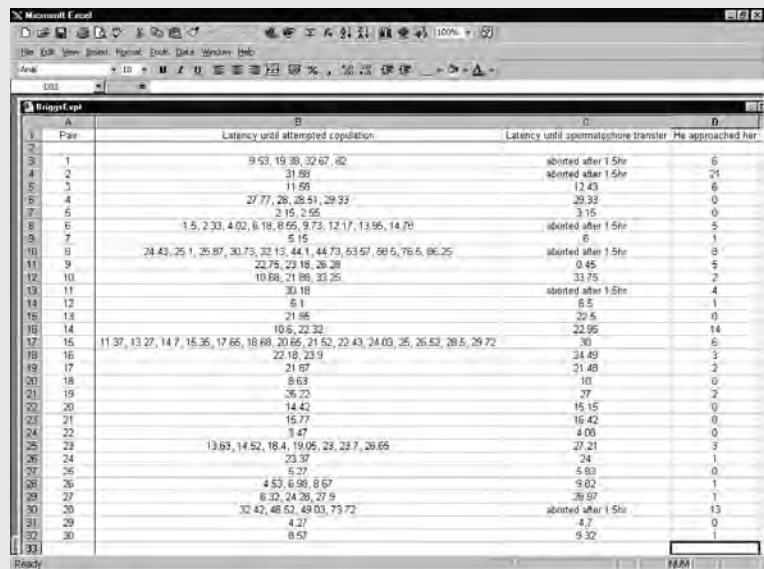


Figure (ii) An example of an Excel worksheet format that would not be suitable for reading into a statistical package. Problems include: multiple data within cells, combinations of numerical data and textual comment within columns, and unsuitable column headings.

row relate. Thus, suppose we had conducted an experiment looking at associations between anogenital sniffing (often a prelude to aggression) and circulating levels of the hormones testosterone (as a measure of sex hormone activity) and corticosterone (as a measure of stress levels) in male mice. Let's also suppose we had allocated the mice to four different experimental treatments, say different degrees of complexity of their group's home cage environment (e.g. (1) bare cage, (2) cage + nest material, (3) cage + nest material + nest boxes, (4) cage + nest material + nest boxes + shelves). The experiment consisted of introducing randomly chosen mice from eight different cages within each treatment into a clean empty cage and allowing them to interact one at a time with another randomly chosen mouse (different in each case) that had been kept on its own in an empty cage to standardise the social experience of the opponent across experimental subjects. Each pair of mice was allowed to interact for five minutes during which the number of anogenital sniffs initiated by the subject towards the opponent was recorded.

To set up our Excel data file for the experiment, we allocate one column to each of the variables in which we are interested. Thus our first column might be *Mouse* and consists simply of numbers 1–32 so we know which mouse each row refers to. So we type 'Mouse' into the first cell in Column A and enter 1–32 in the successive cells below. Next, we want to note the experimental treatment for each mouse, so we head another column

*Treatment*. Here, the best thing is to use a simple numeric code for each treatment, in this case '1' for bare cages, '2' for cage + nest material, and '3' and '4' for the other two cage types. Avoid using letters (e.g. A, B, C, D) or combinations of letters and numbers (e.g. T1, T2, T3, T4), because statistical packages sometimes treat these as text variables and won't allow them to be used in analyses. We can then complete three further columns for the number of sniffs initiated and the testosterone and corticosterone concentrations for each mouse, so the basic data set will appear as in Fig. (iii). Notice how the header row of variable names only occupies a single row, and each variable name is a *single* short word. If you *have* to use more than one word, use the underscore (\_) rather than a space in between the words; doing this avoids all sorts of potential problems when transferring the data to a statistical package, so we strongly recommend that you do.

We could, of course, add other measures if we wished, such as the body weight of each subject mouse, the time of day of each test, and so on, each in its own headed column. Depending on the nature of the variables, we could also calculate new ones within the file itself using Excel's formula capability (see also Boxes 2.3 and 2.4). Suppose, for instance, we wished to express anogenital sniffing in terms of the *rate* of sniffing rather than total number. Since all mice were observed for five minutes, this is simply a matter of dividing the number of sniffs by five to get the rate per minute. To do this, go to the first blank column after the existing data (here Column F), head it *Sniffrate* (or similar) and click on the first empty cell below the heading (Cell F2). Now type an '=' sign, followed by C2/5 (so '=C2/5') and hit 'Return'. This produces the number 0.8 in the cell, which is the number of sniffs (from cell C2) divided by 5. There's no need to repeat this for every cell of the *Sniffrate* column; all we need

Mouse	Treatment	Sniffs	Testosterone	Corticosterone
1	2	4	3.2	21.1
2	1	2	2.1	19.6
3	3	7	3.9	17
4	4	11	6.2	54.3
5	1	3	4.1	20
6	2	6	4.5	32.2
7	3	8	3.2	17.9
8	4	9	4	23
9	4	14	7.5	66.4
10	2	7	6.5	52.1
11	3	5	5.7	46.6
12	1	1	4	18.4
13	2	7	4.5	25.8
14	3	12	5.5	30

Figure (iii) Part of an Excel worksheet in a suitable format for reading into a statistical package.

column after the existing data (here Column F), head it *Sniffrate* (or similar) and click on the first empty cell below the heading (Cell F2). Now type an '=' sign, followed by C2/5 (so '=C2/5') and hit 'Return'. This produces the number 0.8 in the cell, which is the number of sniffs (from cell C2) divided by 5. There's no need to repeat this for every cell of the *Sniffrate* column; all we need

to do is click on cell F2 again, then click on 'Edit' at the top of the screen, followed by 'Copy' (or just click on the 'Copy' symbol itself at the top of the screen). The cell is now highlighted by a shimmering dashed border, which indicates its contents have been copied to the clipboard. To complete the column, click on the cell below F2 and, holding down the 'Shift' key as you do it, the last cell in the column (here F25). The remainder of the column is now highlighted. Click 'Edit' and, this time, 'Paste' (or just click the 'Paste' symbol directly) to fill in the rates for each remaining cell. What Excel does here is to copy the actions of the formula you typed in (rather than the numerical content of Cell F2 itself) to each of the other cells you highlighted.

While calculating new variables like this in Excel is easy, we recommend that you keep it to a minimum. This is because worksheets with variables calculated using formulae may not be read properly by some statistical packages, or require special steps to ensure they are read properly. Since all these calculations can be done just as easily within statistical packages themselves, using menus rather than embedded formulae, our advice is to do them there rather than in Excel.

### Reading Excel data into SPSS

To read the data in Fig. (iii) into SPSS directly from Excel, simply open SPSS, then click 'File' and 'Open' on the opening screen. Select the directory containing the Excel file in the 'Look in' box, and change the 'Files of type' box at the bottom from the default 'SPSS' to 'Excel' in order to display the available Excel files. Double click on the file required, and, in the 'Opening Excel Data Source' box make sure the 'Read variable names' box is ticked. Click 'OK' and the data, complete with variable names, will appear\* (Fig. (iv)a). Note, however, that if you hadn't chosen variable

names of appropriate length, SPSS may have truncated them, depending on the version you are using. You may have to do some editing, in

	mouse	treatment	sniffs	testosterone	corticos	
1	1	2	4	3.2	21.1	
2	2	1	2	2.1	19.6	
3	3	3	7	3.9	17.0	
4	4	4	11	6.2	54.3	
5	5	1	3	4.1	20.0	
6	6	2	6	4.5	32.2	
7	7	3	9	3.2	17.9	
8	8	4	9	4.0	23.0	
9	9	4	14	7.6	65.4	
10	10	2	7	5.5	52.1	
11	11	3	5	5.7	45.6	
12	12	1	1	4.0	16.4	
13	13	3	7	4.6	26.8	
14	14	3	12	5.5	30.0	
15	15	1	2	3.1	18.9	
16	16	4	17	7.7	22.2	
17	17	4	12	7.0	45.0	
18	18	3	3	3.2	36.7	
19	19	2	3	4.7	19.8	
20	20	1	1	4.7	19.0	
21	21	2	4	3.6	34.2	
22	22	3	15	5.4	66.0	
23	23	1	5	4.1	37.8	
24	24	4	19	6.7	71.2	
25	25	3	12	4.9	65.3	
26	26	2	11	5.9	49.8	
27	27	4	17	6.0	57.0	
28	28	1	3	3.0	21.8	
29	29	4	9	5.2	36.2	
30	30	1	1	2.8	33.2	
31	31	2	1	3.6	19.6	
32	32	3	2	4.1	17.3	
33	33	1	6	5.7	38.7	
34	34	2	4	3.4	34.0	
35	35	3	3	3.5	36.0	
36	36	4	12	6.9	52.8	
37	37	2	3	6.0	37.0	
38	38	1	5	4.0	35.0	
39	39	4	18	4.2	67.5	
40	40	3	9	5.0	42.1	

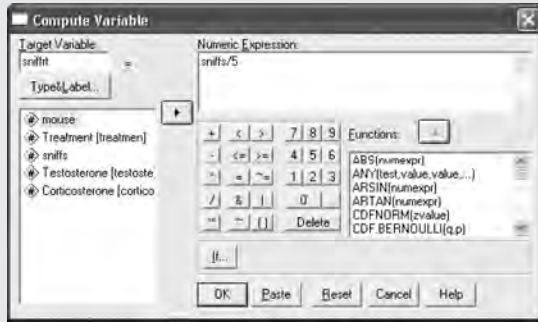
(a)

Name	Type	Width	Decimals	Label	Values	Missing	Columns	Align
1 mouse	Numeric	11	0		None	None	8	Right
2 treatment	Numeric	11	0	Treatment	None	None	8	Right
3 sniffs	Numeric	11	0		None	None	8	Right
4 testosterone	Numeric	11	1	Testosterone	None	None	8	Right
5 corticos	Numeric	11	1	Corticosterone	None	None	8	Right

(b)

**Figure (iv)** (a) Data from Fig. (i) after having been read into SPSS. Note the truncated variable headings (as may happen in some versions). (b) The 'Variable view' screen of the same data, in which the column headings can be edited.

\* If you are using older versions of SPSS, you may need to save your Excel file in version 4 or earlier format before reading it in. The best way of ensuring that your data can be read is to save them in Excel version 3, because this is the last version that only used single worksheets. It is the multiple sheets of later versions that can cause problems.



**Figure (v)** The ‘Compute’ dialogue box in SPSS, set up for calculating a new variable (*Sniffrt*) for sniffing rate from the number of sniffs (*Sniffs*) and the duration of tests (five minutes) (see above).

which case, click on ‘Variable view’ at the bottom of the screen, highlight the relevant variable name(s) and type in your changes (Fig. (iv)b). You can also make other changes here, such as altering the number of decimal places displayed. Click on ‘Data view’ to restore the data screen, now with renamed variables and anything else you changed.

An alternative way of doing things is to cut and paste the data into SPSS. In Excel, highlight

the entire data set, excluding the variable names, and ‘Edit’ and ‘Copy’ (or just click the ‘Copy’ symbol at the top of the screen) to copy it to the clipboard. Open SPSS, click on the first empty cell in the first column and click ‘Edit’ and ‘Paste’ (or click the ‘Paste’ symbol). The data will appear, but not the variable names; instead the columns will be headed ‘var0001, var0002’, etc. To change this, click on ‘Variable view’ and edit the names as above.

To calculate new variables from the entered data, as above in Excel, click on ‘Transform’ at the top of the screen and select ‘Compute’. Using the example of calculating rate of sniffing again, enter the new variable name (e.g. *Sniffrt*) into the ‘Target Variable’ box at top right in the dialogue box (Fig. (v)), transfer *Sniffs* into the ‘Numeric Expression’ box by highlighting it and clicking the forward arrow (or simply type it into the box manually), click on the forward slash (division operator) in the pad of options beneath the ‘Numeric Expression’ box (or again type it in), type ‘5’ (the number of minutes per test) and click ‘OK’. This will return you to the data, with the new variable appearing in the first previously blank column at the end.

### Box 2.3 Exploratory analysis using computer packages

Various spreadsheet and database, and pretty well all statistical, packages written for computers will enable you to explore data sets and calculate standard summary statistics. The means of doing so are usually simplest, and the range of options greatest, in statistical packages because calculating statistics is their stock in trade. In spreadsheets and databases, they’re often something of a frill and can be tricky to get at.

#### Exploratory analysis using Excel

Excel allows various kinds of exploratory analysis, some very straightforward, others requiring more complicated procedures that highlight Excel’s limitations on the statistical side. Again,

our advice would be to use Excel as a simple spreadsheet for preparing raw data and reserve statistical treatment, even of an exploratory kind, for a proper statistical package. Some basic analyses that can be done in Excel, however, can be summarised as follows:

#### *Calculating summary statistics in Excel*

Most useful summary statistics can be calculated by using one or other of the menu functions at the top of the screen (see Box 2.2, Figs (i) and (iii)), or directly by typing in a simple formula. For example, using the sample data in Box 2.2, Fig. (iii), we could calculate a range of basic summary statistics for the columns or rows by

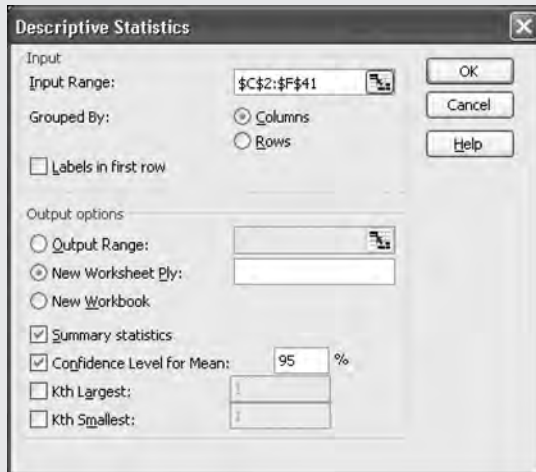


Figure (i) The 'Descriptive Statistics' dialogue box in Excel.

clicking on 'Tools' followed by 'Data Analysis'\* and 'Descriptive Statistics'. This generates the dialogue box in Fig. (i) above. If we wish to obtain summary statistics for each of our columns of data in Box 2.2, Fig. (iii), we can tick the 'Summary statistics' and 'Grouped by: Columns' options in the box in Fig. (i), and enter the first and last data cells of the data we wish to include in the 'Input range:' box. For the data in Box 2.2, Fig. (iii), this would be cell C2 (the first data value in column C [the first cell is taken up with the name of the variable]), and cell F41 (we're not interested in summary statistics for the first two columns, because these are just coding variables). In Excel format, the range is entered as C2:F41. Clicking the 'OK' button now results in a list of summary statistics for each column of data. These include some statistics relating to the shape of the frequency distribution of values in the column, such as *kurtosis* and *skewness*, which we shall deal with in the next chapter. Also included are

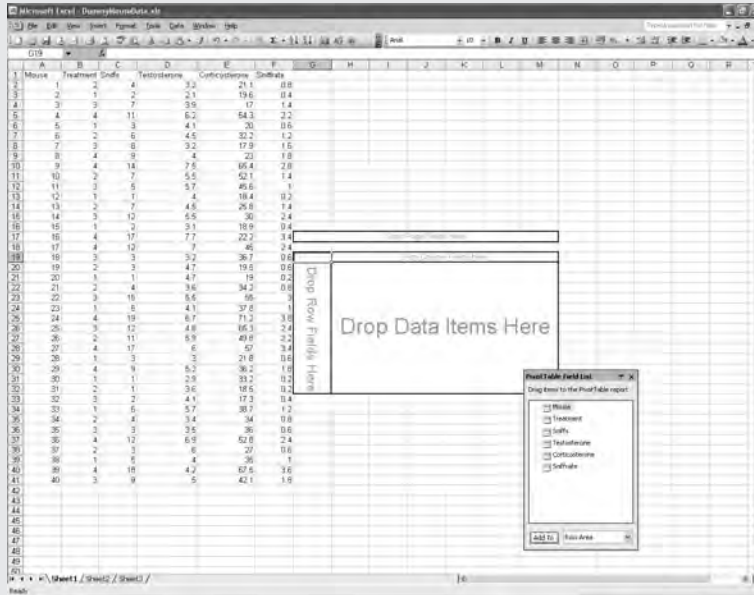
\* If you're using later versions of Excel, you will need to install the 'Analysis toolpack' before 'Data analysis' appears under the 'Tools' menu. To do this, click on 'Add-ins' within 'Tools', select the analysis toolpack you want and click 'OK' - 'Data analysis' will now appear in the drop-down 'Tools' menu.

all the 'standard' statistics commonly used to summarise data: the mean and standard error, the median and mode, the standard deviation and variance, and the range, sum and number of values. Selecting the 'Confidence' option in Fig. (i) will add the 95 per cent confidence limits to the mean to the list. Unless otherwise instructed, Excel will present the summary statistics in a new sheet of the active data workbook. Clicking on the 'Output range' button and specifying a new cell in the current worksheet, however, will result in the summary being presented next to the data themselves.

Alternatively, summary statistics can be obtained individually by selecting an empty cell in the data sheet, clicking the 'Paste function' ( $f_x$ ) button at the top of the screen (Fig. (ii)) and selecting from the 'Statistical' menu (note that Excel uses 'AVERAGE' for the mean). Irritatingly, this menu does not offer the standard error to the mean, but this can be calculated easily from the standard deviation using a simple formula (see Box 2.4). Indeed, such formulae are what the 'Paste function' itself uses; these can be seen by clicking on the cell containing the calculated summary statistic and looking in the '=' row above the data columns. Once people are familiar with them, the direct use of formulae usually becomes the preferred method, since it avoids having to click through the succession of options in the dialogue boxes. Box 2.4 shows some examples.

#### **Using pivot tables to summarise by category.**

Often, we want to summarise according to categories of data within a data set, such as, for example, the different cage treatments in the mouse data in Box 2.2, Fig. (iii). This can be done very simply for total values using the so-called 'pivot table' facility. Click on 'Data' at the top of the screen (Box 2.2, Fig. (iii)) and select 'Pivot Table and Pivot Chart Report...'. Let's assume we're interested in knowing how the total number of sniffs varies with treatment in the data set in Box 2.2, Fig. (iii). Select 'Microsoft Excel list or database' as the source of data in the first pivot table dialogue box and click 'Next'. In the 'Range box' of the new dialogue box, enter the inclusive



(a)

Treatment	Total
1	29
2	50
3	76
4	139
Grand Total	293

(b)

Figure (ii) (a) Setting up a 'pivot table' in Excel by dragging variable headings into the Rows and Columns box; (b) the resulting 'pivot table' in the worksheet.

cells of the data set; here we enter B1:F41 including the variable names, note, because the pivot table will use these to label the output. Clicking 'Next' produces the outline of the output table and a request to drag the 'field buttons' (corresponding to our selectable columns of data) on

the right into the 'Row' or 'Data' compartments of the table (Fig. (ii)a). In this case, therefore, *Sniffs* is dragged into the 'Data' compartment (where it now says 'Sum of Sniffs'), since we're interested in variation in this with respect to treatment, and *Treatment* is dragged into the 'Row' compartment. Click 'Next' again, and select a target cell in the worksheet to receive the table, or select 'New worksheet' if you want it in a separate sheet, then click 'Finish' and the table will appear (Fig. (ii)b).

Sometimes, we might want to summarise by more

than one category within our data (for instance, if we had recorded the number of sniffs by mice with respect to both their cage treatment *and* their social rank [dominant or subordinate] in their home cage). Where totals are concerned, this can be done easily once again using 'pivot tables'. Follow the procedure above, except this time specify three columns (the two categories and the variable of interest) in the 'Range box', and drag *both* category columns into the 'Row' compartment of Fig. (ii)a. Select a target cell for the table and click 'Finish', and a table of totals by one category within the other will appear.

Of course, it would be useful to be able to produce mean or median values along with their respective errors by category. Unfortunately, Excel provides no simple way of doing this, again underlining its limitations as a statistical package. While there are circuitous routes through which it can be done, it is really not worth the effort when the whole operation can be accomplished in a couple of clicks of the mouse in a statistical package like SPSS.

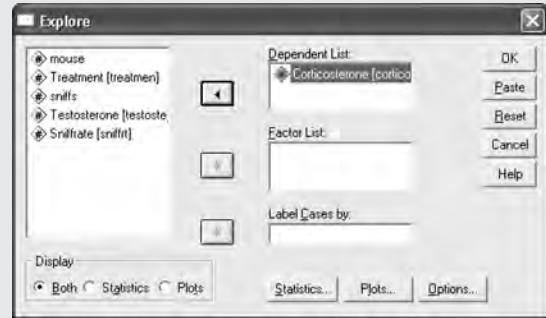
### Exploratory data plots in Excel

Using Excel to produce exploratory data plots like those in Figs 2.1–2.6 is straightforward, with the important exception of producing plots of

mean/median values with their individual error bars. The figure-plotting menu does offer standard errors to means, but calculates a single standard error for the mean of the entire data set, not a separate error for the mean of each category within it, which is usually what is wanted. Again, there are long-winded ways around this, but it is far simpler to do the whole thing in a statistical package, which is what we recommend.

**Plots of totals.** Summary plots of totals by category, like Fig. 2.1, however, are easy to produce via the ‘pivot table’ function. The pivot table in Fig. (ii)b presents the kind of information we need. To produce a graph, all we need to do is click on the ‘Chart wizard’ option (the vertical coloured bars button) at the top of the screen, select ‘Column’, and the first chart subtype within that, and click ‘Next’. Selecting the four treatment totals in the pivot table will enter their cell numbers in the ‘Data range’ box and produce a bar chart of them. Note that Excel defaults to scaling the x-axis as 1-n (here 1–4); if you wish to change this so the axis says, for example, ‘2, 4, 6, 8’ or anything else, you need to enter your new labels into an empty column in your data set (so, here, just four cells of the column), right-click on the figure, select ‘Source Data’, then ‘Series’. Put the cursor in the ‘Category (X) axis labels’ box and highlight the cells containing your desired labels; their specification will appear in the box. Clicking ‘OK’ will take you back to the now amended figure. Click ‘Next’ to edit as desired, or ‘Finish’ for the figure as it is. However you choose to edit the figure, we suggest you get rid of the default grid lines by right-clicking on the plotting area, clicking the ‘Gridlines’ option, and unchecking the relevant checkbox.

**Scatterplots.** To produce a scatterplot, like Fig. 2.2, highlight the two columns of data containing the X and Y data you want to plot, and click again on the ‘Chart wizard’ button and select ‘XY (Scatter)’. Clicking on ‘Finish’ will then produce the scatterplot you want. If the X and the Y data are in columns that are not adjacent, then highlight the data in the first column, and press and keep down the Ctrl key while you highlight the matching



**Figure (iii)** The ‘Explore’ dialogue box under ‘Descriptive Statistics’ in SPSS, with Corticosterone from the data in Box 2.2, Fig. (iii) transferred into the ‘Dependent List’ for summary.

data in the other column. Then press the ‘Chart wizard’ button and follow the procedure above.

**Frequency distributions.** Unfortunately, Excel does not provide a ready means of plotting frequency distributions. You will need to use a statistical package to do this (*see below*).

### Exploratory analysis using SPSS

SPSS is designed as a comprehensive statistical package, as opposed to a spreadsheet with some statistical facilities added on, like Excel. Consequently, life is usually more straightforward when it comes to performing analyses, at least at the simpler end of the spectrum. Data can either be entered into SPSS manually, or imported from other packages, including Excel (*see Box 2.2*).

#### Calculating summary statistics in SPSS

To obtain summary statistics for any of the variables in the data set, go to ‘Analyze’ at the top of the screen, and select ‘Descriptive statistics’ followed by ‘Explore’. Enter the variables for which statistics are required by highlighting them in the list on the left and clicking the arrow to transfer them into the ‘Dependent List’ box (Fig. (iii)). The output can be in terms of a table of numerical statistics, box-and-whisker plots (like Fig. 2.6), or both, depending on what is selected in the ‘Display’ panel on the bottom left. Choosing just ‘Statistics’ provides the table of output, which presents all the standard summary statistics – mean, standard



error, standard deviation, median, interquartile range, etc. – in a fashion similar to Excel above. If more than one variable had been selected and transferred to the ‘Dependent List’ box, the table would present summary statistics for each in turn. Summary statistics for selected variables can also be subdivided by category. So, for example, entering *Treatment* (Box 2.2, Fig. (iii)) into the ‘Factor List’ box in Fig. (iii) will produce the list of summary statistics by treatment for each variable entered into the ‘Dependent List’.

### *Exploratory data plots in SPSS*

Producing exploratory plots of data in SPSS is much more straightforward than in Excel, with most requirements being catered for directly in the menus.

*Plots of totals.* To plot total sniffs by treatment as we did above, for example, simply click on ‘Graphs’ at the top of the data screen and select ‘Bar’. Choose ‘Simple’ and ‘Summaries for groups of cases’, then click ‘Define’. Enter *Sniffs* in the ‘Variable’ box and *Treatmen* in the ‘Category axis’ box. Note that the ‘Variable’ box defaults to mean values; to change this to totals, highlight ‘MEAN(sniffs)’, click on the ‘Change summary’ button below it, select ‘Sum of values’ and click ‘Continue’. Click ‘OK’ and the figure output will appear. Double-clicking on the figure itself then takes you into the chart editor where the appearance and labelling can be altered as required.

*Plots of means or medians.* While the ‘Bar’ option above allows mean values to be plotted, it does not, for some bizarre reason, provide a direct means of adding standard error bars to the figure. Instead, the simplest way to produce a figure of means with

standard errors within the ‘Graphs’ menu is to select ‘Error bar’. Clicking on this, selecting ‘Simple’ and ‘Summaries for groups of cases’ again and clicking ‘Define’ takes you to a similar dialogue box to before. Once again, enter *Sniffs* in the ‘Variable’ box and *Treatmen* in the ‘Category axis’ box, but this time go to the ‘Bars represent’ box, select ‘Standard error of mean’ and change the ‘Multiplier’ setting to ‘1’. Clicking ‘OK’ produces the figure, which can then be edited to requirement.

To plot medians and interquartile ranges as a box plot, go to ‘Boxplot’ in the ‘Graphs’ menu. Again, choose ‘Simple’ and ‘Summaries for groups of cases’ and click ‘Define’. Enter *Sniffs* and *Treatmen* in the ‘Variable’ and ‘Category axis’ boxes, respectively, and click ‘OK’ for the figure.

*Scatterplots.* To generate a scatterplot of the relationship between two variables, select ‘Scatter’ from the ‘Graph’ menu, and click ‘Simple’ then ‘Define’. Enter the desired variables (say, *Sniffs* and *Testoste* in Box 2.2, Fig. (iv)) in the ‘Y-axis’ and ‘X-axis’ boxes and click ‘OK’.

*Frequency distributions.* To produce plots of frequency distributions for variables, click on ‘Analyse’, then ‘Descriptive Statistics’ and ‘Frequencies’. Select the variables of interest and transfer them into the ‘Variable(s)’ box. Select ‘Chart’ from the options at the bottom of the dialogue box and choose either ‘Bar charts’ then ‘Frequencies’ if the data are discrete values, like those in Fig. 2.7, or ‘Histograms’ if they are continuously distributed, like those in Fig. 2.8. Under ‘Histograms’, there is a ‘With normal curve’ option, which, if selected, fits a normal curve to the resulting plot, thus giving an eyeball impression of how well the distribution conforms to a normal distribution (but see Box 3.1).

### Box 2.4 Some basic formulae for summary statistics in Excel

The following are some of the formulae commonly used to calculate summary statistics in Excel. The cell numbers in parentheses are, of course, arbitrary for the purposes of illustration:

Number of items in sample	'=COUNT(A2:A51)'
Arithmetic mean	'=AVERAGE(A2:A51)'
Mode	'=MODE(A2:A51)'
Median	'=MEDIAN(A2:A51)'
Variance	'=VAR(A2:A51)'
Standard deviation	'=STDEV(A2:A51)'
Standard error	'=(STDEV(A2:A51))/(SQRT(COUNT(A2:A51)))'
Interquartile range	'=QUARTILE(A2:A51,1)' for the first quartile (the value 25% up the data set when rank ordered), and '=QUARTILE(A2:A51,3)' for the third quartile (the value 75% up the data set)

As we emphasise in Box 2.2, using statistical packages on a computer will teach you nothing about the underlying assumptions and mechanics of particular tests. As far as the user is concerned, the package simply offers a magic box of tricks that throws up results when data are put into it. While we use our selected packages as the main presentational vehicle for statistical tests in the text – because, for better or worse, packages like these are what students will be recommended to use – we also include worked examples of the main kinds of test in the Appendices, so the reader can have some idea of what they actually do.

Of course, while interesting to do in their own right, exploratory analyses are useful only to the extent that they promote further investigation to confirm that what looks interesting at an exploratory level is still interesting when data are collected more rigorously and subjected to more thorough analysis. This brings us back to hypotheses and predictions, and leads to a consideration of *confirmatory statistics*.

## 2.3 Forming hypotheses

### 2.3.1 Turning exploratory analyses into hypotheses and predictions

Exploratory analyses are generally open-ended in that they are not guided by preconceived ideas about what might be going on. However, they are the first important step on the way to formulating hypotheses that do then guide investigation. As we have seen already, hypotheses can be very general or they can be specific. Both kinds can be generated from our observational notes.

**Example 1****Plants and herbivores**

The observations of leaf damage in the samples of plant species suggest that several factors are influencing the type and amount of damage. A possibility emerging from the Notes is that damage decreases with height off the ground and thus vulnerability to slugs. This can be framed as a readily testable hypothesis:

**Hypothesis 1A** *The type and extent of leaf damage reflects availability to slugs.*

from which some predictions for testing might be:

**Prediction 1A(i)** *Taller plant species will have less leaf damage by slugs than shorter species.*

**Prediction 1A(ii)** *Leaf damage will decrease the further up a plant that samples are taken.*

Of course, this is a very broad hypothesis and many features of a leaf affecting its likelihood of predation are going to change along with its height off the ground. Some of these are suggested by other observations. For example:

**Observation** Larger leaves are often tougher or smell strongly.

**Hypothesis 1B** *The decrease in damage among larger leaves is due to their reduced palatability.*

**Prediction 1B** *For any given size of leaf, damage will decrease the tougher the cuticle or the stronger the odour on crushing.*

**Observation** Larger leaves are sometimes associated with thorns or sticky hairs on the stems.

**Hypothesis 1C** *Reduced damage among larger leaves is due to grazing deterrents on the stems.*

**Prediction 1C(i)** *The incidence of severe damage (suggestive of large herbivores) will be lower on thorny species.*

**Prediction 1C(ii)** *The incidence of less severe damage (suggestive of invertebrate herbivores) will be lower on species with sticky, hairy stems.*

**Example 2****Hosts and parasites**

The notes on the samples of material and parasites taken from voles suggest a number of interesting possibilities, some to do with the age and sex of the voles, others to do with relationships between the different parasites.

**Observation** The number of *Babesia*-infected red cells and faecal egg scores appeared to be higher in male voles than in females, and higher in adults than in juveniles.

**Hypothesis 2A** *Parasite burdens are affected by differences in the levels of reproductive hormones between age and sex classes.*

**Prediction 2A** *Parasite burdens will increase with host testosterone levels.*

**Hypothesis 2B** *Parasite burdens are affected by differences in aggressive behaviour and stress between age and sex classes.*

**Prediction 2B(i)** *Parasite burdens will be greater among dominant territorial males.*

**Prediction 2B(ii)** *Parasite burdens will increase in any individual with the amount of aggressive behaviour shown.*

**Prediction 2B(iii)** *Parasite burdens will increase with host corticosterone (stress hormone) levels.*

**Observation** Sex and age differences in *Babesia* levels appear to be associated with the number of ticks on the host.

**Hypothesis 2C** *The intensity of infection with Babesia depends on the degree of exposure to infected ticks.*

**Prediction 2C** *Babesia burden will increase with the number of infected ticks establishing on the host.*

**Hypothesis 2D** *The intensity of infection with Babesia depends on the degree of resistance to tick infection.*

**Prediction 2D** *The intensity of infection with Babesia will decrease with the host's ability to mount an antibody response to ticks.*

### Example 3

#### Nematodes and pollutants

Observations on the samples of soil-dwelling nematodes suggest a number of things vary with the pollution status of the site of origin – species diversity and fecundity among them.

**Observation** Fewer species were identified in the samples from the two polluted sites compared with the unpolluted site.

**Hypothesis 3A** *Pollution reduces species diversity.*

**Prediction 3A** *The addition of pollutants to identical multi-species nematode cultures will result in a reduction in the number of species supported over time.*

Hypothesis 3A is another very broad hypothesis and could give rise to several more specific hypotheses, each generating their own predictions. For example:

**Hypothesis 3B** *Pollutants are toxic to those species missing from polluted sites.*

**Prediction 3B** *Species present at unpolluted sites but missing from polluted sites will show greater mortality when exposed to pollutants.*

**Hypothesis 3C** *Pollutants affect resource availability for certain groups of nematodes.*

**Prediction 3C** *Species missing from polluted sites will tend to come from certain trophic or microhabitat groups.*

**Observation** Some species are present only in polluted sites.

**Hypothesis 3D(i)** *Such species benefit from relaxed interspecific competition in polluted sites.*

**Prediction 3D(i)** *Increasing the number of species in the culture in an otherwise constant and pollutant-free environment will tend to result in the loss of such species from the community.*

**Hypothesis 3D(ii)** *Pollutants create niche opportunities not available in unpolluted sites.*

**Prediction 3D(ii)** *For any given number of species in the culture in an otherwise constant environment, such species will do better when pollutant is added compared with an unpolluted control.*

**Observation** Fewer juvenile stages were recorded from the organophosphate-polluted site than from the other two sites, but there was no consistent difference in the number of females with eggs.

**Hypothesis 3E(i)** *Organophosphate pollution affects recruitment to nematode populations through reduced egg viability.*

**Prediction 3E(i)** *Females reared in organophosphate-treated, single-species culture will show reduced hatching success per egg compared with those reared in heavy metal or untreated control cultures.*

**Hypothesis 3E(ii)** *Organophosphate pollution affects recruitment to nematode populations through increased juvenile mortality.*

**Prediction 3E(ii)** *Females reared in organophosphate-treated, single-species culture will show comparable hatching success per egg but reduced survival of resultant juvenile stages than those reared in heavy metal or untreated control cultures.*

#### Example 4


#### Crickets

Male field crickets seemed to be aggressive to one another when put together in an arena. Whether or not an encounter resulted in fighting varied with the number of crickets, and individuals differed in their tendency to initiate and win fights. The apparent effects of providing egg box shelters and introducing females suggest that interactions between males are concerned ultimately with gaining access to females.

**Observation** The number of encounters leading to a fight was lower when more crickets were present.

**Hypothesis 4A** *The cost of fighting on encounter increases with population size and the chance of encountering another male.*

**Prediction 4A** *The probability of an encounter's resulting in a fight will decrease with increasing numbers of males and in the same number of males maintained at a higher density.*



**Observation** Larger males initiated more fights per encounter and won in a greater proportion of encounters.

**Hypothesis 4B** *Large size confers an advantage in fights between males.*

**Prediction 4B** *Males will be less likely to initiate a fight when their opponent is larger.*

**Observation** Interactions tended to escalate from chirping and antenna-tapping to overt fighting.

**Hypothesis 4C** *The escalating sequence reflects information-gathering regarding the size of the opponent and the likelihood of winning.*

**Prediction 4C** *Encounters will progress further when opponents are more similar in size and it is more difficult to judge which will win.*

**Observation** Larger males ended up in or near egg box shelters and females tended to spend more time with these males.

**Hypothesis 4D(i)** *Females prefer to mate with males in shelters for protection from predators.*

**Prediction 4D(i)** *Giving a male a shelter will increase the attention paid to him by females and his chances of copulating.*

**Hypothesis 4D(ii)** *Females prefer large males.*

**Prediction 4D(ii)** *Given a choice of males, all with or without shelters, females will spend more time and be more likely to copulate with larger males.*

### 2.3.2 Null hypotheses

In the examples above, we have phrased predictions in terms of the outcomes they lead us to expect. Prediction 1B, for example, leads us to expect that the amount of damage sustained by a leaf will decrease the tougher its cuticle or the more volatiles it contains. We can test this prediction by carrying out a suitable investigation and associated confirmatory analysis. Formally, however, we do not test predictions in this form. Rather, we test them in a null form that is expressed as a hypothesis *against* the prediction. This is known as a *null hypothesis* and is often expressed in shorthand as  $H_0$  (the prediction made by the study is then referred to conventionally as  $H_1$ ). Predictions are tested in the form of a null hypothesis because science proceeds conservatively, always assuming that something interesting is *not* happening unless convincing evidence suggests, for the moment, that it might be. In the case of Prediction 1B, therefore, the null hypothesis would be that tougher cuticles or more volatiles would make no difference to the amount of damage sustained by a leaf. We shall see later what burden of proof is necessary to enable us to reject the null hypothesis in any particular case.

There is a second, and from a practical point of view more crucial, point to make about the predictions. Skimming down them gives the impression of specificity and diversity; each prediction is tailored to particular organisms and circumstances, and those from one example seem to have little to do with those from others. At the trivial level of detail, this is obviously true. However, in terms of the kinds of question they reflect, predictions from the different examples in fact have a great deal in common. Before we can proceed with the problem of testing hypotheses and choosing confirmatory analyses, we need to be aware of what these common features are.

### 2.3.3 Differences and trends

Although we derived some 23 different predictions from our notes, and could have derived many more, all fall (and any others would have fallen) into one of two classes. Regardless of whether they are concerned with nematodes or crickets or with surviving pollutants or fighting rivals, they either predict some kind of *difference* or they predict some kind of *trend*. Recognising this distinction is vitally important because it determines the kind of confirmatory test we shall be looking to perform and therefore the design of our experiments. Surprisingly, however, it proves a stubborn problem for many students throughout their course, with the result that confirmatory analyses often fall at the first fence. Let's look at the distinction more closely.

A *difference* prediction is concerned with some kind of difference between two or more groups of measurements. The groups could be based on any characteristics that can be used to make a clear-cut distinction; obvious examples could be sex (e.g. a difference in body size between males (Group 1) and females (Group 2)), functional anatomy (e.g. a difference in enzyme activity between xylem (Group 1), phloem (Group 2) and parenchyme (Group 3) cells in the stem of a flowering plant), or experimental treatment (e.g. a difference in the number

of chromosomal abnormalities following exposure to a mutagen (Group 1) or exposure to a harmless control (Group 2)). Which of the predictions we derived earlier are difference predictions?

### Example 1

#### Plants and herbivores

In the leaf sample study there are two difference predictions:

- Prediction 1C(i) leads us to expect that the incidence of severe damage will be lower on thorny species (Group 1) than on non-thorny species (Group 2).
- Prediction 1C(ii) leads us to expect a reduction in less severe damage on species with sticky, hairy stems (Group 1) compared with species without (Group 2).

### Example 2

#### Hosts and parasites

One prediction from the example of host/parasite relationships involves a difference:

- Prediction 2B(i) suggests a difference in parasite burden between dominant territorial males (Group 1) and other age and sex categories of host (Group 2).

### Example 3

#### Nematodes and pollutants

Almost all the predictions arising from the soil-dwelling nematode samples turn out to be difference predictions:

- Prediction 3A suggests a difference in the number of species between cultures to which pollutant has been added (Group 1) and those that are pollutant-free (Group 2).
- Prediction 3B predicts a difference in sensitivity to pollutants between species absent from polluted sites (Group 1) and those present at such sites (Group 2).
- Prediction 3C suggests a difference between trophic groups (e.g. bacterial feeders (Group 1), fungal feeders (Group 2), plant feeders (Group 3)) in the tendency to be present at polluted sites.
- Prediction 3D(ii) leads us to expect a difference between cultures to which pollutant has been added (Group 1) and pollutant-free controls (Group 2) in the tendency to support nematode species found only in polluted sites in the field.
- Prediction 3E(i) suggests a difference in egg hatching success between female worms exposed to organophosphate pollutant (Group 1) and those not (Group 2).
- Prediction 3E(ii) is similar to the last prediction except that it suggests a difference in larval mortality instead.



## Example 4

## Crickets

Two predictions from the crickets involve differences:

- Prediction 4B suggests that males will be less likely to initiate a fight when their opponent is larger than them (Group 1) than when it is smaller (Group 2).
- Prediction 4D(i) predicts that females will pay more attention to males with a shelter (Group 1) than to males without a shelter (Group 2).

*Trend* predictions are concerned not with differences between hard and fast groupings but with the relationship between two more or less continuously distributed measures. Thus, for example, a relationship might be predicted between the amount of an anthelmintic drug administered to a rat infected with nematodes and the number of worm eggs subsequently counted in the animal's faeces. In this case, we should expect the relationship to be negative with egg counts decreasing the more drug the rat has received. On the other hand, a positive relationship might be predicted between the number of hours of sunlight received and the standing crop of a particular plant. With trends we can therefore envisage two measures as the axes of a graph. One measure extends along the bottom ( $x$ ) axis, the other up the vertical ( $y$ ) axis. Sometimes it doesn't matter which measure goes along the  $x$ -axis and which up the  $y$ -axis because there is no basis for implying cause and effect and we are interested only in whether there is some kind of association. Thus, we might expect a strong association between the amount of ice cream eaten and the amount of time spent in the sea on a visit to the seaside because both would go up with temperature. Since neither could reasonably be thought of as a cause of the other, it is of no consequence which goes on the  $x$ - or  $y$ -axis. In the two examples above, however, there are reasonable grounds for supposing cause and effect. While it is plausible for the anthelmintic drug to affect faecal egg counts, it is not plausible for the egg counts to have influenced the amount of drug. Similarly, hours of sunlight could influence a standing crop but not vice versa. In these cases, the drug dose and hours of sunlight measures should go on the  $x$ -axis and the egg counts and standing crops on the  $y$ -axis. It is important to stress, however, that by doing this we are not asserting that the  $x$ -axis measure really *is* a cause of the  $y$ -axis measure – as we shall see later, inferring cause and effect from relationships requires extreme caution – merely that if there was a cause-and-effect relationship it would most likely be that way round. This is also clear in the remainder of our example predictions, all of which involve trends.

## Example 1

## Plants and herbivores

In Prediction 1A(i), leaf damage is expected to decrease as the height of plant species increases. Plant height should thus be the  $x$  measure and leaf damage the  $y$  measure:

- Prediction 1A(ii) makes a similar prediction except that the expected relationship is *within* plants. Height up the plant is the  $x$  measure and leaf damage once again the  $y$  measure.

- Prediction 1B suggests a negative relationship between toughness of the cuticle ( $x$  measure) and leaf damage ( $y$  measure).

### Example 2

#### Hosts and parasites

All except one of the predictions from the vole parasites example are trend predictions:

- Prediction 2A predicts an increase in parasite burdens ( $y$  measure) with increasing testosterone levels ( $x$  measure).
- Predictions 2B(ii) and (iii) predict similar increases in parasite burdens ( $y$  measure) but this time as a function of increasing aggression and corticosterone levels ( $x$  measures), respectively.
- Prediction 2C suggests an increase in *Babesia* levels ( $y$  measure) with the number of ticks recovered from the host ( $x$  measure).
- Prediction 2D suggests a reduction in *Babesia* levels ( $y$  measure) with host immune responsiveness ( $x$  measure).

### Example 3

#### Nematodes and pollutants

Only one of the predictions arising from the nematode example suggests a trend:

- Prediction 3D(i) predicts a loss of species found only at polluted sites ( $y$  measure) as the number of species in a culture increases ( $x$  measure).

### Example 4

#### Crickets

Three predictions from the crickets involve trends:

- Prediction 4A first of all predicts that the number of encounters ending up in a fight ( $y$ ) will increase with the number of crickets ( $x$ ), then predicts a similar increase when the same number of crickets are maintained at higher densities (density =  $x$ ).
- Prediction 4C involves a predicted trend in the tendency to escalate an interaction ( $y$ ) with decreasing difference in size between opponents ( $x$ ).
- Prediction 4D(ii) suggests that the time females spend with a male ( $y$ ) and their tendency to copulate ( $y$ ) will increase with male size ( $x$ ).

There is thus a clear distinction between difference and trend predictions. Of course, it is possible to recast some trend predictions as difference predictions (for instance, a continuous measure of group size for use in a trend could always be recast in terms of small groups (groups below size  $w$ ) and large groups (groups above size  $w$ ) and thus be used in a difference prediction). What makes the distinction, therefore, is not the data *per se* but the way data are to be collected or classified for analysis. Thus, while measures such as group size or time intuitively suggest trends, there is nothing to stop their being used in difference


predictions. It all depends on what is being asked. This is often a source of serious confusion among students encountering open-ended data-handling for the first time.

## 2.4 Summary

1. Open-ended observation is a good way to develop the basis for forming hypotheses and predictions about material. It pays to make observations quantitative where possible so that exploratory analyses can highlight points of interest.
2. Exploratory analysis is a useful (often essential) first step in extracting interesting information from observational notes or other sources of exploratory information. It can take a wide variety of forms, such as bar charts, scattergrams, or tables of summary statistics.
3. Exploratory analyses, or raw exploratory information itself, can lead to a number of hypotheses about the material. In turn, each hypothesis can give rise to several predictions that test it. Formally, predictions are tested in the form of null hypotheses.
4. While predictions derived from hypotheses may be diverse and specific in detail to the material of interest, they fall into two clearly distinguishable categories: predictions about *differences* and predictions about *trends*. Which of these categories a prediction belongs to is determined by the way data are to be collected or classified for analysis.

### Reference

Dytham, C. (2003) *Choosing and using statistics: a biologist's guide*. 2nd edition, Blackwell, Oxford.



# 3

## Answering questions


What do the results say?

In the last chapter, we looked at the way hypotheses can be derived from exploratory information. We turn now to the problem of how to test our hypotheses. As we have seen, we begin by making predictions about what should be the case if our hypotheses are true. These predictions then dictate the experiments or observations that are required. However, this may not be as straightforward as it sounds; decisions have to be made about what is to be measured and how, and how the resulting data are to be analysed. The questions of measurement and analysis are, of course, interdependent. This is obvious both at the level of choosing between difference and trend analyses – there is little point collecting data suitable for a difference analysis if what we’re looking for is a trend – and at the choice of analyses *within* differences and trends. While at first sight it might seem like putting the cart before the horse, therefore, we shall introduce confirmatory analysis *before* dealing with the collection of data so that the important influence of choice of analysis on data collection can be made clear.

### 3.1 Confirmatory analysis

#### 3.1.1 The need for a yardstick in confirmatory analysis: statistical significance

Take a look at the scattergram in Fig. 3.1. It shows a relationship between the concentration of a fungicide sprayed on a potato crop and the percentage of leaves sampled subsequently that showed evidence of fungal infection. A plot like this was presented to a class of undergraduates. Students in the class were



**Figure 3.1** A scattergram of the relationship between concentration of a fungicide applied to a potato crop and the percentage of leaves subsequently found to be infected.

asked whether they thought it suggested any effect of fungicide concentration on infection. The following are some of their replies:

*Yes, fungicide concentration obviously has an effect because infection goes down with increasing concentration.*

*It's hard to say. It looks as though there is some effect, but it's pretty weak. More data are needed.*

*Fungicide concentration is reducing infection but there must be other things affecting it as well because there's so much scatter.*

*I don't think you can say anything from this. Yes, there is some downward trend with increasing concentration but several points for high concentrations are higher than some of those for low concentrations. Totally inconclusive.*

*Yes, there is a clear negative effect.*

Clearly there are different, subjective, reactions to the plot. To some it is unequivocal evidence for an effect of fungicide concentration; to others it doesn't suggest much at all. Left to 'eyeball' impressions, therefore, the conclusion that emerged would be highly dependent on who happened to be doing the eyeballing. What is required, quite clearly, is some independent yardstick for deciding whether or not we can conclude anything from the relationship. Since the scenario above could be repeated with any set of data – difference or trend – the need for such a yardstick arises in all cases. One or two idiosyncratic departures notwithstanding (one well-known ornithologist used to advocate the yardsticks 'not obvious', 'obvious' and 'bloody obvious'), the yardstick used conventionally in science is *statistical significance*. There is nothing magical or complicated about statistical significance. It is simply an arbitrary criterion accepted by the international scientific community as the basis for accepting or rejecting the null hypothesis in any given instance and thus deciding whether predictions, and the hypotheses from which they are derived, hold. If the criterion is reached, the difference or trend in question is said to be *significant*; if it is not, the result is *non-significant*. The term 'significant' thus has an important, formal meaning in the context of data analysis and its use in a casual, everyday sense should be avoided in discussions relating to scientific interpretation. How do we decide whether differences or trends are significant? By using the most appropriate of the vast range of significance tests at our disposal. Before we introduce some of these tests, however, we must say a little more about significance itself.

## 3.2 What is statistical significance?

A *statistic* is a measure, such as a mean or a correlation, derived from samples of data that we have collected. Our expectation is that it relates closely to an equivalent, real value (*parameter*) in the population from which the samples were drawn. Of course it might or it might not. Our sampling technique (*see later*) may have been impeccable and produced a very accurate reflection of the real world. More than likely, however, and for all kinds of forgivable reasons, it will have produced a somewhat biased sample and our calculated statistics will differ from their population parameters. This is what makes statistical inference tricky. If we detect an apparent difference between two sets of data, or an apparent trend in the relationship between them, is the difference or trend real or is it just an artefact of the chance bias in our sampling? Or, to put it another way, is it statistically significant?


The criterion that determines significance is the probability (usually denoted as  $\alpha$ ) that a difference or trend as extreme as the one observed would have occurred if the null hypothesis – that there is really *no* difference or trend in the population from which the sample came – was true. Confused? An example makes it clear. Let's take one of our earlier predictions, say Prediction 2C. This predicts an increase in *Babesia* burden in voles with increasing degree of tick infestation. Suppose we had tested this by infesting each of ten sets of five parasite-free voles with a different number of *Babesia*-infected ticks, measured

the subsequent number of infected blood cells in each vole and found what looked like a convincing positive trend: *Babesia* burden goes up with increasing numbers of ticks. The null hypothesis in this case, of course, is that *Babesia* burden will *not* increase with the number of ticks. What, then, is the probability of obtaining a positive relationship as extreme as the one we got if this null hypothesis is really true and the apparent trend a chance effect? A helpful analogy here might be the probability of obtaining the apparent trend by haphazardly throwing darts at the scattergram. An appropriate significance test will tell us (we shall see how later). By convention in biology, a probability of 5 per cent (= 0.05 when expressed as a proportion) is accepted as the threshold of significance. If the probability of obtaining a relationship as extreme as ours by chance turns out to be 5 per cent or less, we can regard the relationship as significant and reject the null hypothesis. If the probability is greater than 5 per cent we do not reject the null hypothesis and the relationship is regarded as non-significant. If the null hypothesis is not rejected, we effectively assume that our apparent relationship was due to a chance sampling effect. As a matter of interest, the negative trend in Fig. 3.1 is significant at the 5 per cent level, so the optimists have it in this case!

The 5 per cent threshold is, of course, arbitrary and still leaves us with a one-in-twenty chance of rejecting the null hypothesis incorrectly (falsely accepting there is a difference or a trend when there isn't). Under some circumstances, for instance when testing the effectiveness of a drug, a one-in-twenty risk of incorrect rejection might be considered too high. In certain areas of research such as medicine, therefore, the arbitrary threshold of significance is set at 1 per cent (= 0.01). In other disciplines it is sometimes relaxed to 10 per cent (= 0.1). Although we have talked of threshold probability ( $p$ ) values ( $p < 0.05$ ,  $p < 0.01$ , etc.), most computer statistical packages now quote *exact* probabilities ( $p = 0.0425$ ,  $p = 0.1024$ , etc.) for the outcome of significance tests. If the package doesn't tell you whether the exact probability it quotes is significant, simply apply the threshold value rule as before. Thus, on the 5 per cent criterion,  $p = 0.0425$  is significant, because it is less than 0.05, but  $p = 0.1024$  is not, because it is greater than 0.05\*.

An important point must be made here regarding the inference to be drawn from achieving different levels (10, 5, 1 per cent, etc.) of significance. A high level of significance is not the same as a large effect in the sense of a large difference or a steep trend. The magnitude of an effect – difference or trend – is usually

\* There are differing opinions about the form in which the  $p$ -value from a statistical test should be reported. Some take the view that, since most packages give the exact probability, this should be reported as the result. However, some statisticians think that these probabilities are really correct only if the assumptions of the test are absolutely fulfilled by the data, which is rarely true, and hence it is misleading to give the appearance of accuracy by citing the exact probability. For many scientists the threshold is the important thing, and they report only whether (a) it has not been equalled or crossed (in which case the outcome is reported as 'non-significant' or 'ns'), or (b) if it has been crossed, at what level of probability, either by quoting the thresholds themselves ( $p < 0.05$ ,  $p < 0.01$  or  $p < 0.001$ ) or using a conventional indicator such as asterisks ('\*', '\*\*' or '\*\*\*') to denote the same thing (*see also* p. 155). However, see the interesting paper by Ridley *et al.* (2007) which suggests scientists sometimes pursue these critical thresholds rather too zealously.



**Figure 3.2** Scattergrams showing the effect of sample size on apparent trends (see text).

known as an *effect size*\*\* . This is quite different from the level of significance. The limitations of significance in this sense are made clear in Fig. 3.2. The figure shows two trends. In Fig. 3.2a, the *y* measure increases in close relationship with the *x* measure, yielding what looks like a clear positive trend. Figure 3.2b, on the other hand, shows a scatter of points in which it is more difficult to discern a trend. However, if we perform a suitable significance test for a trend on the two sets of data, the trend in Fig. 3.2b turns out to be significant at the 1 per cent level while that in Fig. 3.2a isn't even significant at the 10 per cent level. The crucial difference between the two trends, of course, is the sample size. The number of data points in Fig. 3.2a is low, so a few inconsistencies in the trend are enough to push it below significance. Figure 3.2b, however, has a large number of points; so even though there is a wide scatter, the trend is still significant. Exactly the same sample size effect would operate in the case of difference analyses.

Because the level of significance by itself gives little indication of the magnitude of a difference or trend, it is always important to provide such an indication, usually in the form of summary statistics and their associated sample sizes. We shall return to this point later.

\*\* 'Effect size' is a name given to a family of indices that measure the magnitude of a difference or trend. Unlike significance tests, effect sizes are independent of sample size. They are usually measured in two ways: (a) as the standardized difference between means (difference analyses) or (b) as the correlation between the independent variable classification and the individual scores of the dependent variable (trend analysis) (see also Box 3.14).



## 3.3 Significance tests

So far, we have seen how to get round the problem of subjective impression in interpreting data by using the criterion of statistical significance, and we have looked at some caveats on the interpretation of significance. We can now turn to the statistical *tests* that enable us to decide significance.

Right at the beginning we said that this book was not about statistics. It isn't. At the same time, statistical significance tests are an essential tool in scientific analysis and the rules for using them must be clearly understood. This does not necessarily require a knowledge of statistical theory and the mathematical mechanics of tests any more than using a computer program requires an appreciation of electronics and microcircuitry. As with most tools, it is competent *use* that counts rather than theoretical understanding. But while our aim here is simply to introduce the use of significance tests as a basic tool of enquiry, we stress again (*see* Box 2.2) that acquaintance with statistical theory is strongly recommended, and it is envisaged that many users of this book will also be pursuing courses in statistics and have at their disposal some of the many introductory and higher-level textbooks now available (e.g. Siegel and Castellan 1988; Sokal and Rohlf 1995; Grafen and Hails 2002; Dytham 2003; Hawkins 2005; Ruxton and Colegrave 2006).

### 3.3.1 Types of measurement and types of test

The test we choose in a particular case may depend on a number of things. The following are three important ones.

#### 1 Types of measurement

The first is the kind of measurement we employ. Without getting too bogged down in jargon, we can recognise three kinds.

**Nominal or classificatory measurement.** Here, observations or recordings are allocated to one of a number of discrete, mutually exclusive categories such as male/female, mature/immature, red/yellow/green/blue, etc. Thus if we were to watch chicks pecking at red (R), green (G) and orange (O) grains of rice and recorded the sequence of pecks with respect to the colour targeted, we might end up with a string of data as follows:

**R O R R G G O R G G G G O O G R O**

Such data are measured purely at the level of the category to which they belong and measurement is thus *nominal* or *classificatory*.

**Ordinal or ranking measurement.** In some cases, it may be desirable (or necessary) to make measurements that can be *ranked* along some kind of scale. For instance, the intensity of the colour of a turkey's wattles might be used as a guide to its state of health. The degree of redness of the wattles of different birds could be scored on a scale of 1 (pale pink) to 10 (deep red). The allocation of scores to

wattles is arbitrary and there is no reason to suppose that the degree of redness increases by the same amount with each increase in score. Thus the difference in redness between scores of 8 and 9 might be greater than the difference between scores of 2 and 3. All that matters is that 9 is redder than 8 and 3 is redder than 2; the absolute difference between them cannot be quantified meaningfully.

**Constant interval measurements.** In other scale measurements, the difference between scores can be quantified so that the difference between scores of 2 and 3 is the same as that between scores of 8 and 9. Such measurements may have arbitrarily set (e.g. scales of temperature) or true (e.g. scales of time, weight, length) zero points. Such *constant interval measurements* can in fact be split into two categories on the basis of arbitrary versus true zero points and their scaling properties (e.g. Martin and Bateson, 1993), but this is not important here.

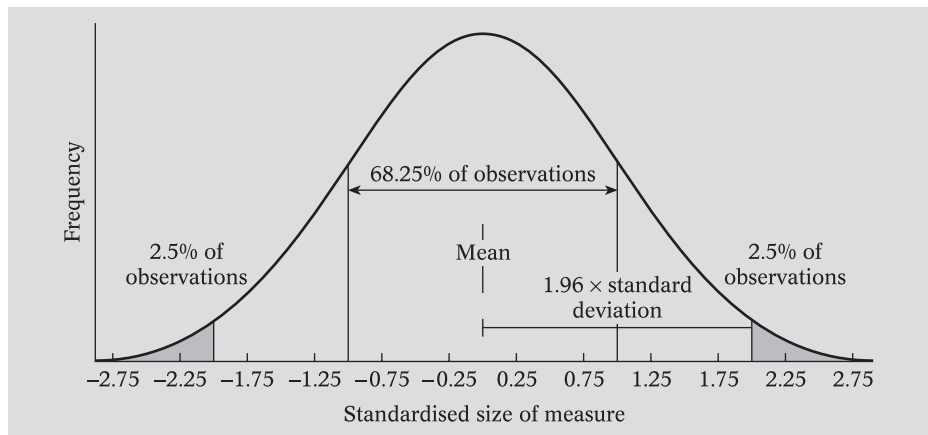
While defining measurements seems rather dry and theoretical, we need to be aware of the kind of measurement we use because, as we shall see, some significance tests are very restrictive about the form of data they can accept. Another reason for highlighting it, is that we should always seek the measures that give us the maximum amount of information for the cost required (usually in time) to obtain them. Usually this means constant interval measurements, because they are on a continuous, non-arbitrary scale. However, nominal or ordinal data are sometimes more appropriate and not infrequently the best that can be achieved.

## 2 Parametric and non-parametric significance tests

The second thing we must keep an eye on is the nature of the data set itself, in particular the sample size and the distribution of values within the sample. Again, detailed consideration of this is unnecessary but it is a factor that determines the range of tests we shall be introducing so a brief discussion is warranted.

**Parametric tests.** These make a number of important assumptions that are frequently violated by the kinds of data sets collected during practical exercises. The most critical concerns the distribution of values within samples. Parametric tests generally assume that the data conform (reasonably closely at least) to what is known as a *normal* distribution.

As Dytham (2003) puts it, the normal distribution is the most important distribution in statistics (*but see* Box 3.1a) and it is often assumed (all too frequently without checking) that data are distributed in this way. We've already encountered it in our discussion of frequency distributions and it is illustrated again in Fig. 3.3. Essentially a normal distribution demands that most of the data values fall in the middle of the range (cluster about the mean) with the number tapering off symmetrically either side of the mean to a few extreme values in each of the two tails. The height of the adult male or female population in a city would look something like this: most people would be around the average height for their sex, some would be quite tall or quite short and a few would be extremely tall or extremely short. While normality is not the only assumption underlying parametric tests, the arithmetic of such tests is based on the parameters describing this symmetrical, bell-shaped curve (hence the term *parametric*). Therefore, if you are testing for group differences, parametric tests assume that each group



**Figure 3.3** A standardised normal distribution (with a mean of 0) and showing one standard deviation (vertical lines) either side of the mean (modified from Dytham, 2003).

has the same normal distribution of values around its mean; if for trends, that there is a normal distribution of Y-values at each X-value (for regression), or a normal distribution along both X and Y axes (for correlation). The more distorted (less normal) the distribution becomes, therefore, the less meaning the calculations of parametric tests have. We can convince ourselves of this by briefly considering some basic features of the distribution. We shall then consider what to do if the data are not normal.

*The standard deviation and probability.* We've talked about the arithmetic mean as a measure of central tendency. This is a useful parameter that tells us one thing about the nature of the data set. What it doesn't tell us, of course, is anything about the *variation* in the data. This is where we come back to the *standard deviation*, first encountered in calculating the standard error to the mean in Box 2.1. The standard deviation (*see* calculation in Box 2.1) is a measure of the spread of data values about the mean, but instead of simply being the full range of *actual* values from smallest to largest (which will vary with sample size), it reflects the theoretical spread of the majority of values (68.25 per cent of them in a perfectly normal distribution) in the true population. We can represent this as two vertical lines (one either side of the mean) in Fig. 3.3, which is actually a standardised normal distribution (where the mean is subtracted from each data value and the result divided by the standard deviation, thus giving a mean of zero and a variance of 1). An easy way to visualise the standard deviation is as the point of inflection either side of the mean (where the curve of the normal distribution changes from convex to concave). Taking the majority spread like this avoids the distortion that might be inflicted by odd outlier values at the extremes and provides a convenient standard yardstick of variability within data sets. Moving two or three standard deviations away from the mean includes predictably greater percentages of the data set: 94.45 per cent in the case of two standard deviations and 99.97 per cent in the case of three. The important point here is that we can use this property when it comes to significance testing.

Since we know what percentage of the data points is included within different multiples of the standard deviation, we can easily work out how many standard deviations would include 95 per cent of the data or, more to the point, exclude the 5 per cent at the extremes (2.5 per cent at each extreme). The answer is 1.96, represented by the horizontal bar in Fig. 3.3. Thus 95 per cent of data values lie within the range ‘mean  $\pm$  1.96 standard deviations ( $\bar{x} \pm 1.96$  s.d.)’, which means that the probability of encountering a value larger or smaller than this range is 5 per cent or less – the conventional threshold of statistical significance! If we found a value with a probability of occurrence as low as this we’d be justified in concluding that it was unlikely to have come from the population that generated the curve. Of course, as the fact that the calculation of the standard deviation in Box 2.1 is based on squared deviations makes clear, all this works *only* as long as the curve is symmetrical and conforms respectably to a normal distribution. Since the calculation of parametric significance tests generally employs the same squared deviations procedure, as we shall see shortly, the restrictions apply to all these tests.

One additional point. Quite a number of significance tests are based on calculating the deviation of an observed mean from the null expectation of a standardised normal distribution (Fig. 3.3). Conventionally this is called a *z*-value and, from the above, we can see that if *z* exceeds 1.96 the result is significant at the 5 per cent level. (The value 1.96 is actually the threshold value when both sides – ‘tails’ in the jargon – of the distribution are taken into account, a so-called two-tailed test; when only one tail is considered (a one-tailed test), the threshold value is 1.64. We shall discuss one- and two-tailed tests in detail later.)

*Departures from normality.* Frequency distributions can depart from normality in a number of ways. Two broad kinds of departure, however, are *skewness* and *kurtosis*. Skewness is a synonym for asymmetry, i.e. one or other tail of the distribution is more drawn out than the other. The distribution in Fig. 2.8b is said to be skewed to the right (towards the *y*-axis), while the opposite bias would be skewed to the left. Kurtosis refers to the flatness of the distribution, which can be *leptokurtic* (more values are concentrated around the mean and in the tails and fewer in the ‘shoulders’ of the distribution), or *platykurtic* (where the reverse is true). Bimodal distributions, like that in Fig. 2.8c, are thus extremely platykurtic.

Percentages and proportions present their own problems for normality. Because they range between 0 and 100, or 0 and 1, the distribution of values is artificially truncated at either end. This may not present a serious problem if most of the values in the data set occur in the middle two-thirds or so of the distribution, but if they approach 0 or 100/1 there is cause for concern.

*Transformations.* So what do we do if we suspect our data may not be normally distributed? Happily, and as long as our sample size is big enough (> 50 as a rough guide) to make a comparison meaningful, the wide range of statistical packages now available for personal computers makes the answer simple. Test it! There are some well-established significance tests (such as the chi-squared and Kolmogorov–Smirnov tests) that allow comparisons between frequency distributions of data and various theoretical distributions of which the normal is the commonest. But there is a catch, and Box 3.1b explains this in showing how to

test for normality using SPSS. If we test our distribution and find it does not differ significantly from normal, then we're at liberty to use any appropriate parametric test at our disposal. If it does differ, we can do one of two things. We can abandon the idea of using parametric statistics and choose an appropriate *non-parametric* test instead (*see below*), or we can *transform* the data to see whether it can be normalised. Several transformations are available but the most widely used are probably *logarithmic* ( $\log x$  or, where there are zeros in the untransformed data,  $\log(x + 1)$ ) or *square root* transformations. Simply log or take the square root of each data value, recast the distribution and test it for normality again. Where percentages or proportions stray below about 30 per cent (or 0.3) or above 70 per cent (or 0.7), an *arcsine square root* transformation (calculated by taking the square root of the proportion – so divide percentages by 100 first – then the inverse sine ( $\sin^{-1}$  on many calculators) of the result) will stretch out the truncated tails and prevent undue violation of the normality assumption of parametric tests or *square root* transformations (Box 3.1b shows how to transform data in SPSS).

Of course, even transformation may not succeed in normalising our data, in which case we must seriously consider using non-parametric statistics. Indeed, we may not even get as far as worrying about the normality of our data before opting for a non-parametric approach. Among their various other requirements, parametric tests demand that measurements are of the constant interval kind, so cannot usually deal with the other types of measurement we might be forced to use. Non-parametric tests are much less restrictive here.

***Non-parametric tests.*** Non-parametric tests are sometimes referred to as distribution-free, *ranked* or *ranking* tests because they do not rely on data being distributed normally, and generally work on the ranks of the data values rather than the data values themselves. While they may be distribution-free, however, they are not entirely assumption-free. They assume the data have *some* basic properties, such as independence of measurement (*see later*) and a degree of underlying continuity (*see* Martin and Bateson, 1993): crucially, they also make the same assumption of equal dispersion among groups as do parametric tests (which assume equal variances among groups). In most cases, however, these assumptions are easily met. In the jargon of statisticians, non-parametric tests are thus more *robust* because they are capable of dealing with a much wider range of data sets than their parametric equivalents. While they can deal with the same constant interval measurements as parametric tests, they can also cope with ordinal (ranking) and classificatory measurements. Non-parametric tests are especially useful when sample sizes are small and assumptions about underlying normality particularly troublesome. There are a couple of drawbacks, however. The first, arguably overstated (*see* Martin and Bateson, 1993), weakness is that non-parametric tests are generally slightly less powerful (*power* here meaning the probability of properly rejecting the null hypothesis – we shall return to this shortly) than their parametric equivalents. The second, which is slowly being addressed (*see for example* Meddis, 1984), is that the range of tests for more complex analyses involving several variables at the same time is very limited.

**Box 3.1a** Types of distribution

The normal (or Gaussian) distribution is a form of *continuous* distribution, where, in principle, data can take a continuum of values. There are other forms of continuous distribution, but the normal is by far the most relevant to biological situations. However, there is another family of distributions which is important in biology. These are *discontinuous* or *discrete* distributions: two are particularly relevant.

**The Poisson distribution**

The Poisson distribution describes occurrence in units of time or space, for instance the number of solitary bee burrows scored in a quadrat or the number of hedgehogs dropping into a cattlegrid overnight. The key assumptions are: (a) the mean number of occurrences per unit is small in relation to the maximum number possible – i.e. occurrences are rare, (b) occurrences are independent of each other – i.e. there is no influence of one on the likelihood of another, (c) occurrences are random. Indeed, the reason a Poisson distribution is fitted to data is to test for independence or randomness in time or space. Poisson data are characterised by the mean being equal to the square of the standard deviation (known as the *variance*). If the variance is bigger than the mean, the data are more clustered than random; if it is smaller, they tend towards uniformity. A glance at these two summary statistics thus gives a good indication of the kind of distribution we're dealing with. When analysing count data with small numbers close to zero, the Poisson distribution also makes the crucial assumption that no values can be *below* zero. It can be very misleading to use statistical tests based on the normal distribution here, because this assumes the data can take any value, including negative ones. It is only relatively recently that statistical methods have been available that can make the assumption

that the data are Poisson-distributed – packages such as S-Plus, R, SAS, Genstat and GLIM are able to do this.

**The binomial distribution**

The binomial distribution is a discrete distribution of occurrences where there are two possible outcomes for each so that the probability of one outcome determines the probability of the other. Thus if the probability of an egg's hatching is 0.75, the probability that it won't hatch is  $1 - 0.75$ , i.e. 0.25. This logic can be extended to calculate the probability of various combinations of hatching (H) and failed (F) eggs in clutches of different size. Thus if two eggs are laid, there are four possible outcomes: HH, HF, FH, FF. Applying the values above gives a probability of 0.56 ( $0.75 \times 0.75$ ) for both eggs hatching successfully, 0.19 ( $0.75 \times 0.25$ ) for each of the two mixed outcomes (HF and FH) and 0.06 for two failures. The same could be done for clutches of three, four or however many eggs. The binomial distribution thus gives a baseline chance probability against which outcomes in various situations can be judged. So, if we found a population in which the proportion of two-egg clutches failing completely was 30 per cent instead of 6 per cent, we might become suspicious about the health of the birds or their environment. As with the Poisson distribution, it makes little sense to analyse probabilities assuming an underlying normal distribution if the data are distributed in a binomial fashion, since values cannot go below zero or above 1. Logistic regression is the technique that (with some restrictions) allows the assumption of the binomial distribution, and is available in many packages, including SPSS. It can also be incorporated easily into other statistical packages such as S-Plus, R, SAS, Genstat and GLIM.

Sophisticated multivariate analysis is still the undisputed province of parametric statistics. Nevertheless, for our purposes, and with a few exceptions, there are perfectly good parametric and non-parametric equivalents and we shall introduce them both in our discussion of significance testing.

**Box 3.1b** Testing whether data conform to a normal distribution

Given the importance of the normal distribution to parametric statistical analyses, you might be forgiven for imagining that testing whether your data conform to it would be made easy. Actually it isn't straightforward or automatic in any package: you have to understand what it is you are trying to test, and hence how to go about it.

The key point to understand is that all statistical tests involve fitting a *model*. If you are testing for differences, this model usually involves the fitting of mean values to each group, and testing whether these mean values are different among the groups. If you are testing for a trend, then the model usually involves fitting a straight line to the data, and then testing whether the slope of the line is different from zero.

The assumption of normality in these statistical tests is about the *residuals*, *after* you have fitted the model (*see* discussion in Sokal and Rohlf, 1995). A residual is the difference between an individual measurement and the mean value for its group (for analyses of differences between groups), or from the value predicted by a regression line (for analyses of trends). Thus usually you are not testing whether the raw data are normally distributed, but whether the residuals are\*. The process therefore usually involves (a) fitting the model, (b) saving the residuals, and (c) testing whether the residuals are normally distributed.

Once we have the residuals, a variety of tests can be used to detect departures from a normal distribution. Those used most commonly are the one-sample Kolmogorov-Smirnov test (for large samples, where  $n > 2000$ ), chi-squared, and, for

\* We can understand intuitively why testing the raw data for normality would be a mistake. Suppose, for example, that the raw data reflect two widely separated groups of values, perhaps body height from males and females where males are much bigger than females. Since the two groups differ significantly in mean value, the distribution of the overall raw data will be bimodal, i.e. having two peaks of value with a trough in between. This will fail any test for normality, even though the assumption of normality is almost certainly correct (height is the classic case of the normal distribution).

small- to medium-sized samples where  $n$  lies between 3 and 2000, the Shapiro-Wilk test. The 'Explore' option under 'Descriptive Statistics' in SPSS offers the first and third of these.

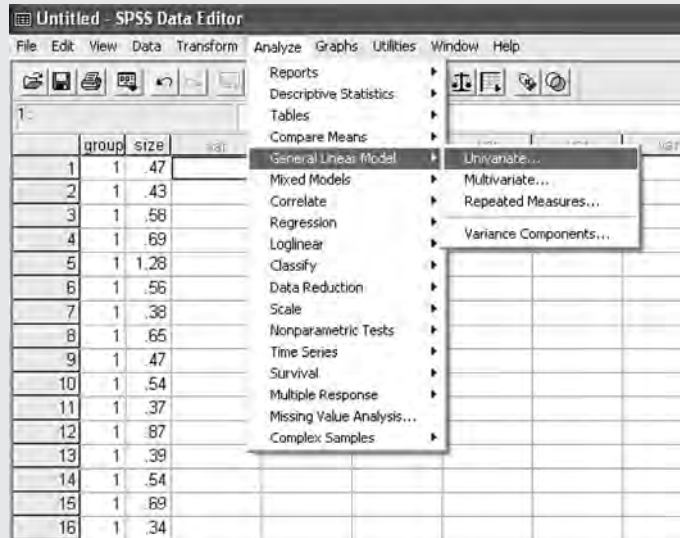
**Testing for normality as part of a test for differences**

Assume that we are interested in whether the mean sizes of seven groups of grasshoppers are different, and that we have 100 measurements for each group. One of the assumptions of the test is that the residuals, the differences of each value from the mean for its group, are normally distributed. Another assumption is that the residuals of each group all have the same normal distribution (i.e. the variances of the groups all have the same value –referred to as the 'homogeneity of variances' assumption): we shall show how to test for this assumption, along with the details of exactly how to carry out tests of difference, later on (*see* Box 3.3a on p. 73). For now we are concerned simply with testing for normality.

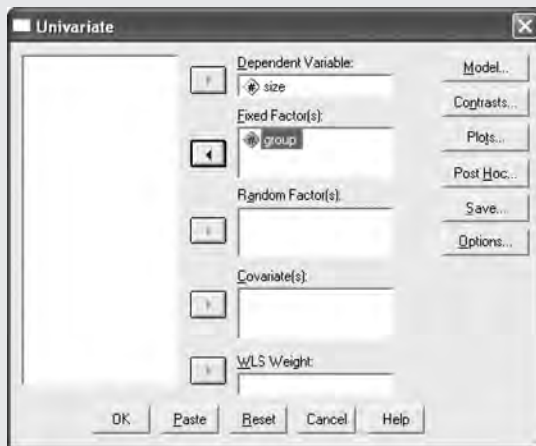
In SPSS, the data ('*size*') are in a single column, and a second column ('*group*') indexes the group to which each value belongs (Fig. (i)a). Click 'Analyse' and then 'General Linear Model' and 'Univariate'. This will bring up the dialogue box of Fig. (i)b, where we enter the variable being analysed ('*size*') into the 'Dependent Variable' box, and the index variable for the groups ('*group*') into the 'Fixed Factor(s)' box. So far this is a standard test for a difference among groups, which we shall meet in detail in Box 3.3a.

Now, in order to save the residuals, click the 'Save' button, bringing up the dialogue box of Fig. (i)c. Here, tick 'Unstandardized' in the box headed 'Residuals', and then click 'Continue'. Standardised residuals merely multiply each value of your data by a constant so that the standard deviation is 1.00, but it does not change them relative to one another, and so the shape of the distribution is unchanged.

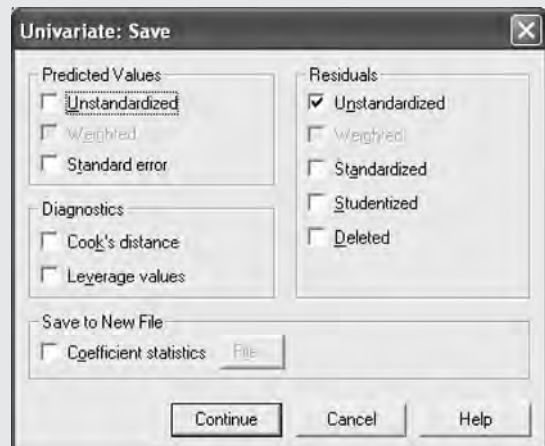
When we click 'OK' on the 'Univariate' dialogue, you will see that a new variable will appear in the data spreadsheet, called RES\_1 [Residual



(a)



(b)



(c)

**Figure (i)** How to obtain residuals from a difference test before testing them for normality. (a) Data entry, (b) selecting the routine and variables, (c) choosing the residuals to save.

for size]: these are the residuals. An Output window will open, but for the moment we are not interested in what it contains.

Now to test the residuals for normality. Click on 'Graph' and 'Histogram' (Fig. (ii)a), and in the resulting dialogue box (Fig. (ii)b), enter the residuals variable RES\_1 into the 'Variable' box, and make sure that the 'Display normal curve' box

is ticked. Clicking 'OK' produces a histogram of the distribution of the residuals, together with a fitted normal distribution (Fig. (ii)c). Notice how, relative to the normal distribution, there are too few very small values, too many just below the mean, too few above the mean, and a long tail out to the right. This looks questionable as a normal distribution.

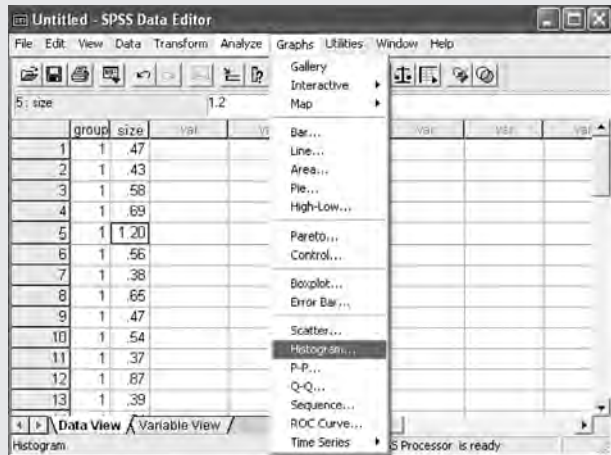


To test statistically whether it is, click on 'Analyze', 'Descriptive Statistics' and 'Explore'. Enter the residuals into the 'Dependent Variable' box of the resulting dialogue (Fig. (ii)d). Choose 'Display Plots', and then click on the 'Plots' button. From the dialogue (Fig. (ii)e) uncheck the 'Stem-and-leaf' descriptives, and ask for 'None' of the boxplots, but make sure that the 'Normality plots with tests' is checked. Pressing 'Continue' and then 'OK' brings the Output window (Fig. (ii)f) with a couple of tables and a couple of graphs. The first table (not shown) just summarises the available data, but the second contains two statistical tests, a Kolmogorov-Smirnov test and a Shapiro-Wilk test. Even with large sample sizes, the Kolmogorov-Smirnov test is very conservative, and tends to reject the normal distribution too often: thus the Shapiro-Wilk test is the preferred method. Here, both tests show that our data differ significantly from the normal distribution, because the  $p$ -values (the 'Sig.' column in the table) are less than 0.05. The two graphs show the same thing in two different ways: the clearest one is the second (Fig. (ii)f), which shows departure from normality as the degree of deviation from a horizontal line at zero. In the figure here it is obvious that the data for size differ most from the normal distribution at the small and large values, as we suspected.

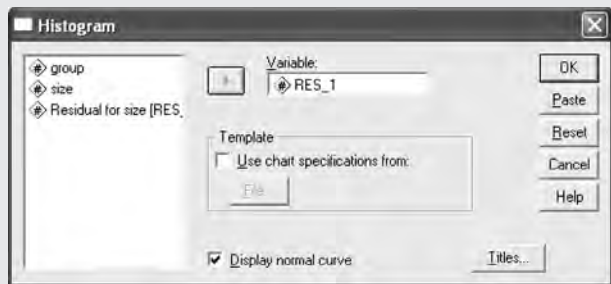
### Testing for normality as part of a test for a trend

Assume we have applied five concentrations of enzyme to a substrate with 10 replicates of each, and we have measured the amount of product produced. We are testing whether changing the enzyme concentration causes linear changes in the amount of product. The assumption of the test is that the residuals of the  $Y$ -variable (amount of product) from the regression line are normally distributed.

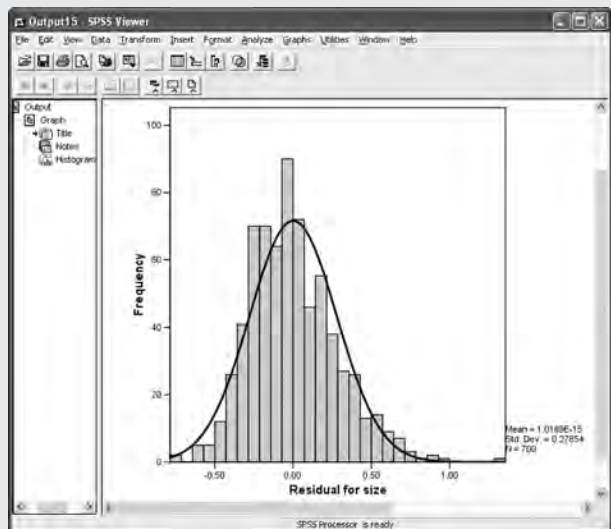
In SPSS, the enzyme concentration data go in one column (*conc*), and the amount



(a)



(b)



(c)

Figure (ii) Testing residuals for normality. (a) The menu option, (b) entering the residuals variable to obtain a histogram of the distribution together with a normal curve, (c) the resulting histogram, (d) entering the residuals variable into Explore, (e) choosing normality tests, (f) results of the normality testing.

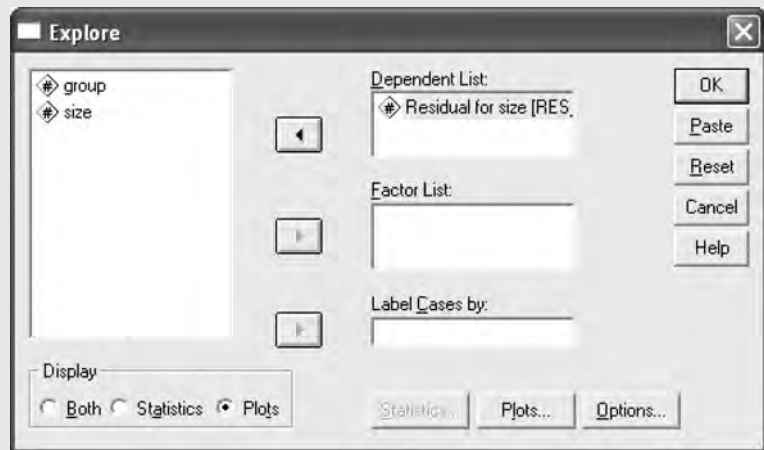
of product generated (*product*) in another (Fig. (iii)a). Click on ‘Analyze’, ‘Regression’ and ‘Linear’, and enter ‘*product*’ into the ‘Dependent’ box and ‘*conc*’ into the ‘Independent(s)’ box of the resulting dialogue box (Fig. (iii)b). Click ‘Save’, and check the ‘Unstandardized’ box (Fig. (iii)c), as before. Clicking ‘Continue’ and ‘OK’ will again result in the appearance of an extra column in the dataset, labelled ‘RES\_1 [unstandardized residual]’.

Now follow exactly the same procedure as for tests of difference above, to obtain the histogram plot and the test for normality (Fig. (iii)d). In this case we have fewer data ( $n = 50$ ) and so it is harder to tell whether it really differs from the normal distribution or not. Despite some indications of systematic departures from normality at the ends, the pattern is not suspicious enough to be able to reject the normal distribution by the Shapiro-Wilk test (Fig. (iii)d). Thus we can assume these data do conform to the assumption of the test.

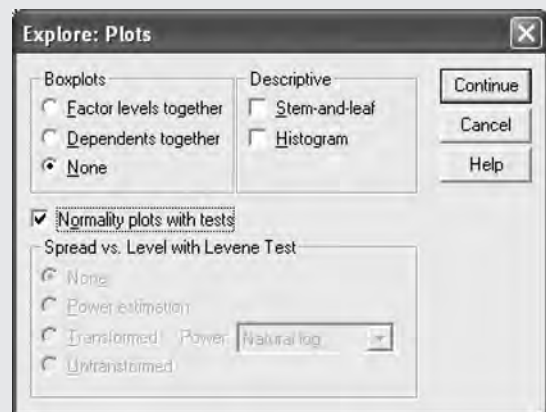
### Normalising data by transformation

Where data do not conform to a normal distribution, we can, of course, try some transformations (*see text*) to see whether we can normalise them. We can try this for the variable ‘*size*’, which showed departure from normality in the tests for a difference above. To transform ‘*size*’, click on ‘Transform’ at the top of the screen, and select ‘Compute’. Scroll down the ‘Functions’ box (Fig. (iv)a) to select a suitable transformation operator – we’ll try a  $\log_{10}$  (‘lg10(numexpr)’ in the operator list) and a square root (‘sqrt(numexpr)’ transformation as two widely used ones. Double clicking on the operator of choice puts it into the ‘Numeric Expression’ box at the top, with a highlighted question mark in the brackets. With the question mark highlighted, double click on the variable desired (here ‘*size*’) and it will transfer

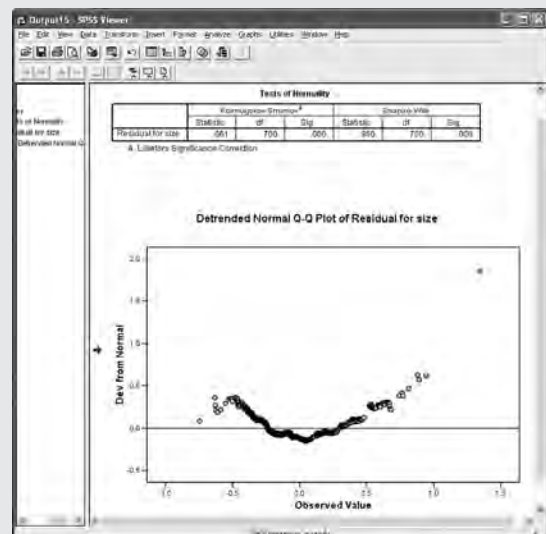
Figure (ii) cont’d



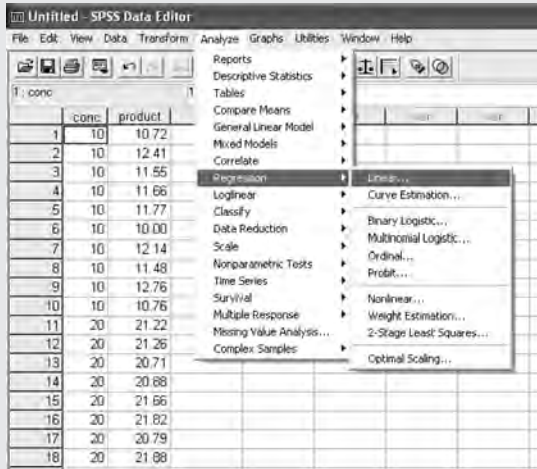
(d)



(e)



(f)



Unfiled - SPSS Data Editor

File Edit View Data Transform Analyze Graphs Utilities Window Help

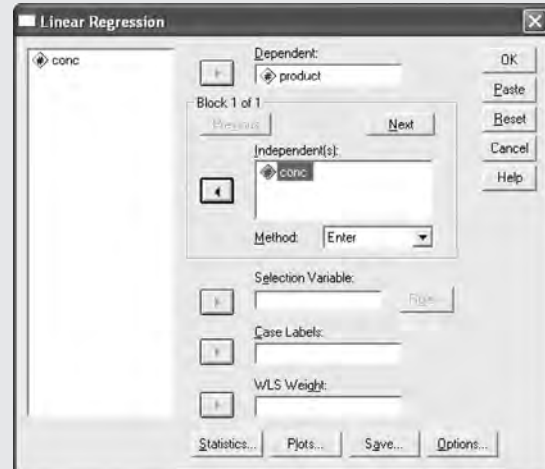
Reports  
Descriptive Statistics  
Tables  
Compare Means  
General Linear Model  
Mixed Models  
Correlate  
Regression  
Loglinear  
Classify  
Data Reduction  
Scale  
Nonparametric Tests  
Time Series  
Survival  
Multiple Response  
Missing Value Analysis...  
Complex Samples

1. conc

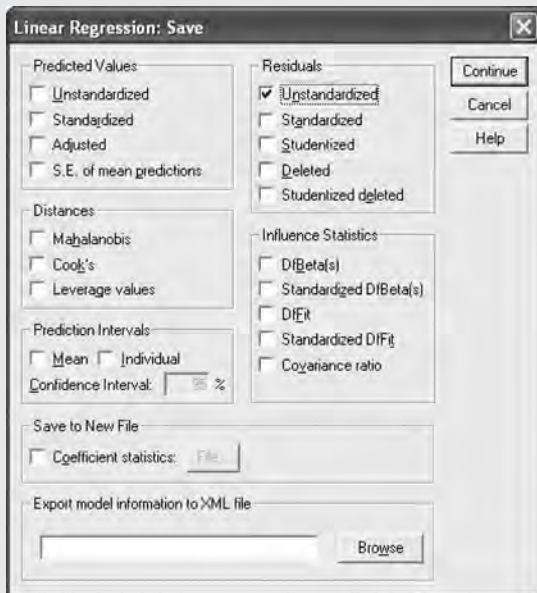
	conc	product
1	10	10.72
2	10	12.41
3	10	11.55
4	10	11.66
5	10	11.77
6	10	10.00
7	10	12.14
8	10	11.48
9	10	12.76
10	10	10.76
11	20	21.22
12	20	21.26
13	20	20.71
14	20	20.88
15	20	21.66
16	20	21.82
17	20	20.79
18	20	21.88

Regression  
Linear...  
Curve Estimation...  
Binary Logistic...  
Multinomial Logistic...  
Ordinal...  
Probit...  
Nonlinear...  
Weight Estimation...  
2-Stage Least Squares...  
Optimal Scaling...

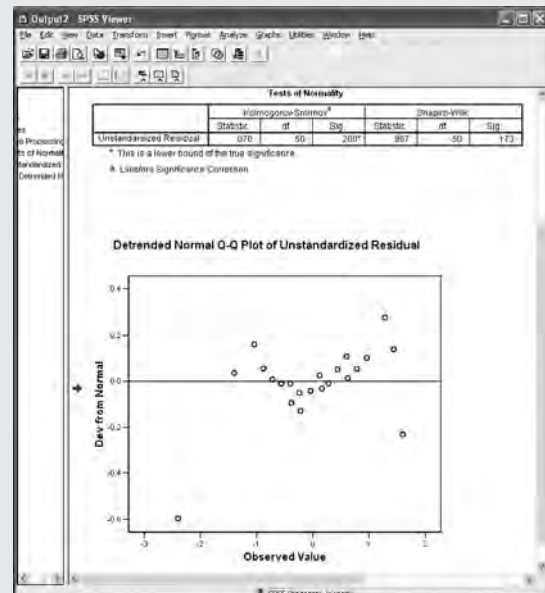
(a)



(b)



(c)

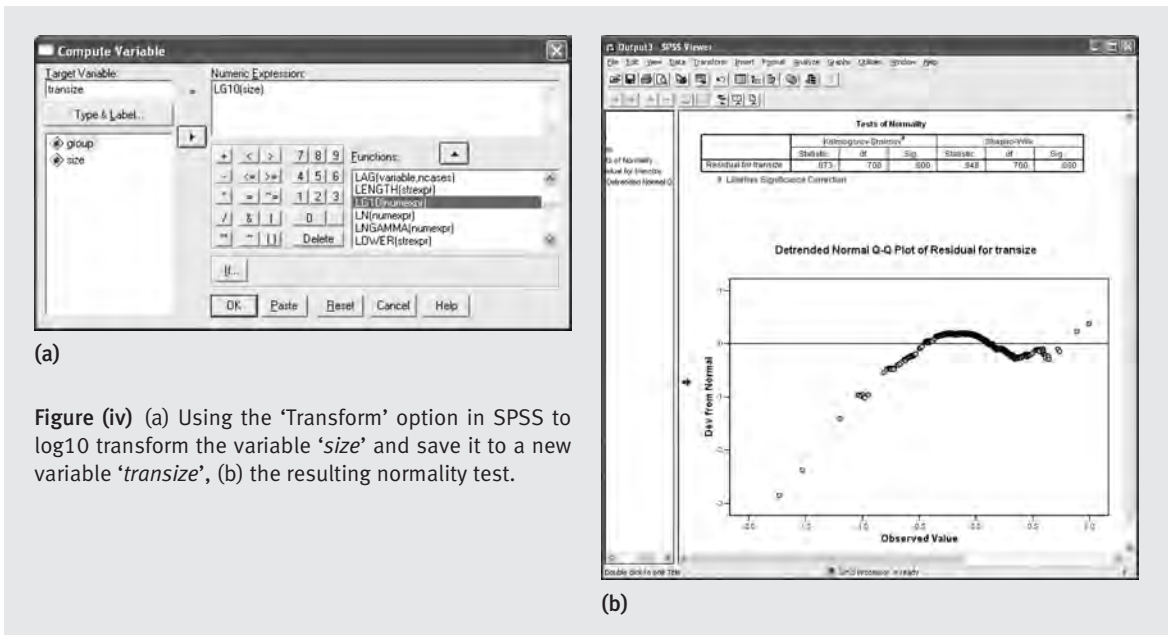


(d)

Figure (iii) Testing for normality for a trend. (a) Data entry, (b) entering variables to obtain the residuals, (c) choosing the residuals to save, (d) results of the normality testing.


into the brackets. We now have to choose a new variable name (here 'transize') to accept the transformed variable, so enter a suitable name in the 'Target Variable' box at top left (Fig. (iv)a) and click 'OK' to return to the data screen where the new variable will appear in the last column.

We now simply repeat the procedure for a test of difference above, but for the transformed variables rather than the raw ones. If we do this (Fig. (iv)b), it turns out that neither  $\log_{10}$  nor square root transformations help in the case of 'size'; both still show a significant departure from normality.



### 3 One-tailed versus two-tailed, and general versus specific tests

The third important factor we must consider relates to the prediction we are trying to test. Suppose we are predicting a difference between two sets of data, say a difference in the rate of growth of a bacterial culture on agar medium



containing two different nutrients. We could make two kinds of prediction. On the one hand, we could predict a difference without implying anything about which culture should grow faster. In this case, we wouldn't care whether culture A grew faster than culture B or vice versa. This is a *general* prediction. On the other hand, we might predict that one particular culture would grow faster than the other, e.g. A would grow faster than B; this is a *specific* prediction. Which of these kinds of prediction we make affects the way we test the predictions.

The same distinction arises with trend predictions. Imagine we want to know whether there is a trend between the size of a male cricket and the number of fights he wins over the course of a day. We can make a general prediction (there will be a trend, positive or negative), or we can make a specific prediction (larger males will win more fights; i.e. the trend will be positive). We can think of the general prediction as incorporating both positive and negative trends; either would be interesting. The specific prediction is concerned with only one of these.

In cases like those above, where there are only two possible specific predictions within the general one, we can use the same significance test for either general or specific predictions but with different threshold probability levels for the test statistic (*see below*) to be significant. Because here the specific and general predictions are concerned with one and two directions of effect respectively, the threshold value of the test statistic at the 5 per cent level in the general test becomes the threshold value at the 10 per cent level in the specific test. In statisticians' jargon, we thus do either a one-tailed (specific) or a two-tailed (general) version of the same test. There is, of course, an obvious, and dangerous, trap for the unwary here. The trap is this: if the value of a test statistic just fails to meet the 5 per cent threshold in a two-tailed test, there is a sore temptation to pretend the test is really one-tailed so that the test statistic value appropriate to the 10 per cent probability level in a two-tailed test can be used, thus increasing the likelihood of achieving a significant result. *It must be stressed that this is tantamount to cheating. A one-tailed test is legitimate **only** when the prediction is made in advance of obtaining the result **and** when results in the opposite direction can reasonably be regarded as equivalent to no difference or trend at all\*. It is completely inadmissible as a fallback when a two-tailed test fails to yield a significant outcome.* A one-tailed test should thus be used only when there are genuine reasons for predicting the direction of a difference or trend *in advance*.

However, the distinction between general and specific is not simply that between two-tailed and one-tailed tests. The latter distinction is normally used for two-group tests such as *t*-tests, where there are clearly only two possibilities for a one-tailed test ( $A > B$  or  $B > A$ ) that together make up the two-tailed case. But the world of specific testing is much richer than this!

There are in fact two different ways of making specific hypotheses, and the distinction is important. You can either predict the *rank order* of the means

\* If you predict from theory that A will have a greater mean value than B (i.e.  $H_1$  is  $A > B$ ), you are also assuming that both of the alternative results ( $A = B$  and  $A < B$ ) are *equivalent* and *together* form the null hypothesis,  $H_0$ . Thus if you find that, contrary to your prediction, the mean value of B is much greater than that of A, and would have been significant had you framed your hypothesis as a general one ( $H_1$  is  $A \neq B$ ), you are not allowed to conclude *anything* other than that the null hypothesis has *not* been rejected.

(values or ranks) of *all* the groups in your experiment, or you can pick out *particular* sets of groups to make a *contrast*. The testing of these two kinds of specific hypothesis differs.

Imagine, for instance, that we are predicting a difference between three (or more) groups – say the weight of fruits produced in a season from trees given three different maintenance treatments: A, B and C. The general prediction is that there will be differences between the three treatments, i.e. that  $A \neq B \neq C$ . Unlike in the two-group case, this general prediction is ‘made up of’ *six* potential specific predictions about the *rank order* of the treatment means:

$$A > B > C, B > C > A, C > A > B, A > C > B, B > A > C \text{ and } C > B > A$$

As long as we make the specific prediction *before* we collect the data, we can test any *one* of these in our experiment. Thus, as generally understood, the one-tailed/two-tailed distinction is only a special, and limited, case of the difference between general and specific tests.

There is an alternative way of making specific predictions that uses *contrasts*. These are ways of dividing up (‘decomposing’ in statistical parlance) the differences between groups into particular patterns. Each contrast consists of one subset of the groups ( $\alpha$ ) contrasted against a different subset of groups ( $\beta$ ), effectively creating *two* groups out of the data and comparing them using a two-group test such as a *t*-test (*see below*). A contrast can be general ( $\alpha \neq \beta$ ) or specific (e.g.  $\alpha > \beta$ ), but must be made a priori, i.e. in advance of collecting the data. However, each contrast should be independent of the others, so the general rule is a maximum of one fewer contrasts than the number of groups that could potentially be compared. In the example of the effect of three treatments on fruit production, for example, we could make two independent contrasts (because there are three groups): one of A versus B + C, and the other of B versus C. We shall make use of this approach later (*see*, for example, Boxes 3.9b,d).

### 3.3.2 Simple significance tests for differences and trends

Having discussed the general principle of statistical significance, we come now to some actual tests that allow us to see whether differences or trends are significant at an appropriate level of probability. A glance at any comprehensive statistics textbook will reveal a plethora of significance tests for both kinds of analysis. These cater for the various subtleties of assumption and requirement for statistical power under different circumstances. Many of these tests, however, are sophistications of more basic tests that are suitable for a wide range of analyses. Here, we introduce a selection of such tests that can be used with most kinds of data that are likely to be collected in practical exercises. Where appropriate, and the package performs the test(s) in question, parametric and non-parametric equivalents are presented in the on-screen format of SPSS. There are some extremely useful non-parametric tests of difference, however, that have been developed independently (Meddis, 1984) and are not available on any existing statistical package as far as we are aware. We have therefore produced routines for these in Excel as the package AQB, which is available free on the

Internet/World Wide Web, and for which we include examples of on-screen input and output for the relevant tests in the present chapter. As in previous editions of the book, we also include calculations for most of the simpler tests (Appendix II) so their mechanics are clear and they can be done by hand if required.

### Test statistic

Significance tests calculate a test statistic that is usually denoted by a letter or symbol:  $t$ ,  $H$ ,  $F$ ,  $\chi^2$ ,  $r$ ,  $r_s$  and  $U$  are a few familiar examples from various parametric and non-parametric tests. The value of a test statistic has a known probability of occurring by chance for any given sample size or what are known as *degrees of freedom* (see later). A calculated value can thus be checked to see whether it concurs with or passes (positively or negatively, depending on the test) the threshold value appropriate to the level of probability chosen for significance. This used to mean comparing the value with a table of threshold values but such comparisons are now made automatically for the tests in many statistical computer packages.

### 3.3.3 Tests for a difference

Tests of differences between groups involve asking whether there are differences in overall counts, or in mean (parametric tests) or median (non-parametric tests) values. Since non-parametric tests replace the data with their ranked values, these tests actually test for differences in the mean ranks, which, when decoded back to the original data, represent the medians. Apart from tests of counts, all these tests involve comparing the variation in values within groups with the differences in the central values (mean or median) between groups: they are therefore all forms of '*analysis of variance*'. We shall introduce three types of difference test, with parametric and non-parametric equivalents as appropriate: first  $\chi^2$  (chi-squared, pronounced 'ky-squared'), then the  $t$ -test and Mann-Whitney  $U$ -test as parametric and non-parametric (respectively) tests for differences between *two* groups, and finally other analyses of variance (parametric and non-parametric) dealing with differences between two or more groups.

#### Tests for a difference between two groups

We shall start with the relatively simple situation of comparing two groups. Here, two mutually exclusive groups (e.g. male/female, small/large, with property  $a$ /without property  $a$ , etc.) have been identified and measurements made with respect to each (e.g. the body length of males versus the body length of females, the number of seeds set by small plants versus the number set by large plants, the survival rate of mice on drug A versus the survival rate on drug B). Depending on the kind of measurement made, we can use one of a number of tests to see whether any differences are significant.

**Chi-squared ( $\chi^2$ ).** A chi-squared test can be used if data are in the form of counts, i.e. if two groups have been identified and observations classified in terms of the number belonging to each. Chi-squared can be used *only* on raw counts; it cannot be used on measurements (e.g. length, time, weight, volume) or proportions,

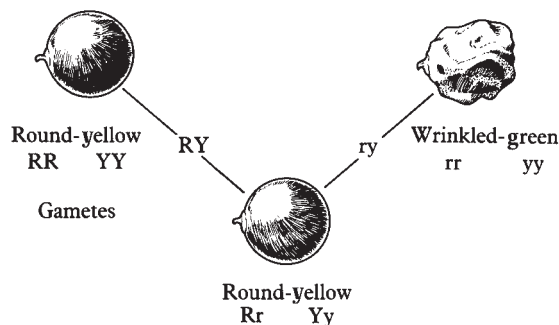
percentages, or any other derived values. The test works by comparing observed counts with those expected by chance or on some prior basis. As an example, we can consider a simple experiment in Mendelian inheritance. Suppose we crossed two pea plants that are heterozygous for yellow and green seed colour, with yellow being dominant. Our expectation from the principles of simple Mendelian inheritance, of course, is that the progeny will exhibit a seed colour ratio of 3 yellow : 1 green. We can use the chi-squared test to see whether our observed numbers of yellow and green seeds differ from those expected on a 3 : 1 ratio. The expected numbers are simply the total observed number of progeny divided into a 3 : 1 ratio. Thus:

	Seed colour		Total
	Yellow	Green	
Number observed	130	46	176
Number expected	132	44	176

To find our  $\chi^2$  test statistic using SPSS and AQB, we can follow the procedure in Box 3.2.

In the pea example, the expected numbers were dictated by the Mendelian theory of inheritance; we have good reason to expect a 3 : 1 ratio of yellow : green seeds and thus a difference between groups. In many cases, of course, we should have no particular reason for expecting a difference and our expected numbers for the two groups would be the same (half the total number of observations each). Thus if our yellow and green groups had referred to pecks by chicks at one of two different coloured grains of rice on a standard background instead of the inheritance of seed colour, our chi-squared table would have looked very different:

	Grain colour	
	Yellow	Green
Number of pecks observed	130	46
Number of pecks expected	88	88





### Box 3.2 A test comparing counts classified into two groups ( $1 \times 2 \chi^2$ test)

In our example of Mendelian ratios in seed colour (*see text*), the data are nominal, in that each seed is classified into one of two colour groups (yellow or green). The nominal factor *colour* forms the groups, with two levels, *yellow* and *green*. We then arrive at the total numbers for the groups (frequencies of each colour) and analyse them using  $\chi^2$ . In this case, the null hypothesis  $H_0$  is that the totals will be no different from those expected under a Mendelian ratio of 3:1.

The data will be in one of two formats: either raw data in the form of a list of the colours of each individual seed, or the total numbers of seeds calculated for each colour (as in the table in the text). AQB deals with the ready-calculated totals, whereas SPSS expects the former, although you can get it to use the totals (*see below*).

#### Calculating a $1 \times 2 \chi^2$ using AQB

In AQB (Fig. (i)), click on the tab at the bottom of the screen labelled 'chi-sq' and use the upper input table marked '1-way' (since your data are categories [*yellow/green*] of a single classification [*colour*]). Enter the calculated totals in the *data: counts* line. Since you expect a 3:1 ratio, divide your total number into this ratio ( $\frac{3}{4}N$  and  $\frac{1}{4}N$ ) to generate the expected numbers, and enter these in the *expected* line. The result appears automatically in the 'RESULTS' table on the right. In this case, the value of  $\chi^2$  with one degree of freedom (the number of groups minus 1) is 0.12. The probability (*p*) associated with a value as low as this is  $p = 0.73$ , which is much greater than 0.05, so we can conclude that the ratio of yellow:green seeds does not differ from our

expected Mendelian ratio and we do not reject the null hypothesis.

#### Calculating a $1 \times 2 \chi^2$ using SPSS

Instead of using pre-calculated totals, SPSS expects to count the numbers up itself, so usually uses data in the format shown in Fig. (ii)a. This is a pretty long-winded way to go about it, though, and, normally, you will have the totals rather than a list of the raw data. To get SPSS to deal with these, put the totals in two columns as shown in Fig. (ii)b: colour categories in one column (*colour*) and totals in the second (*frequency*). Make sure you use numerical codes for the colour groups, as in Fig. (ii)b; don't use letters (e.g. 'y', 'g') or names (e.g. 'yellow', 'green') because SPSS will sometimes not permit these to be used in the analysis (*see also* Box 2.2). Then click 'Data' followed by 'Weight Cases' from the menus at the top of the screen, and select the 'Weight cases by' option, entering the column containing the data, in this case *frequency* (Fig. (ii)c). SPSS will then read the data as totals rather than constituent numbers (but be sure to uncheck the 'Weight cases by' option afterwards, otherwise it will remain to affect all subsequent analyses). Then click 'Analyze' and 'Non-parametric Tests' followed by 'Chi-Square' (Fig. (ii)b). In the resulting dialogue box (Fig. (ii)d), put the variable *colour* into the 'Test Variable List' (if *colour* consists of names rather than numbers, it will not be listed as available). Then check the 'Values' option under 'Expected Values' in the order of the values in frequency, clicking 'Add' after each one to enter it into the 'Expected Values' box. The order in which you enter them is crucial

as SPSS will assume you have entered the expected values to match up to the right data values in calculating the output. Clicking 'OK' will then produce the analysis (Fig. (ii)e) (and you can double-check SPSS has read the data in correctly by checking that Expected values matching the correct Observed ones). Once again, a  $\chi^2$  value of 0.12 is returned, with an associated probability of 0.73.

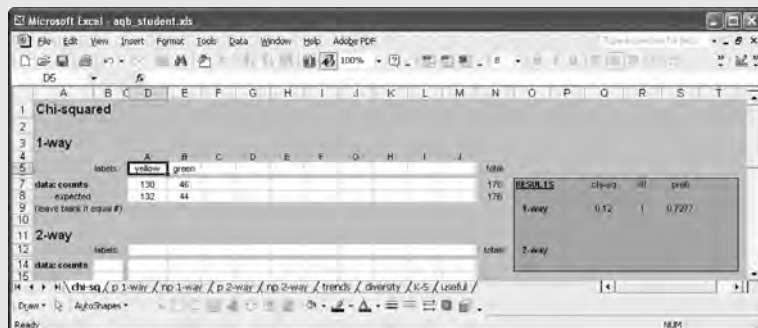
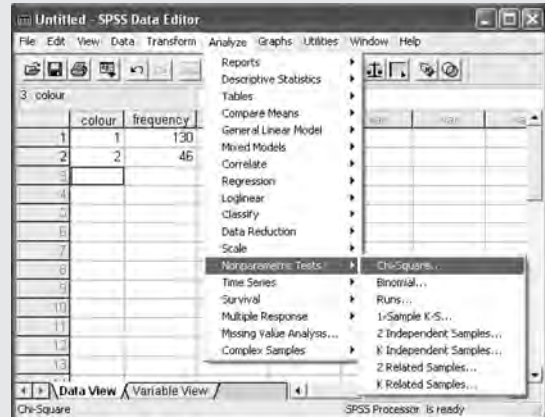



Figure (i) Performing a  $1 \times 2 \chi^2$  test in AQB.



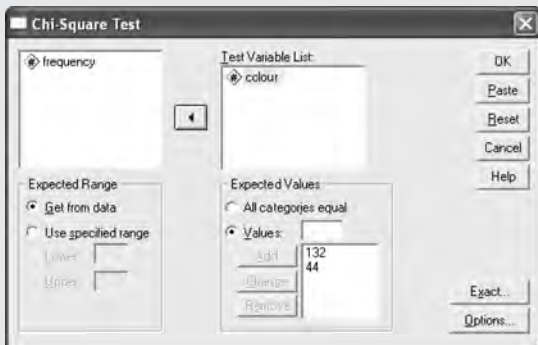
(a)



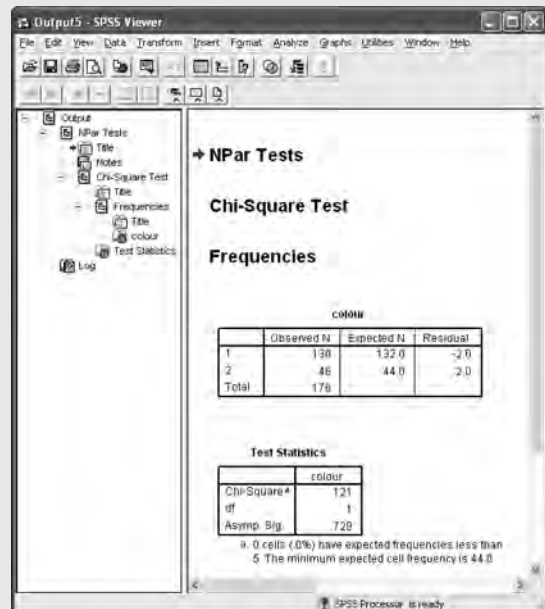
(b)



(c)



(d)



(e)

Figure (ii) (a) Standard data entry for a  $1 \times 2 \chi^2$  test in SPSS, (b) a short cut data entry for a  $1 \times 2 \chi^2$  test in SPSS, (c) the 'Weight cases' dialogue box for setting up the analysis, (d) the 'Chi-Square' dialogue box, (e) output from the analysis.

and the result would have been a  $\chi^2$  value of 40.09, which exceeds even the threshold value (10.83) for a significance level of 0.1 per cent (0.001) (*see* Appendix III, Table A). In this case we could safely reject the null hypothesis of no difference in the number of pecks to different coloured grains and infer a bias towards yellow on the part of the chicks.

An important point to bear in mind with  $\chi^2$  is that it is not very reliable when very small samples are being tested. While a correction can be brought into play here, it is a good idea not to use  $\chi^2$  when the sample size is smaller than 20 data values or when any expected value is less than 5.

***t*-tests and Mann–Whitney *U*-tests.** *t*-tests and Mann–Whitney *U*-tests can both be performed on raw counts like chi-squared, except that they deal with each contributing data value in the two groups separately instead of as a single total. Thus if the pecking data in our last example of chi-squared were derived from ten chicks given the opportunity to peck at yellow grains and ten more given the opportunity to peck at green grains on a standard background, the values that would be used in the chi-squared and the *t*- or *U*-tests can be indicated as follows:

Chick	Pecks to yellow	Chick	Pecks to green
a	12	k	2
b	14	l	3
c	13	m	10
d	3	n	6
e	23	o	4
f	13	p	5
g	11	q	3
h	15	r	1
i	9	s	7
j	17	t	5
Total	130		46

} A *t*- or *U*-test uses these values

} A chi-squared test uses these values (or, in principle, any subtotal of data values)

In addition, however, *t*- and *U*-tests can deal with data other than counts. We can use them to compare two groups for constant interval data such as body size, time spent performing a particular behaviour, percentage of patients responding to different drug treatments or a host of other kinds of data. The non-parametric *U*-test can also cope with ordinal (rank) data. In addition, unlike some kinds of two-group test working on individual data values, the *U*-test and most forms of the *t*-test do not require equal sample sizes in the two groups. However, and we stress this again, these tests can be used *only* for comparing *two* groups and only if each data value is *independent* of all other data values. Furthermore, the *U*-test can only test a general (or two-tailed) prediction (i.e. there is a difference, but not in any one predicted direction). Analyses for differences between two groups using AQB and SPSS are shown in Boxes 3.3a–d.

### Box 3.3a Mean values: a *general* parametric test for two groups (two-tailed *t*-test)

We shall use data for chicks pecking at yellow and green grains of rice. The data being analysed are thus measured on a constant interval scale – the number of pecks. Testing the residuals for normality (*see* Box 3.1) shows that they do not depart significantly from a normal distribution (Shapiro-Wilk = 0.98, d.f. = 80, ns), so we can analyse them perfectly reasonably using a parametric test. The nominal factor *seed colour* constitutes the grouping, with two categories, *yellow* and *green*.

First we frame our prediction, in this case the general prediction  $H_1$  that chicks will peck at different rates towards the two colours of grain; i.e. that the two groups *yellow* (A) and *green* (B) will have different mean values ( $A \neq B$ ). Thus the null hypothesis  $H_0$  is that the two mean values will not differ.

#### Testing for a general (two-tailed) difference in the means of two groups using AQB

In AQB (Fig. (i)), select the ‘two groups’ worksheet from the options at the bottom of the screen, and

enter a ‘2’ (for a general test) in the top portion of the top right-hand square in the first (‘RESULTS, *t*-test’) box. Then enter the data for each group in a separate column. The mean values, standard errors and sample sizes will automatically appear, together with the test statistic ‘*t*’ in the box labelled ‘*t*-test’. This value for *t* allows for the possibility that each group may have a different variance. The value of *t* is high (12.88), with a *p*-value of less than 0.0001, so we can reject the null hypothesis of no difference between the two groups.

#### Testing for a general (two-tailed) difference in the means of two groups using SPSS

In SPSS, we need to enter the data in the standard format, with the data values in a single column, and the group to which each value belongs coded in another column (Fig. (ii)a). Then click ‘Analyze’ and select ‘Compare Means’ and ‘Independent-samples T test’ (SPSS ignores the convention of using the lower case for ‘*t*’) to obtain the dialogue box in Fig. (ii)b. Transfer the name of the data column (*pecks*) into the ‘Test Variable’ box, and the name of the factor (*colour*) into the ‘Grouping Variable’ box. Then click on ‘Define Groups’ to declare to SPSS that your groups are indexed by the numbers ‘1’ and ‘2’ (in this example). Click ‘OK’, and the result will appear (Fig. (ii)c). The ‘Group Statistics’ box presents the mean, sample size, standard deviation and standard error for each group, and the test results themselves appear in the ‘Independent Samples Test’ box below. The latter first gives a test for the equality of the variances between the two groups (called Levene’s test, with an ‘*F*’ value as its test statistic). If the *p*-value associated with this test is significant ( $p < 0.05$ ), then use the *t* values quoted

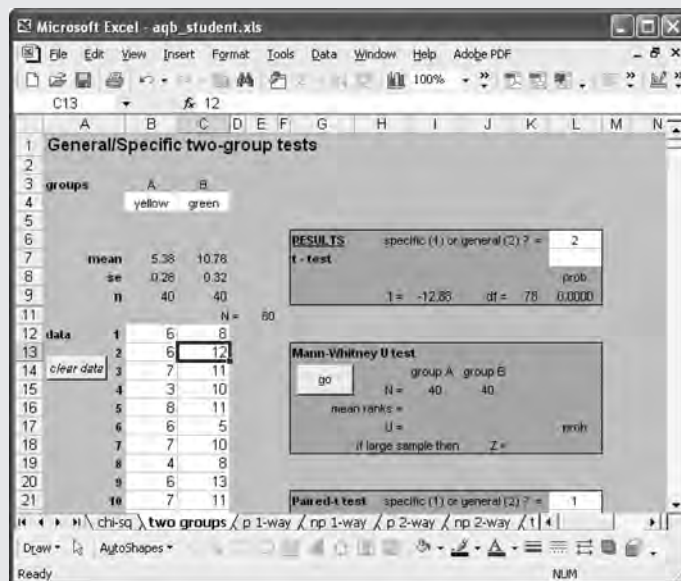
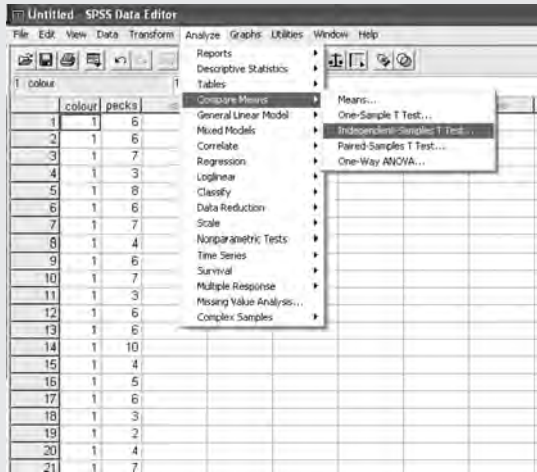
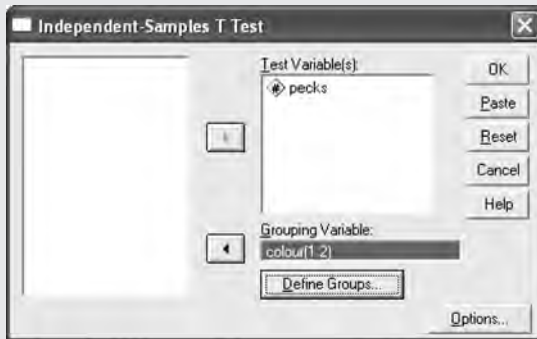


Figure (i) Data entry and analysis output for comparing the means of two groups in a two-tailed test in AQB.



(a)



(b)

The screenshot shows the SPSS Output Viewer window displaying the results of an Independent-Samples T Test. The output is organized into two main sections: 'Group Statistics' and 'Independent Samples Test'.

**Group Statistics**

colour	N	Mean	Std. Deviation	Std. Error Mean
pecks 1	40	5.38	1.750	.277
pecks 2	40	10.78	1.893	.315

**Independent Samples Test**

	Levene's Test for Equality of Variances	t-test for Equality of Means								
		F	Sig.	t	df	Sig. (2-tailed)	Mean Difference	Std. Error Difference	95% Confidence Interval of the Difference	
									Lower	Upper
pecks	Equal variances assumed	.010	.922	-12.677	78	.000	-5.400	.418	-6.235	-4.565
	Equal variances not assumed			-12.677	76.708	.000	-5.400	.418	-6.235	-4.565

(c)

in the second line of the box; otherwise use the values in the first line. In this case, the  $F$ -value is not significant, so the first line is the correct one to use. Once again (of course), the conclusion is that the two groups show evidence of having different mean values ( $t = 12.88$ ,  $d.f. = 78$ ,  $p < 0.001$  – unfortunately it is another irritating habit of computer statistical packages to cite very low probabilities as '.000', as both AQB and SPSS have done here; the probability in these cases is not actually zero, so one is left with a meaningless value that can't be cited as it stands in reports. One must therefore cite a threshold value above this, so  $p < 0.001$  would do in this case).

Figure (ii) (a) The data format and selection of test for comparing the means of two groups in a two-tailed test in SPSS, (b) transferring the data and defining the groups for analysis, (c) the output of the analysis.

### Box 3.3b Mean values: a *specific* parametric test for two groups (one-tailed *t*-test)

As in Box 3.3a, the data concern the number of pecks by chicks to yellow and green grains of rice, but this is a different experiment and hence different data. The residuals were checked for normality (following Box 3.1), and they could be assumed to be normally distributed (Shapiro-Wilk = 0.97, d.f. = 80, ns).

Once more, we frame our prediction; this time, though, we are going to predict that one particular group will have a larger mean than the other – i.e. we shall decide *in advance of collecting the data* which of the two groups (A and B) is predicted to have the greater mean value. There needs to be some a priori reason (theory, or previously published or gathered data) for this prediction; it can't just be made on a whim. Suppose that on the basis of your knowledge, you predict that chicks should peck more at yellow than at green seeds, i.e. that  $A > B$  ( $H_1$ ). The null hypothesis ( $H_0$ ) is that the mean value of A is not greater than that of B.

#### Testing for a specific (one-tailed) difference in the means of two groups using AQB

In the 'two groups' worksheet of AQB again (Fig. (i)), choose a specific test by inserting '1' in the top half of the box in the top right corner of the first 'RESULTS' box, and enter the group expected to have the larger mean value (in this case, 'A') in the lower half. Notice that the value of  $t$  is large (8.85) but negative (Fig. (i)), indicating that, contrary to prediction, the mean of B is larger than that of A. Because of this, the  $p$ -value says 'ns', i.e. not significant. Had you predicted the order the other

way round ( $B > A$ ), then you would have obtained a significant result (if you replace 'A' by 'B' in the lower half of the test selector box, you will see that the  $t$ -value becomes positive, and the probability changes from 'ns' to 0.0000, actually half the two-tailed probability), but as it is, the result is not significant, *and you most certainly cannot change the prediction retrospectively to make it significant.*

#### Testing for a specific (one-tailed) difference in the means of two groups using SPSS

In SPSS, there is no facility for doing one-tailed tests. You therefore need to look at whether the pattern of mean values falls into the predicted pattern, and if it does, then halve the two-tailed probability generated for  $t$  to obtain the one-tailed probability of getting your  $t$ -value. Thus, had you predicted chicks would peck more at green grain than at yellow in the example of Box 3.3a, the two-tailed probability (Box 3.3a, Fig. (ii)c)

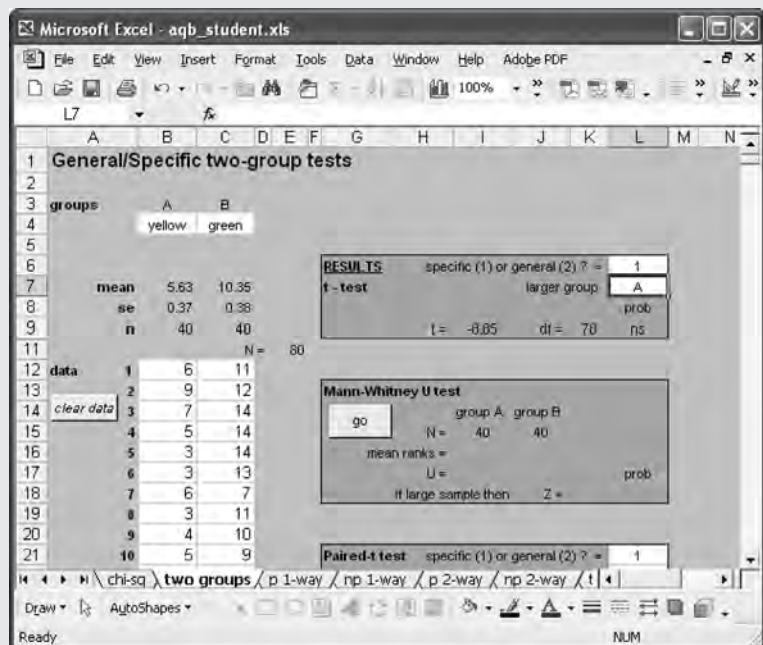


Figure (i) Data entry and analysis output for comparing the means of two groups in a one-tailed test in AQB.

stands at  $p < 0.001$ , so the one-tailed probability, at half the value, will also be highly significant. *Once again, though, it must be stressed that you can apply a one-tailed test like this only if you had a sound reason for predicting the direction of*

*the difference in advance.* In the current example, however, the conclusion would once again be that there is no evidence of the mean values falling into the predicted pattern ( $t = -8.9$ , d.f. = 78, ns), and the null hypothesis cannot be rejected.

### Box 3.3c Mean ranks (medians): a *general* non-parametric test for two groups (Mann–Whitney $U$ )

Suppose an ecologist was interested in the effect of microhabitat on the distribution of periwinkles (*Littorina* spp.) on a rocky shore. Two habitats – a boulder/shingle beach and crevices in a rocky stack – were compared for the prevalence of the commonest periwinkle species measured as the percentage of the total number of all invertebrate individuals recorded within quadrat samples.

The data being analysed are measured on a constant interval scale – percentage – but these are bounded by zero and 100, and are therefore unlikely to be normally distributed unless all values are well away from the boundaries. Testing the residuals for normality (as in Box 3.1) shows that the distribution is indeed far from normal (Shapiro-Wilk = 0.42, d.f. = 80,  $p < 0.001$ ), which remains so even when the data have been arcsine-square-root transformed (*see text*) to ‘stretch out’ the tails: hence a non-parametric test is required. The nominal factor *habitat* forms the level of grouping, with two categories: shingle *beach* and *rock* crevices.

Frame the prediction, in this case the general prediction  $H_1$  that the prevalence of periwinkles will differ between the two kinds of habitat. In terms of our non-parametric test, we are predicting that the two groups *beach* (A) and *rock* (B) will have different mean ranks ( $A \neq B$ ). Thus the null hypothesis  $H_0$  is that the two mean ranks will not differ.

#### Testing for a general difference in the medians of two groups using AQB

In AQB (Fig. (i)), choose the ‘two groups’ worksheet as in Boxes 3.3a,b, enter the data for each group in a separate column, and click the ‘go’ button of the lower box labelled ‘Mann–Whitney  $U$ -test’. This will give you the sample size in each group, the mean ranks, and the value of the test statistic,  $U$  (for small sample sizes  $< 20$  in both groups) and  $Z$  (for large sample sizes  $> 20$  in either group). The value of  $U$  needs to be compared with the critical values for  $U$  in Appendix III, Table B. If the value is less than the threshold value for a probability of 0.05, we can reject the null hypothesis that there is no difference between the groups. Note that in this test we

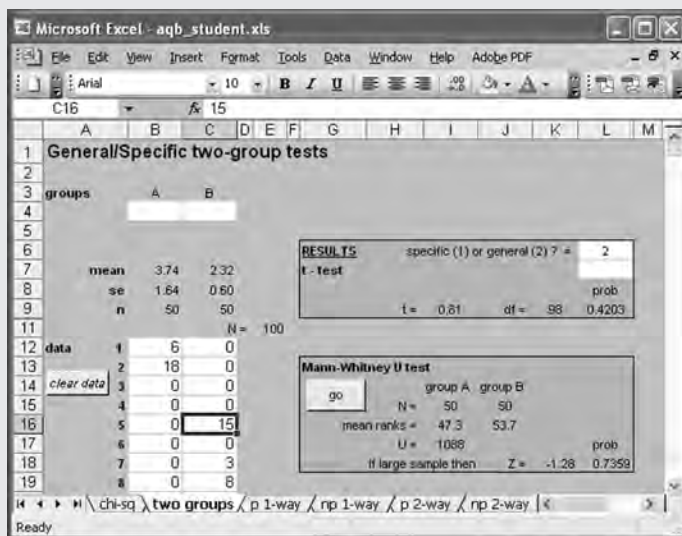
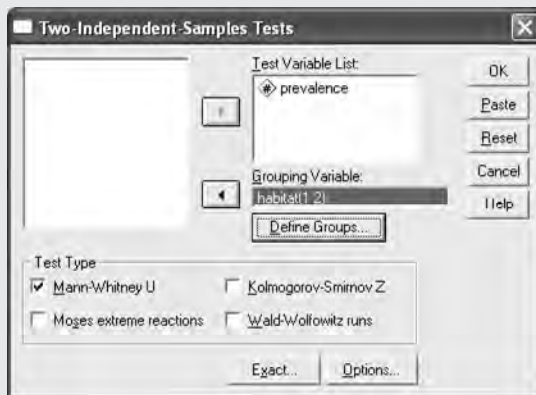


Figure (i) Data entry and analysis output for comparing the medians of two groups in a general test in AQB.

	habitat	prevalence
1	1	6
2	1	18
3	1	0
4	1	0
5	1	0
6	1	0
7	1	0
8	1	1
9	1	1
10	1	0
11	1	9
12	1	0
13	1	1
14	1	57
15	1	1
16	1	0
17	1	2
18	1	0
19	1	0
20	1	58
21	1	0

(a)



(b)

Ranks				
	habitat	N	Mean Rank	Sum of Ranks
prevalence	1	50	47.26	2363.00
	2	50	53.74	2687.00
	Total	100		

Test Statistics*	
	prevalence
Mann-Whitney U	1068.000
Wilcoxon W	2363.000
Z	-1.281
Asymp. Sig. (2-tailed)	.200

(c)

use the sample sizes  $N_A$  and  $N_B$  rather than the degrees of freedom to determine our threshold value. In this case, we have large sample sizes, and AQB has calculated  $z$ , and provided automatically the probability under the null hypothesis. Since  $z = 1.12$  (ns), the difference in mean ranks is obviously not significant. We ignore the sign of the  $z$ -value here because we have posed a general rather than a specific prediction, and the sign is only relevant to a test of a specific prediction.

### Testing for a general difference in the medians of two groups using SPSS

In SPSS, the data are entered in the standard format of the measured values in a single column, and the group to which each value belongs in another column (Fig. (ii)a). The habitat codes must be numbers rather than letters: if SPSS thinks that *habitat* is a string (text) variable rather than a numeric one, the variable will not be available in the test. Then click 'Analyze' followed by 'Non-parametric Tests' and '2 Independent Samples' to obtain the dialogue box in Fig. (ii)b. Transfer the data variable (here, *prevalence*) into the 'Test Variable List' box, and the factor (here, *habitat*) into the 'Grouping Variable' box, and tell SPSS what the group category codes are by clicking on 'Define Groups' and typing in the codes you have used (in the example, these are '1' and '2'). Make sure that the 'Mann-Whitney U' box is checked, and click 'OK'.

The resulting output (Fig. (ii)c) shows three test statistics. They are all different ways of testing the same hypothesis – for small sample sizes we want the 'Mann-Whitney  $U$ ', whereas (as here) when the sample sizes are large we use the  $z$ -value. Once again, we conclude that there is no significant difference between habitat types in the prevalence of periwinkles ( $z = 1.28$ , ns).

**Figure (ii)** (a) The data format and selection of test for comparing the medians of two groups in a general test in SPSS, (b) transferring the data and defining the groups for analysis, (c) the output of the analysis.



**Box 3.3d Mean ranks (medians): a *specific* non-parametric test for two groups**

We shall analyse the same data as in Box 3.3c for this test. Assume that from the literature the ecologist is aware that periwinkles are said to form a greater proportion of the invertebrate community in rocky habitats. The prediction is therefore that periwinkles will show a higher prevalence in the *rock* habitats than in the *beach* habitats, i.e.  $H_1$  is that  $B > A$ . Thus the null hypothesis is that  $B$  is not greater than  $A$ , i.e.  $H_0$  is that  $B = < A$ .

**Testing for a specific difference in the medians of two groups using AQB**

In AQB, unless there are large enough sample sizes ( $n$  for either group  $> 20$ ) so that the  $z$  test statistic is calculated, it is not possible to do a specific non-parametric version of the Mann–Whitney. If a  $z$ -score is available, and the difference is in the predicted direction, then merely halve the probability to get the specific (one-tailed) probability. In the dataset here, notice that while the mean rank for  $B$  (53.7) is indeed greater than the mean rank for  $A$  (47.3), the arithmetic mean values (shown above the data) are actually the other way round. This situation underlines the necessity for checking the residuals for conformity to the normal distribution, since a specific (one-tailed)  $t$ -test of this hypothesis would be very misleading. Here, the mean ranks are in the predicted pattern, and hence we can halve the probability of the  $z$ -score (which is available because of the large samples). Even so, the result is not significant.

**Testing for a specific difference in the medians of two groups using SPSS**

In SPSS we can perform a specific (one-tailed) test by choosing to calculate the exact test by clicking the ‘Exact’ button in the dialogue box of Box 3.3.c, Fig. (ii)b. Then check the ‘Exact’ option, and ‘Continue’. Clicking ‘OK’ yields the ‘Test statistics’ box (Fig. (i)) which contains some extra information relative to Box 3.3c, Fig. (ii)c, including ‘Exact Sig. (one-tailed)’. As long as the mean ranks fall into the predicted pattern, you can use this probability as the result. The conclusion is that the data do not support the prediction that periwinkles are at a greater prevalence in rocky habitats (one-tailed exact probability = 0.101).

**Test statistics<sup>a</sup>**

	Prevalence
Mann–Whitney $U$	1088.000
Wilcoxon $W$	2363.000
$Z$	−1.281
Asymp. Sig. (two-tailed)	0.200
Exact Sig. (two-tailed)	0.202
Exact Sig. (one-tailed)	0.101
Point probability	0.001

<sup>a</sup> Grouping variable: habitat

**Figure (i)** SPSS output for a one-tailed exact probability comparing the medians of two groups in a specific test.

**Tests for two related samples**

The tests for two groups introduced so far assume that data in the different groups are independent of each other. In many cases, the values of the data in each group are not independent, but are related in some way. For example, we might have tested one leaf of a plant with one treatment, and another leaf of the same plant with a different treatment. Or we might have provided a sample of ten female Siamese fighting fish (*Betta splendens*) with a red male as a potential mate on one occasion, and a blue male on another occasion, and measured their reactions to each of them. In order to test for the effect of different treatments with these kinds of data, we need a method that takes the *non-independence* of the data into account. In these kinds of experiments, the data usually come in

replicated pairs (as in the ‘pair’ of males presented to each female above), and the appropriate statistical tests are known as ‘paired-sample’ tests. The sample sizes of the two groups are therefore necessarily the same.

The methods that deal with this design work by taking the differences between the measurements of each pair, and then analysing these differences. In the parametric case, the *paired t-test* tests whether the mean of these differences is significantly different from zero. In the non-parametric case, the *Wilcoxon matched-pairs signed ranks test* ranks all the differences ignoring whether they are positive or negative, and then tests whether the mean rank of the positive differences is the same as that of the negative ones. A further test is also possible here, since the null hypothesis is that there are equal numbers of positive and negative differences: the *binomial test* tests for a departure from this equality. Boxes 3.4a–d show how to do some of these tests in AQB and SPSS.

### Box 3.4a Mean values: a *general* parametric test for *related* samples in two groups

A botanist wanted to assess whether leaves treated with nitrogen differed in length from untreated leaves. She therefore treated one leaf with water and another leaf of the same plant with a solution of nitrogen, replicating this across 40 plants. After a suitable length of time, she measured the length of each leaf.

The data being analysed are measured on a constant interval scale – length. The nominal factor ‘treatment’ forms the groups, with two categories: water-treated (*control*) and nitrogen-treated (*N*). The (nominal) blocking factor (*see text*) that relates values in the two groups is plant individual, with 40 levels. A related-samples test looks at the differences between each pair of values, and it is these differences that should be normally distributed for a parametric test. The differences need to be calculated in SPSS (using ‘Transform’ and ‘Compute’) (see Box 3.1) into a new variable, and then this tested in the usual manner. When we do this, we find that the differences can indeed be assumed to be normally distributed (Shapiro-Wilk = 0.98,  $n = 40$ , ns).

Frame the prediction. Here it is a general prediction that the treatment will change the length of the leaf within plants, so that water-treated controls (A) and nitrogen-treated (B) leaves will differ in length, i.e.  $H_1$  is that  $A \neq B$ .  $H_0$  is therefore that  $A = B$ .

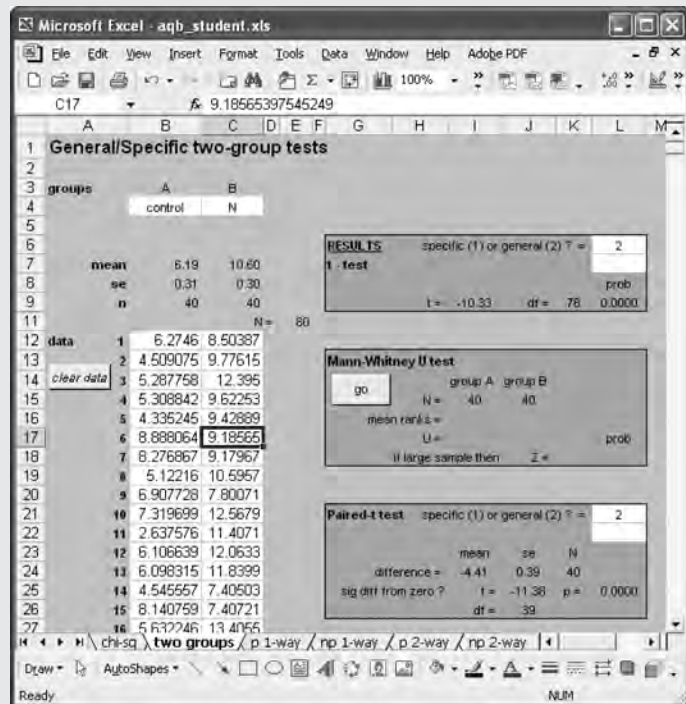
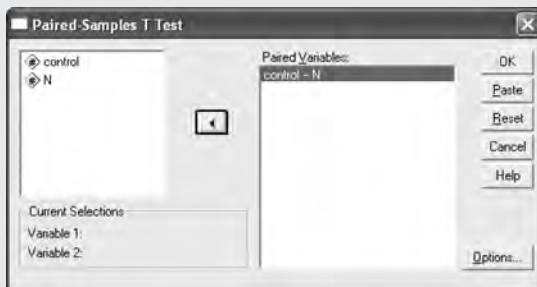


Figure (i) Data entry and output for a two-tailed paired *t*-test in AQB.

	control	N
1	6.2746	8.5039
2	4.5091	9.7761
3	5.2678	12.395
4	5.3088	9.6205
5	4.3352	9.4289
6	8.8881	9.1857
7	8.2769	9.1797
8	5.1222	10.596
9	6.9077	7.8007
10	7.3197	12.588
11	2.6376	11.407
12	6.1066	12.063
13	6.0963	11.840
14	4.5456	7.4050
15	8.1408	7.4072
16	5.6322	13.405
17	5.0666	10.010
18	7.5832	8.1025
19	6.8974	8.1979
20	7.7269	12.400

(a)



(b)

Paired Samples Statistics				
Pair	Mean	N	Std. Deviation	Std. Error Mean
1 control	6.193136	40	1.9404639	.3068143
1 N	10.601218	40	1.8749778	.2983177

Paired Samples Correlations			
Pair 1	control & N	r	Sig.
1	control & N	.176	.278

Paired Samples Test							
Pair 1	control - N	Paired Differences			95% Confidence Interval of the Difference		Sig. (2-tailed)
		Mean	Std. Deviation	Std. Error Mean	Lower	Upper	
1	control - N	-4.40806	2.4453461	.3872756	-5.19142	-3.62474	.11302
							df = 38

(c)

Figure (ii) (a) Data entry and test selection for a two-tailed paired  $t$ -test in SPSS, (b) transferring paired variables for analysis, (c) output for the analysis.

### Testing for a general difference between related samples in two groups using a parametric test (two-tailed paired $t$ -test) in AQB

In the 'two groups' worksheet of AQB, choose the output table labelled 'Paired  $t$ -test' and enter 2 (for a general test) in the top right corner box, then enter the data for each group in separate columns. Since the data are paired, make sure that the data for each pair (in this case, each individual) are in the same row, and therefore that there is exactly the same number of data values in each column. The paired  $t$ -test result will appear automatically in the output box (Fig. (i)). The difference between the paired values amounts to an average of  $4.41 \pm 0.39$  mm ( $n = 40$  pairs), which is significantly different from zero ( $t = 11.4$ , d.f. = 39,  $p < 0.001$ ). The conclusion, therefore, is that treating leaves with nitrogen significantly changes the length of leaves.

### Testing for a general difference between related samples in two groups using a parametric test (two-tailed paired $t$ -test) in SPSS

In SPSS too the data for paired samples go into two separate columns (note the difference compared with the independent samples  $t$ -test), and then click 'Analyze' followed by 'Compare Means' and 'Paired-Samples T test' (Fig. (ii)a). Select (by highlighting them simultaneously) and transfer the two variables together into the 'Paired Variables' box (Fig. (ii)b): SPSS will assume they are paired by rows. Click 'OK' and the output box (Fig. (ii)c) shows the summary statistics for each variable, then the correlation between them (which is irrelevant here), and then the paired-sample  $t$ -value of 11.4, leading to same conclusion as above.

### Box 3.4b Mean values: a *specific* parametric test for *related* samples in two groups

Using the same example of leaf length as Box 3.4a, decide *in advance of collecting the data* which of the two groups (A or B) is predicted to have the greater mean value. As always, there needs to be some a priori reason (theory, or previous published or gathered data) for this prediction. Suppose that, on the basis of her knowledge, the botanist in Box 3.4a predicted that a leaf treated with nitrogen would be *longer* than the control leaf on the same plant, i.e. that  $B > A$  ( $H_1$ ). The null hypothesis ( $H_0$ ) is that the mean value of B would not be greater than that of A.

#### Testing for a specific difference between related samples in two groups using a parametric test (one-tailed paired *t*-test) in AQB

In AQB, using the ‘two groups’ worksheet, choose a specific test by inserting ‘1’ in the top right corner of the ‘Paired *t*-test’ output box (see Box 3.4a, Fig. (i)), and select which group is expected to have the larger mean value (‘B’). Notice that the mean difference changes from negative to positive, reflecting the fact that the differences are now taken to be B-A rather than the default A-B. The value of *t* also changes from negative to positive, and the significance changes from ‘ns’

(caused by the combination of a specific test but a negative *t*-value) to ‘0.0000’ (or  $< 0.0001$ , as we should cite it). The difference is clearly highly significant in the predicted direction.

#### Testing for a specific difference between related samples in two groups using a parametric test (one-tailed paired *t*-test) in SPSS

Again, in SPSS there is no facility for doing specific tests other than noting that the *t*-value it provides (Box 3.4a, Fig. (ii)c) is for a two-tailed test. We therefore need to look at whether the pattern of mean values falls into the predicted pattern, and if it does, halve the two-tailed probability to obtain the one-tailed probability of getting the declared *t*-value (*but again, note the health warning on this in Box 3.3b*). Clearly, since the difference in the example (Box 3.4a, Fig. (ii)c) is in the predicted direction, and the two-tailed probability is already so low as to be displayed as ‘.000’ in the output, the one-tailed probability will also be highly significant.

The conclusion from both tests is that the data support the idea that nitrogen-treated leaves grow to be longer than control leaves on the same plant (paired  $t = 11.4$ , d.f. = 39,  $p < 0.001$ ), and the null hypothesis can be rejected.

In fact, having data values in pairs like this is only a particular case of data occurring as *related samples*, where data values are not independent of each other for a wide range of reasons, for instance because subjects have been kept in the same cage and have been exposed to the same social experience that is different from the social experience in any other cage. Such sources of non-independence lead to data being treated in so-called *blocks*, where the blocking designation allows the non-independence to be taken into account in subsequent analyses. The experimenter may be completely uninterested in whether or not there is a significant block effect, but it must be allowed for in the analysis in order to see the true impact of the treatment. Box 3.5e shows an example of a more general related samples (or ‘repeated measures’) analysis.

### Box 3.4c Mean ranks: a *general* non-parametric test for related samples in two groups (Wilcoxon matched-pairs signed ranks test)

In an experiment on two mouse strains, the average number of offspring was compared over 50 years. The data being analysed are measured on a constant interval scale – the average number of offspring. The nominal factor *strain* forms the level of grouping, with two levels: *strainA* and *strainB*. The (nominal) blocking factor that relates values in the two groups is year, with 50 levels. As mentioned before (Box 3.4a), the normality assumption in related-samples tests concerns the differences between the related values. When the differences are tested here, there is strong evidence of a non-normal distribution (Shapiro-Wilk = 0.89,  $n = 50$ ,  $p < 0.001$ ). Thus a non-parametric test is appropriate.

Frame the prediction. The general prediction is that there will be a difference in the number of offspring between the two strains across years. Thus, allowing for the related samples (i.e. the differences among years), the prediction ( $H_1$ ) is that  $A \neq B$ .

#### Testing for a general difference between related samples in two groups using a non-parametric test in AQB

In the data in the ‘two groups’ worksheet in AQB (Fig. (i)), each year is a row in the table. Clicking ‘go’ in the ‘Wilcoxon matched-pairs test’ box shows that the mean of the positive ranks is 23.3, and the mean of the negative ranks is 29.1. The test statistic is  $T$ , and it has a value of 552 (note that  $T$  is different from  $t$  of the  $t$ -test). However, since the sample size is large enough, a  $z$ -value is calculated and this is the value we use. It is not large enough (0.83) to reject the null hypothesis, and thus the result is not significant.

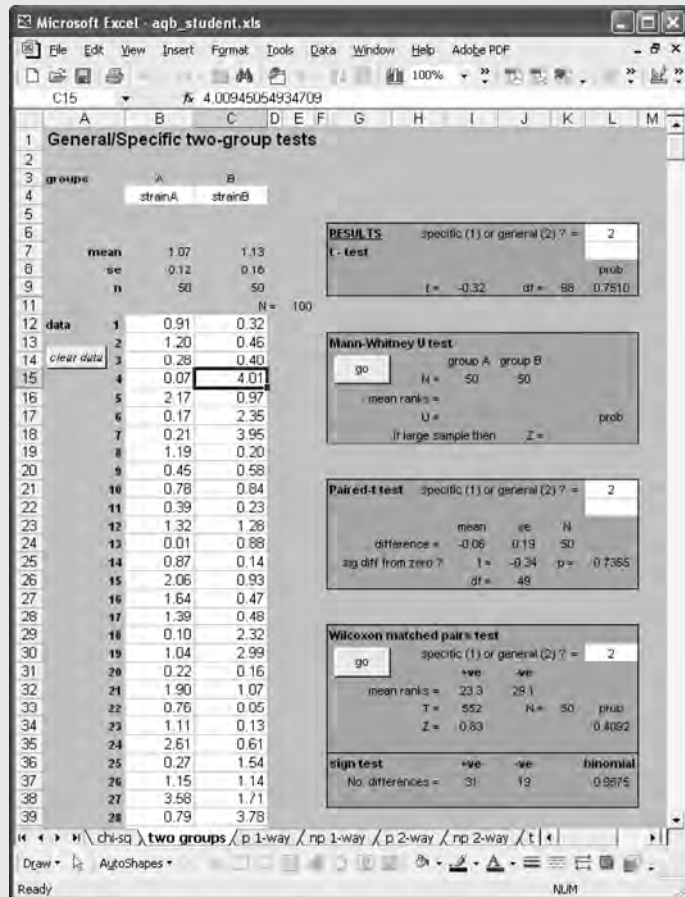



Figure (i) Data entry and output for a general Wilcoxon matched-pairs signed ranks test in AQB.

#### Testing for a general difference between related samples in two groups using a non-parametric test in SPSS

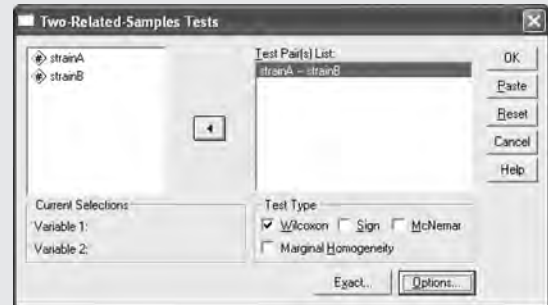
In SPSS (Fig. (ii)a), the data are again in the same format, with separate columns for the strains, and each row representing a block (year). Click ‘Analyze’, followed by ‘Non-parametric Tests’ and ‘2 Related Samples’, giving the dialogue box in Fig. (ii)b. Enter the two variables into the ‘Test Pair(s) List’ box *together* (as in Box 3.4a), so SPSS assumes they are matched by row. The Wilcoxon matched-pairs test is selected by default, so just



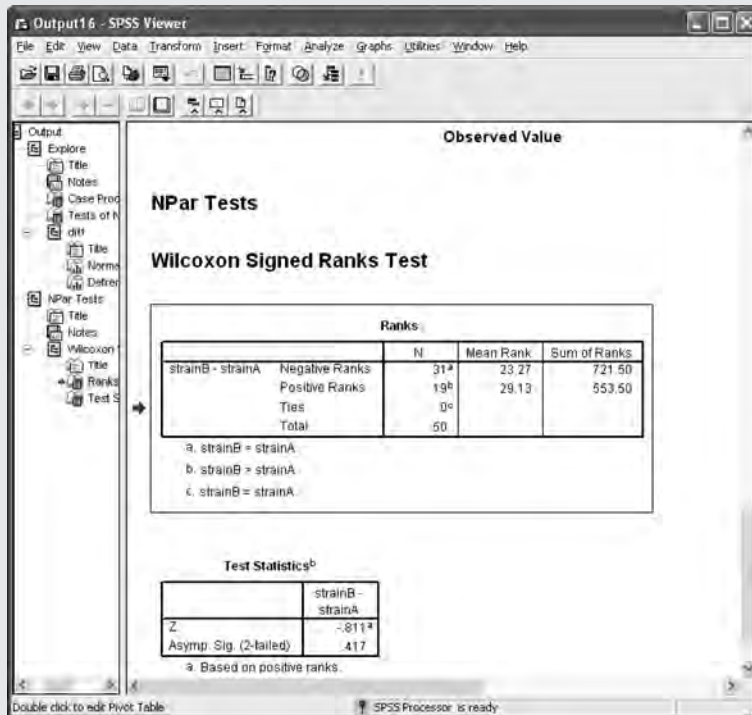
SPSS Data Editor window showing a data table with columns 'strainA' and 'strainB' and rows 1 through 20. The 'Analyze' menu is open, and 'Nonparametric Tests' > '2-Related Samples...' is selected.

	strainA	strainB
1	.91	.32
2	1.20	.46
3	.28	.40
4	0.7	4.01
5	2.17	.97
6	.17	2.35
7	.21	3.95
8	1.19	.20
9	.45	.58
10	.76	.84
11	.39	.23
12	1.32	1.28
13	.01	.88
14	.87	.14
15	2.06	.93
16	1.64	.47
17	1.39	.48
18	.10	2.32
19	1.04	2.99
20	.22	.16

(a)



(b)



Output 16 - SPSS Viewer window showing the results of a Wilcoxon Signed Ranks Test. The output includes a table of ranks and a table of test statistics.

Ranks				
		N	Mean Rank	Sum of Ranks
strainB - strainA	Negative Ranks	31 <sup>a</sup>	23.27	721.50
	Positive Ranks	19 <sup>b</sup>	29.13	553.50
	Ties	0 <sup>c</sup>		
	Total	50		

a. strainB = strainA  
b. strainB > strainA  
c. strainB = strainA

Test Statistics <sup>b</sup>	
	strainB - strainA
Z	-.811 <sup>a</sup>
Asymp. Sig. (2-tailed)	.417

a. Based on positive ranks.

(c)

click 'OK' to get the output (Fig. (ii)c). With small sample sizes (< 25) the asymptotic significance ('Asymp. Sig.') is unreliable, and you can choose to calculate the Exact probability by clicking 'Exact' in the dialogue box in Fig. (ii)b and selecting the 'Exact' option in the resulting dialogue

### Test statistics<sup>b</sup>

	StrainB - strainA
Z	-0.811 <sup>a</sup>
Asymp. Sig. (two-tailed)	0.417
Exact Sig. (two-tailed)	0.422
Exact Sig. (one-tailed)	0.211
Point probability	0.001

<sup>a</sup>. Based on positive ranks

<sup>b</sup>. Wilcoxon Signed Ranks Test

(d)

Figure (ii) (a) Data entry and test selection for a general Wilcoxon matched-pairs signed ranks test in SPSS, (b) transferring paired variables for analysis, (c) output for the analysis of large samples, (d) selecting the 'Exact' option, and the consequent output, for analysing small samples.

box. Click 'Continue' then 'OK', and the new output box gives the general (two-tailed) and specific (one-tailed) exact probabilities for the test (Fig. (ii)d). Our  $z$ -value ( $-0.81$ ) is clearly not significant, whether we use the approximate (asymptotic) or the exact probability.

### Box 3.4d Mean ranks (medians): a *specific* non-parametric test for related samples in two groups

There is no non-parametric test specially designed to test for this kind of analysis for two groups, but you can use the Exact probability option of SPSS to give you a one-tailed  $p$ -value (see Box 3.4c, Fig. (ii)d). Alternatively, use the specific non-parametric analysis of variance (see Box 3.5c), because this can cope with any number of groups.

### Tests for a difference between two or more groups

So far, we have introduced significance tests that can test for a difference between two groups. In many cases, of course, we shall be faced with more than two groups. For instance, Fig. 2.4 suggests that the amount of damage around the edges of leaves increases with leaf size class, perhaps indicating a predilection for big leaves by particular kinds of pest. If we wanted to know whether the difference between the size classes was significant, we should have to deal with three groups of data. How do we do it? The temptation to which many succumb is to do a *round robin* comparison of pairs of groups using a  $t$ - or  $U$ -test. In our Fig. 2.4 example, this would mean testing for a difference between small and medium leaves, then for a difference between medium and large leaves, and finally for a difference between small and large leaves. *The error of this cannot be emphasised too strongly.* The most serious problem arising from such a practice is that it increases the likelihood of obtaining a significant difference by chance when really none exists. To take an extreme example: if we carried out 100 two-group comparisons, then, just by chance, five of them stand to be significant at the 5 per cent level. Even if we made only 20 comparisons, one is likely to be significant by chance. While this may not seem a serious difficulty when we are dealing with only three or four groups, these examples illustrate the error in principle. To get round the problem, we need tests that can cope with comparisons between several groups at the same time.

**One-way analysis of variance.** Where we have series of data values falling into several groups (in a similar fashion to data values in the two groups of a  $t$ - or  $U$ -test), a one-way analysis of variance (often expressed as the acronym *one-way ANOVA*) is a suitable significance test. There are both parametric analyses of variance and non-parametric tests that have an equivalent function. We shall introduce both kinds here. While parametric analyses of variance are by far the more widely used, the non-parametric version we shall describe has the advantage not only of being robust to a wider range of data but also of allowing two kinds of specific predictions about the *direction* of differences between groups to be tested. This test can thus be used to test general or specific predictions for two or more groups. For two groups it is therefore more flexible than the  $U$ -test and, as a result, we recommend it even in these cases. Procedures for the two kinds of analysis of variance in AQB and SPSS are outlined in Boxes 3.5a–e. Procedures where there is a related samples design with more than two groups are shown in Box 3.5e.

### Box 3.5a Mean values: a *general* parametric one-way ANOVA for differences between two or more groups

We are interested in developing artificial techniques for rearing bumblebees (*Bombus* spp.) for pollinating greenhouse plants. Suppose we measure the weights of bumblebee queens that have been overwintered in the soil under four different conditions. The first condition is the normal rearing environment (soil), while different components (stones, leaves or cotton wool) are added to the rearing environment of the other groups. We think that any of these additional components might increase the weight of the resulting queens, and heavier queens in spring have a better chance of producing successful colonies.

The data being analysed are the weights of queens. The nominal factor *treatment* forms the level of grouping, with four categories (1 = normal soil, 2 = plus stones, 3 = plus leaves, 4 = plus cotton wool). A test of the residuals (see Box 3.1) shows that we can assume that they are normally distributed (Shapiro-Wilk = 0.99, d.f. = 160, ns).

Frame the prediction, in this case the general one that there are differences among the mean values of the treatment groups ( $H_1$ ). The null hypothesis ( $H_0$ ) is therefore that there are no differences among the groups.

#### Testing for a general difference in the means of two or more groups using a parametric one-way ANOVA in AQB

In AQB the data are entered as four columns, the data for each group in a separate column (Fig. (i)). The mean values and their standard errors appear above each column, and can be pasted into Excel for plotting (see Box 2.3). The results appear automatically in the 'one-way ANOVA table (general)' in the RESULTS box to the right. It is set out in the con-

ventional format of an analysis of variance output table, which it is standard practice to reproduce in reports and papers as the result of the test. The test statistic is  $F$ , which reflects the ratio of between-group to within-group variation (hence it is usually referred to as an  $F$ -ratio): since each of these sources of variation has a degree of freedom value, the test statistic has two sets of degrees of freedom. Thus, in this case, we should report  $F$  as  $F_{3,156} = 4.38$ ,  $p < 0.001$  (or  $F = 4.38$ , d.f. = 3,156,  $p < 0.001$ ).

#### Testing for a general difference in the means of two or more groups using a parametric one-way ANOVA in SPSS

In SPSS, the data should be in the usual format for independent samples, i.e. the data values in one column, indexed by a group code in a second column: then click 'Analyze', followed by 'Compare Means' and 'One-way ANOVA' (Fig. (ii)a). In the dialogue box, transfer the data variable (*weight*) into the 'Dependent List' box, and the grouping factor (*treatment*) into the 'Factor' box (Fig. (ii)b). Click 'Options', and, in the next dialogue box (Fig. (ii)c), make sure that the

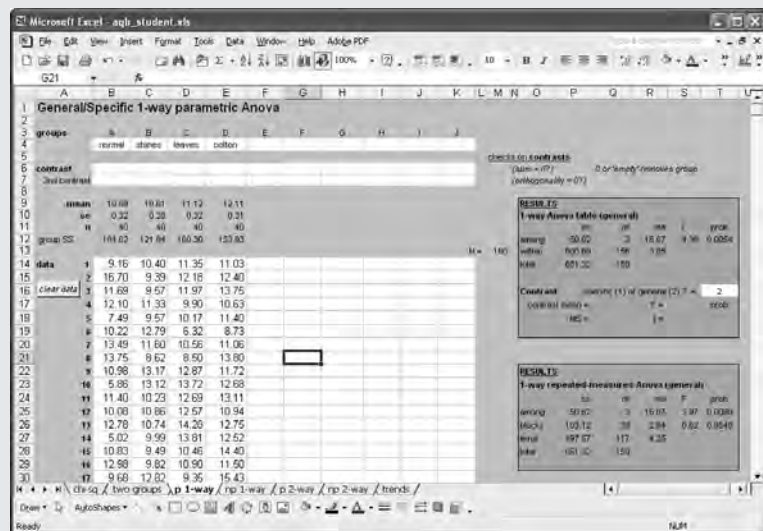
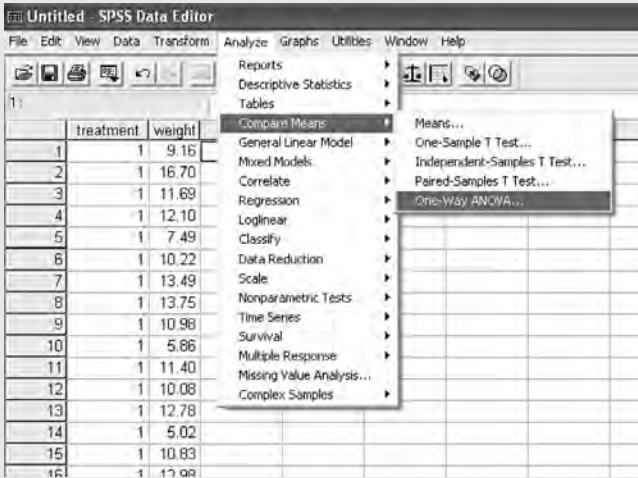


Figure (i) Data entry and results output for a general one-way ANOVA in AQB.

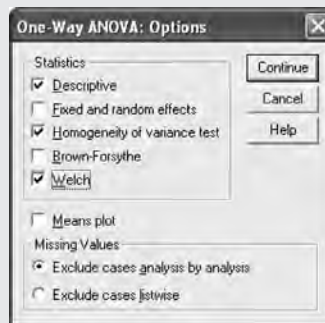




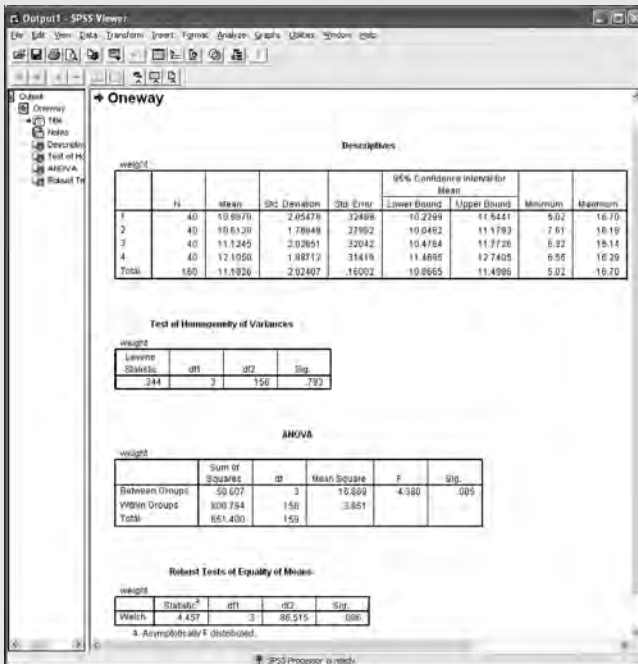
(a)



(b)



(c)



(d)

Figure (ii) (a) Data entry test selection for a parametric one-way ANOVA in SPSS, (b) entering variables for the analysis, (c) selecting supplementary tests and statistics, (d) output from the analysis.

'Descriptive', 'Homogeneity of variance test' and 'Welch' boxes are checked. The 'Descriptive' choice produces summary statistics for each group. We need to test for homogeneity of variance because this is an important assumption of ANOVA; if the variances turn out to differ significantly between groups, we shall need to use a modified ANOVA, and one of these is called a Welch test, hence this is checked just in case. Clicking 'Continue' and 'OK' produces the output (Fig. (ii)d). This contains the summary statistics for the groups first. Then there is a Levene's test for homogeneity of variances, which here is not significant (since the  $p$ -value is above 0.05) – hence the variances are sufficiently similar for the normal ANOVA to be valid. The ANOVA table comes next, which is identical to the one produced by AQB, with the  $F$ -ratio of 4.38. Finally, the Welch test is given in case the variances had been different among groups, but in this case we do not need it.

The conclusion in both cases is therefore that there is good evidence that the treatments affect the weight of overwintering queen bumblebees ( $F = 4.38$ , d.f. = 3,156,  $p < 0.01$ ).

### Post-hoc multiple comparisons

A general one-way ANOVA will tell us whether or not there is a significant difference between our groups, but it does not tell us where the difference lies. We cannot conclude which pairs of mean values are significantly different (although we can certainly suggest them from the group means), because we did not make any specific prediction beforehand, in advance of obtaining the data, about the ordering of the mean values, nor did we set up planned contrasts between particular sets of groups. Of course, it is entirely natural to want to delve further and find out where the differences lie, and many researchers would like to be able to make such *post hoc* tests. A plethora of different tests is available that test

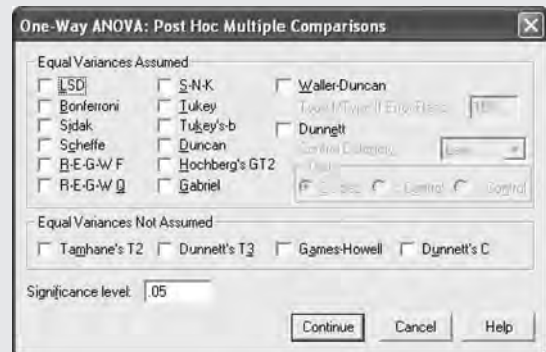


Figure (iii) The range of *post hoc* multiple comparison tests available in SPSS.

every mean value against every other one, and many of these are available in SPSS by clicking the 'Post Hoc' button in Fig. (ii)b above and taking our pick from the options offered in the resulting dialog box (Fig. (iii)).

The very diversity of methods ought to warn us that any one of them might be inadequate. Many biologists use them routinely, but others, and many statisticians, avoid them altogether. Detailed study shows that none of them can be relied upon to give accurate  $p$ -values for differences (Day and Quinn, 1989): many are too conservative, but others are not conservative enough, suggesting significance unjustifiably. Thus our view is that, if they are used at all, they should be taken only as a rough guide to where the differences might lie (so are not really much better than simply inspecting the group means). None can be recommended over the others. Many people use the Least Significance Difference (LSD). All this does is to take the within-groups ('error') mean square from the analysis of variance and to construct a 95 per cent confidence limit with it. However, the LSD 'test' is only accurate for a priori contrasts (Day and Quinn, 1989).

**Box 3.5b** Mean values: a *specific* parametric one-way ANOVA

Suppose we have measured the light absorbance of an indicator of a particular biochemical reaction under each of four temperature conditions. We were unable to arrange the experiment so as to have equal numbers of replicates (which is always desirable), but we do have several replicates for each group. The data being analysed – light absorbance – are on a constant interval scale. The nominal factor *temperature* forms the level of grouping, with four levels (A = 10 °C, B = 15 °C, C = 20 °C and D = 25 °C). Note that although the temperature levels would be treated as nominal groups in a general ANOVA, here they not only form a rankable series, but they actually consist of constant-interval measurements that are equally spaced (by 5 degrees). This will become important in a moment. Testing the residuals for normality (Box 3.1) shows that we can assume normality (Shapiro-Wilk = 0.97, d.f. = 35, ns).

When it comes to framing our predictions, theory (and past practice) tells us that temperature speeds up biochemical reactions, and, furthermore, that a 10 °C rise in temperature should double the rate of reaction. We therefore predict that the rank order of the mean values of light absorbance should be  $A < B < C < D$ . Unfortunately, such a simple specific prediction is hard to test because it says nothing about the magnitude of the changes between the groups. In theory it should be testable (see Gaines and Rice, 1990), and indeed this can be done in some statistical packages (such as R), but it is complicated and therefore not available in AQB or SPSS, or indeed in most generally available packages.

Theory says there should be a particular relationship between temperature and the rate of reaction. In this experiment with light absorbance, which is measured on a logarithmic scale, as our measure, this translates into a prediction of a *linear positive* relationship between temperature and light absorbance. Because (and only because) the temperature groups are evenly spaced, we can test this prediction using so-called polynomial coefficients (which require even spacing between the levels) in SPSS, but not in AQB.

**Testing for a specific difference in the means of two or more groups using a parametric one-way ANOVA in AQB**

In AQB (and SPSS) we can use the alternative of independent contrasts to test almost the same prediction. These make a contrast between two subsets of the groups (the factor levels). If the linear positive relationship between temperature and light absorbance is correct, then the mean light absorbance of the 10 °C group will be less than the average of the other groups. This can be tested using a contrast.

We have three degrees of freedom to use, and, since each contrast involves one degree of freedom, we can frame three independent (so-called orthogonal) contrasts. We can therefore create an hierarchical set of contrasts that encapsulates the linear relationship that actually we would like to test. We therefore predict that:

$$\begin{array}{l} \text{and} \\ \text{and} \end{array} \quad \begin{array}{l} A < (B + C + D)/3 \\ B < (C + D)/2 \\ C < D \end{array}$$

We enter the data as columns in the usual manner (Fig. (i)a). However, we now need to enter some coefficients (weightings) for the first contrast in the line marked ‘Contrast’, and this requires a little explanation.

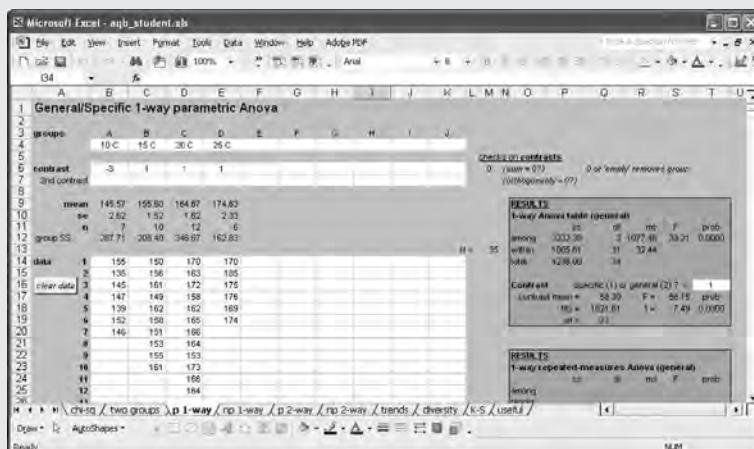
Basically, the contrast coefficients need to be integers that sum to zero, as indicated by the ‘checks on contrasts’ (cell M6) to the right of the coefficient input box. In the first contrast (for A), the prediction is:

$$\begin{array}{l} \text{which is} \end{array} \quad \begin{array}{l} A < (B + C + D)/3 \\ 3A < B + C + D \end{array}$$

which, rearranged to express the inequality as ‘0<’, becomes:

$$0 < -3A + B + C + D$$

therefore the coefficients are -3, +1, +1 and +1 on the groups A, B, C and D, respectively. Placing these in the appropriate boxes brings up



(a)

### One-way ANOVA table

	ss	df	ms	F/t	prob
Among	3232.39	3	1077.46	33.21	< 0.001
A < B + C + D	1821.61	1	1821.61	t = 7.49	< 0.001
B < C + D	997.33	1	997.33	t = 5.54	< 0.001
C < D	413.44	1	413.44	t = 3.57	< 0.01
Within	1005.61	31	32.44		
Total	4238.00	34			

(b)

Figure (i) (a) Data entry and results output in AQB for a parametric one-way ANOVA using specific contrasts, (b) the resulting ANOVA and specific contrasts.

the result for the contrast ('RESULTS' box in Fig. (i)a). The ANOVA in the box shows a significant outcome ( $F_{3,34} = 33.21$ ,  $p < 0.0001$ ) for the general test among all the groups. Here we are not merely making a contrast, but are also predicting its *direction*, so we choose the specific version of the contrast test by entering '1' in the small white box to the lower right of the 'RESULTS' box. When we do this, a  $t$ -value and associated probability for the specific test is then generated, showing a significant fit to our directional prediction ( $t_{33} = 7.49$ ,  $p < 0.0001$ ).

Then we go through each of the remaining contrasts, entering the coefficients and obtaining the general ANOVA table and specific-contrast  $t$ -value (Fig. (i)b). By the same logic as above, the coefficients are 0,-2,1,1 (for  $2B < C + D$ ) and 0,0,-1,1 (for  $C < D$ ).

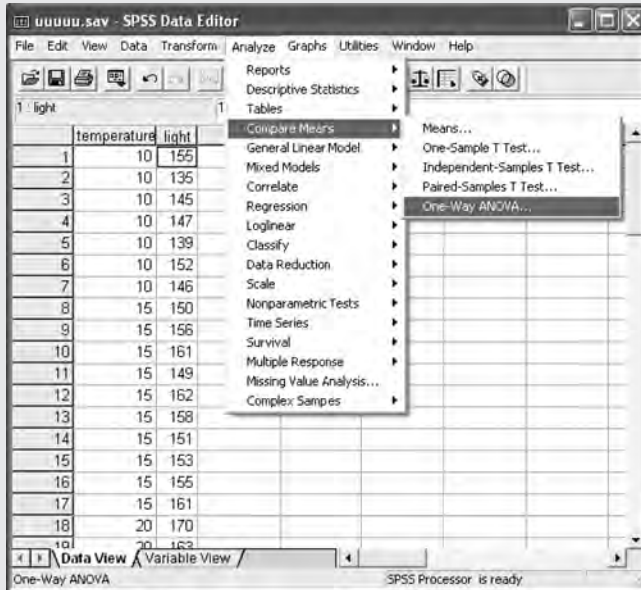
### Testing for a specific difference in the means of two or more groups using a parametric one-way ANOVA in SPSS

In SPSS, we enter the data in the usual manner, with the data in a single column and the group codes in another column. Then click 'Analyze', 'Compare Means' and 'One-Way ANOVA' (Fig. (ii)a). In the dialogue box (Fig. (ii)b), transfer the data variable (*light*) into the 'Dependent

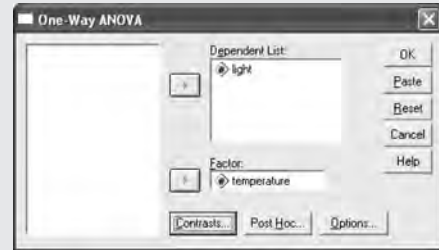
List' box and the grouping factor (*temperature*) into the 'Factor' box. Then click the 'Contrasts' button. In the resulting dialogue box (Fig. (ii)c) we could enter the same contrast coefficients as we did in AQB. In SPSS we are not allowed to enter the contrasts ourselves, but must pick from a predetermined set – in our case the pattern of coefficients corresponds to one called 'Helmert' contrasts (click on the 'Help' button to find out all the various possibilities).

However, we shall take advantage of the equally spaced treatments, so check the 'Polynomial' box and click 'Continue' and 'OK'. The output (Fig. (ii)d) tests for linearity of the response to the evenly spaced factor levels using one degree of freedom [1 d.f.] (for the slope of the line), with the remaining two degrees of freedom ('Deviation') reflecting deviations from linearity. Since there are unequal sample sizes for each treatment, we need the weighted version of the linear term. The linear term is very highly significant, whereas the remainder is barely significant. This means that the vast majority of the treatment effect is accounted for within the linear model.

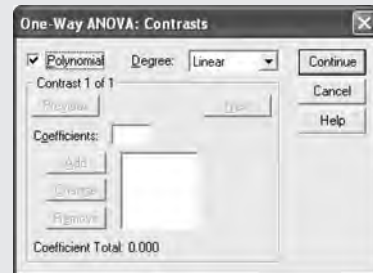
The conclusion in both cases is that there is good evidence that temperature increases light absorbance, and hence also the rate of the biochemical reaction. Furthermore (from the analysis in SPSS), this effect seems to be linear, at least within the range of temperatures studied here.



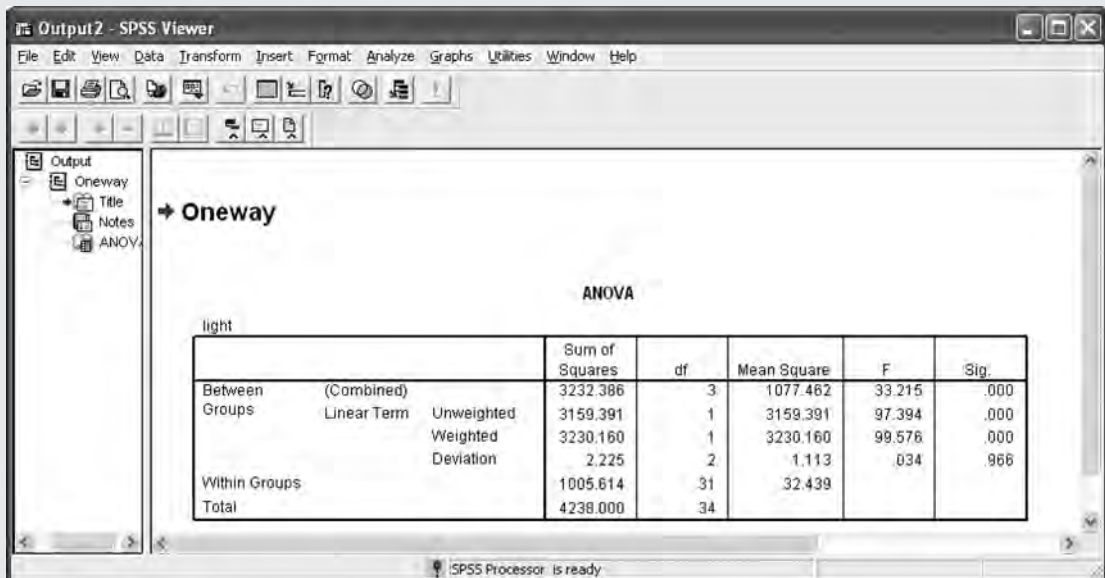
(a)



(b)



(c)



(d)

Figure (ii) (a) Data input and test selection for a specific one-way ANOVA in SPSS, (b) the dialogue box for inputting variables for the analysis, (c) selecting the polynomial linear fit option (see above), (d) output from the analysis.

### Box 3.5c Mean ranks (medians): a *general* non-parametric one-way ANOVA for two or more groups (Kruskal-Wallis test)

An experimenter measured how much alcohol students in all-female (F), mixed (MF), and all-male (M) halls of residence drank in one month. The data for the variable being analysed (*Alcohol*, the units of alcohol consumed in one month by a student) departs significantly from a normal distribution (Shapiro-Wilk = 0.96, d.f. = 300,  $p < 0.001$ ), so we can opt to use a non-parametric test. The grouping factor, *Hall type*, has three levels. The experimenter's prediction was a general one; she merely asked whether these types of hall of residence differed in their alcohol consumption per student, so that the expectation was that  $F \neq MF \neq M$ .

**A general non-parametric one-way ANOVA in AQB**  
In AQB, choose the 'np 1-way' worksheet and enter the data in groups, one group per column (Fig. (i)). This option will perform a Kruskal-Wallis non-parametric one-way ANOVA. Clicking the 'go' button in the 'RESULTS' box to the right leads to the outcome being displayed in the box.  $H$  is the Kruskal-Wallis test statistic, which here has two degrees of freedom (the number of groups minus one), so there is good evidence of a difference between the halls in alcohol consumption ( $H = 143.7$ , d.f. = 2,  $p < 0.001$ ).

**A general non-parametric one-way ANOVA in SPSS**  
In SPSS (Fig. (ii)a) the alcohol measures are entered in one column (*Alcohol*); the types of hall are coded as 1–3 in another column (*Hall type*). Again, the code needs to be a number, otherwise you will find it impossible to select *Hall type* as a factor in the analysis. Choose the Kruskal-Wallis test, the same as in AQB, by clicking 'Analyze' followed by 'Non-parametric Tests' and 'K Independent Samples' (Fig. (ii)a). The dialogue box in Fig. (ii)b will then appear. Transfer *Alcohol* into the 'Test Variable List' box, and *Hall type* into the 'Grouping Variable' box: we need to define the range of code values for

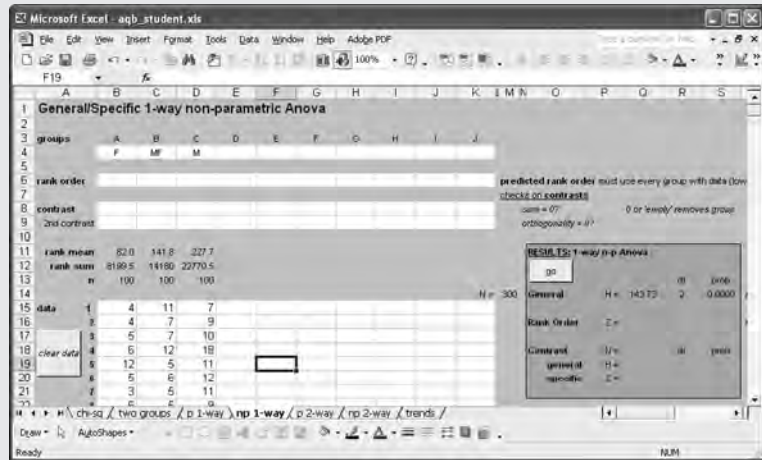


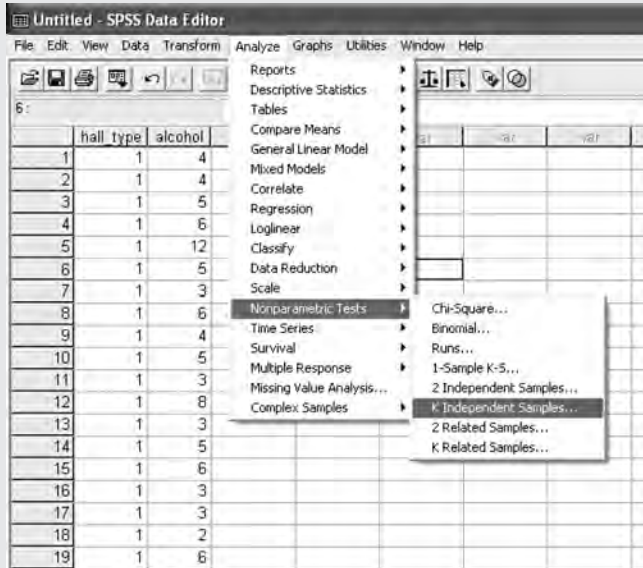
Figure (i) Data entry and results output for a general non-parametric one-way ANOVA in AQB.

*Hall type* as well, by clicking the 'Define Range' button and entering '1' and '3' as the minimum and maximum values. If your sample size is low (no group > 25) then click the 'Options' button and check 'Exact' for the exact test option (Fig. (ii)c). This is because, with large samples, an accurate approximation to the true probability is estimated, but this becomes inaccurate at low sample sizes. In our case, all sample sizes are large, and so the approximate (standard) test is appropriate (AQB can only cope with the approximate test). Click 'Continue' and make sure the 'Kruskal-Wallis H' box (Fig. (ii)b) is checked, then click 'OK'. The resulting Output is shown in Fig. (ii)d. Note that SPSS calls the Kruskal-Wallis test statistic  $\chi^2$  instead of  $H$ ; this is because as sample size increases, the distribution of  $H$  comes to approximate that of  $\chi^2$  – they thus refer to the same thing here (hence the value is the same as that for  $H$  in AQB).

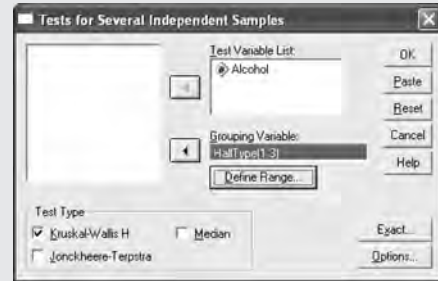
The conclusion therefore is that there is good evidence of differences in alcohol consumption between halls ( $H$  (or  $\chi^2$ ) = 143.7, d.f. = 2,  $p < 0.001$ ).

#### Post-hoc multiple comparisons

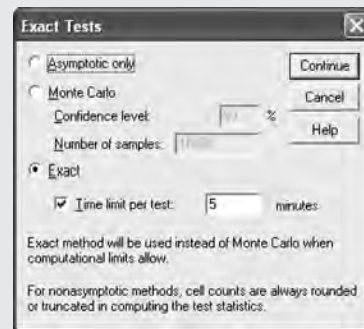
There is no generally accepted *post hoc* multiple comparisons test for non-parametric analyses of variance, and few if any are implemented in the standard statistical packages. However, they do exist.



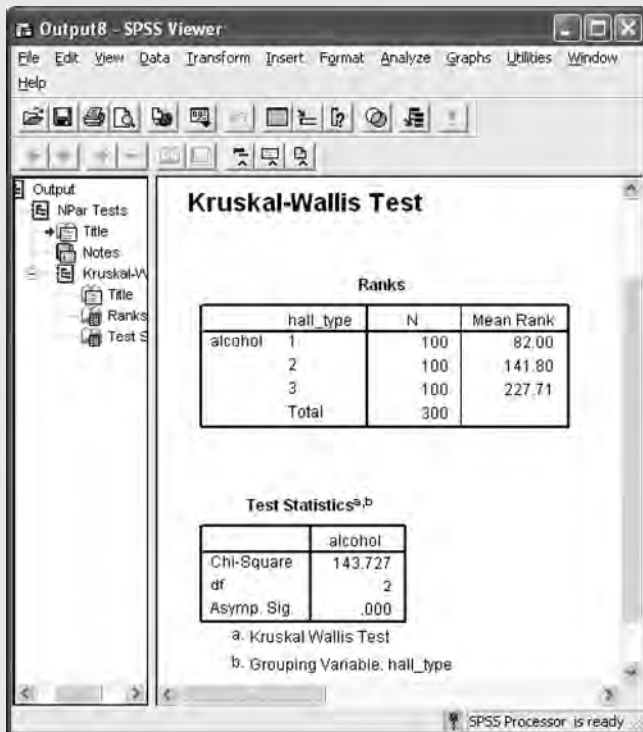
(a)



(b)



(c)



(d)

Figure (ii) (a) Data entry and test selection for a non-parametric one-way ANOVA in SPSS, (b) entering and defining variables for the analysis, (c) selecting the exact test option, (d) output of results.

### Box 3.5d Mean ranks (medians): a *specific* non-parametric one-way ANOVA

Again, we can use non-parametric analysis of variance to test specific as well as general predictions of difference. To illustrate this, we shall use a similar dataset to that of Box 3.5c, on the amount of beer drunk by students at all-female, mixed or all-male halls of residence during the course of one week. Once again, the residuals are highly non-normal (Shapiro-Wilk = 0.94, d.f. = 300,  $p < 0.001$ ), and so a non-parametric test is again appropriate.

Since it is well known that young men are more likely to be heavy drinkers, and to egg one another on in a drinking environment, than women, an obvious specific prediction is that all-female halls (A) will have lower alcohol consumption than mixed-sex halls (B), which, in turn, will have lower consumption than all-male halls (C).

#### A specific non-parametric one-way analysis of variance in AQB

In AQB we can test this prediction explicitly using a predicted rank order. We can also test it indirectly using contrasts.

##### Testing rank order

In the 'np 1-way' worksheet again, enter the groups and their predicted rank order in the 'groups' and 'rank order' boxes (Fig. (i)) and click 'go' in the 'RESULTS' box to the right. You can see that the probability is quite a bit lower for the test for a specific rank order (0.0015) than for the general test for differences from Box 3.5c (0.005). This demonstrates that the specific test, if the prediction is supported by the data, is much more powerful than the equivalent general test.

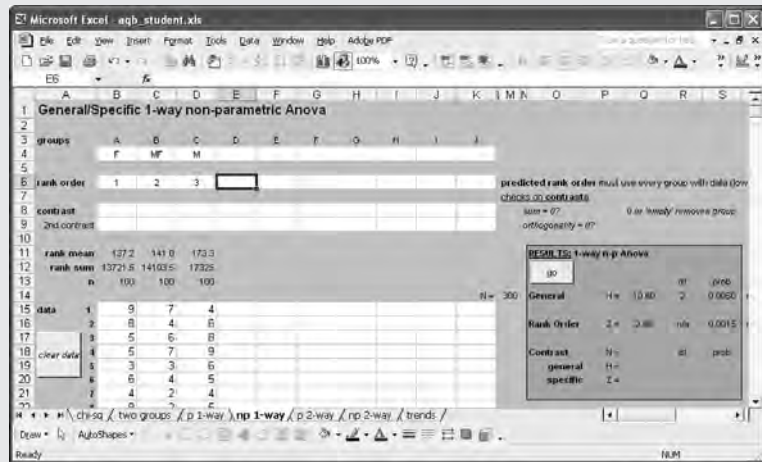


Figure (i) Data entry, rank order prediction and results output for a non-parametric specific one-way ANOVA in AQB.

##### Using contrasts

We can also test the specific prediction using contrasts, since we can (see Box 3.5b), with our two degrees of freedom, predict that:

$$A < (B + C)/2, \text{ i.e. } 2A < (B + C),$$

$$\text{or } 0 < -2(A) + 1(B) + 1(C)$$

and  $B < C, \text{ i.e. } 0 < -1(B) + 1(C)$

Entering the first set of contrast coefficients into the 'contrasts' box in AQB (Fig. (ii)) and clicking 'Go' gives the outcome in the 'RESULTS' box. We can see that the contrast is significant ( $z = 1.89, p < 0.05$ ). Repeating this for the second set of coefficients again shows that the contrast is in the predicted direction and is significant ( $z = 2.57, p < 0.01$ ).

#### A specific non-parametric one-way analysis of variance in SPSS

As before in SPSS (see Box 3.5c, Fig. (ii)a) we click 'Analyze', 'Non-parametric Test' and 'K Independent Samples'. This time, however, we check the test called 'Jonckheere-Terpstra' in the resulting dialogue box (Fig. (iii)a), which is a test



for a predicted rank order. It can *only* test for gradually increasing or decreasing mean ranks across the levels of the factor, and so our groups need to be *already entered in the order of the expected mean ranks at the outset*. If the sample sizes are small, or we are making a specific prediction (as here), then we can click the ‘Exact’ button and ask for either Exact (only for small sample sizes  $< 25$  in each group) or Monte Carlo (larger sample sizes – as here) (Fig. (iii)b).

The Standardised Jonckheere-Terpstra statistic (‘Std J-T Statistic’ in the output) is the same thing as the  $z$ -value of AQB. There is a two-tailed (decreasing or increasing mean ranks) and a one-tailed (one of the two possibilities only) probability. Because we are asking a specific rank order here, then the one-tailed probability is the correct value if the mean ranks fall into the expected pattern. Although in the text (p. 65), we’ve pointed

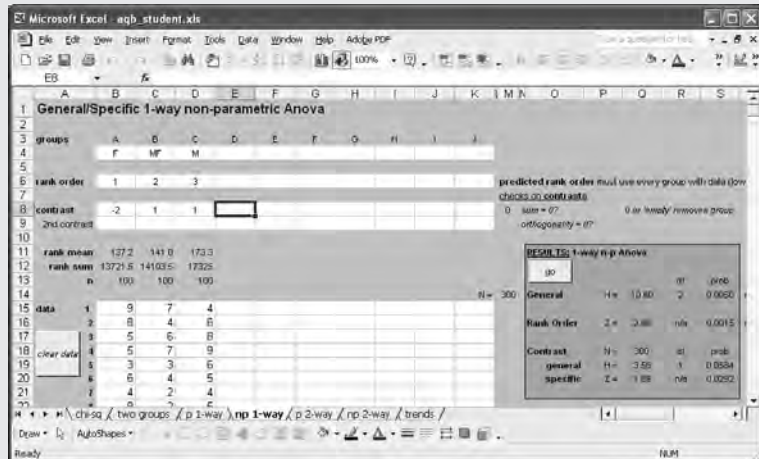
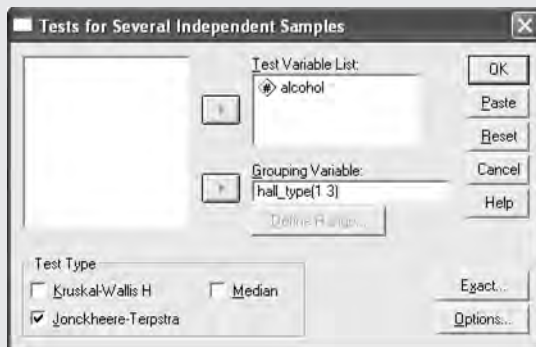


Figure (ii) Use of contrasts and results output for a non-parametric specific one-way ANOVA in AQB.

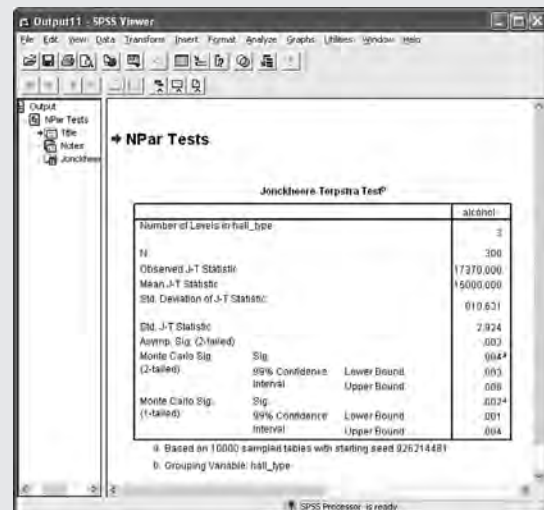
out that specific tests are usually more than just one- versus two-tailed tests, here there are only two possibilities because of rigid requirement of ordering the groups beforehand.

In both cases, we conclude that the data on alcohol consumption fall into the expected pattern of all-female  $<$  mixed  $<$  all-male halls ( $z = 2.96, p < 0.002$  [AQB] OR  $z = 2.92, p < 0.002$  [SPSS]).



(a)

Figure (iii) (a) Entering and defining variables, and selecting the appropriate test for a specific non-parametric analysis of variance in SPSS, (b) the output of results.



(b)

### Box 3.5e A parametric repeated-measures ANOVA

A molecular biologist wanted to study the impact of two different selection regimes on the production of oxygen free radicals. He therefore distributed one culture of liver cells (which he obtained from various sources) to 100 of his colleagues – a different culture going to each – and asked them to divide each into three selection treatments (randomly selected control, selected for high growth rate, and selected for low growth rate). After 10 rounds of cell division under these conditions, his colleagues sent back their cultures and the molecular biologist measured the average production of free radicals of 500 cells (a measure of metabolic activity) of each treatment from each culture.

This is a one-way repeated-measures design because cells from each culture are exposed to all treatments, and there is only one sort of treatment (selection type). Cultures may differ in lots of ways from each other, including their natural level of free radicals; ‘culture’ is therefore a block effect that needs to be allowed for in the analysis in order to see the impact of the treatments. The researcher only obtained a single measurement of the dependent variable (level of free radicals) from each culture and each treatment.

First we frame the prediction: that the selection regimes change the level of free radicals from the randomly selected control condition ( $H_1$ ), without specifying the direction of change (this is therefore a general hypothesis). As with the usual parametric one-way ANOVA (see Box 3.5b), it is not possible to test a specific version of this hypothesis, although, as before in AQB, one can approach it using contrasts. The null hypothesis is that selection does not change the level of free radicals ( $H_0$ ).

#### Repeated-measures ANOVA using AQB

Select the ‘p 1-way’ sheet in AQB and enter the measurement for each treatment as in Fig. (i), with the data for each culture on a single row. Therefore there will be equal numbers of values in each column (treatment), and if this is so, the analysis will automatically appear in the RESULTS box labelled ‘1-way repeated-measures ANOVA’. There are two  $F$ -values. The lower one is a test for the significance of differences among the ‘blocks’, i.e. the repeated-measures factor. Here it is huge and highly significant ( $F = 2880.9$ , d.f. = 99,198,  $p < 0.001$ ). The overall mean values for each of the columns are quite similar (about 71–72) with a fairly high variation (ca. 4.5). However, once this block effect is allowed for, the treatment effect (the upper  $F$ -value) is also significant ( $F = 30.1$ , d.f. = 2,198,  $p < 0.001$ ).

In the RESULTS box for a simple ANOVA (i.e. *not* repeated measures), which also appears automatically, you can see that if there had been no repeated-measures factor, then the large

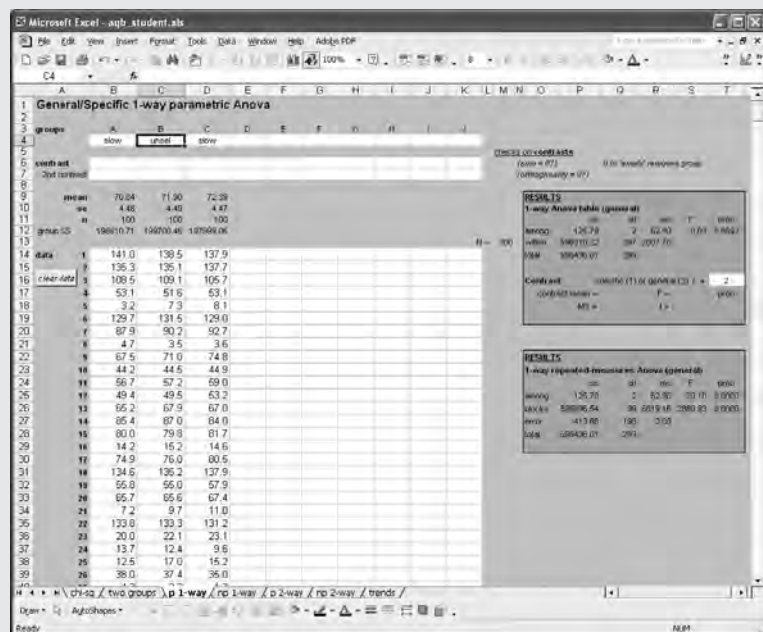
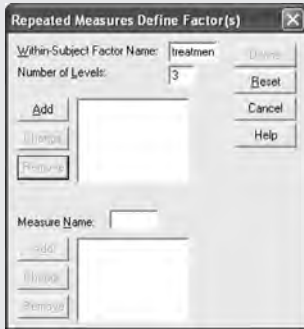
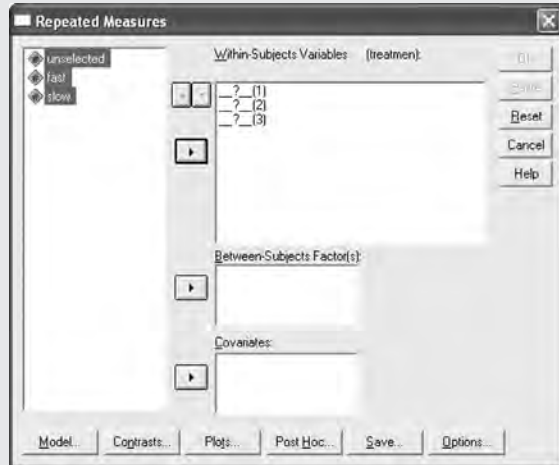


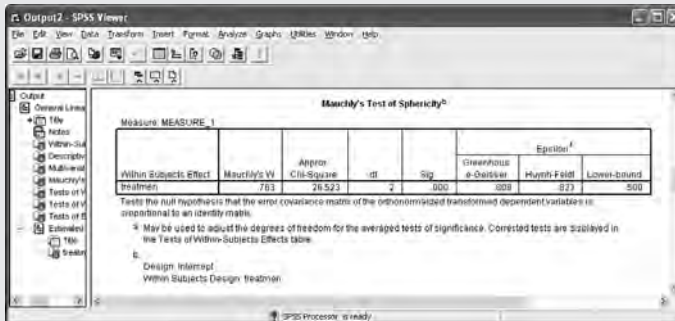
Figure (i) A general parametric repeated-measures ANOVA in AQB.



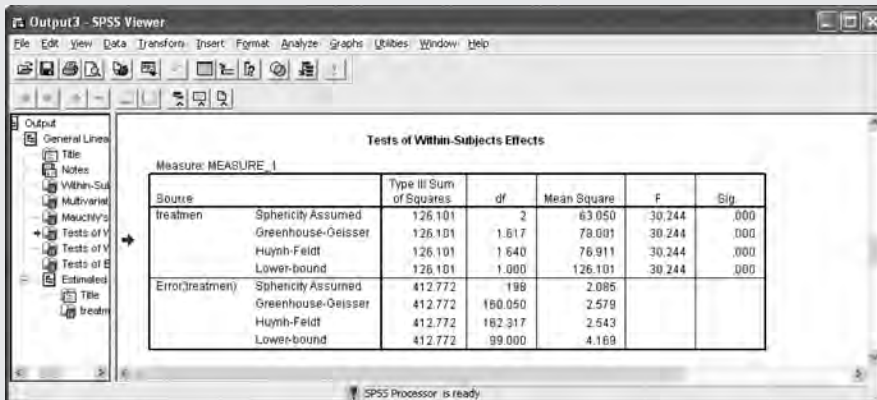
(a)



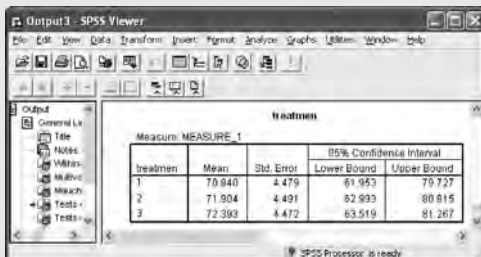
(b)



(c)



(d)



(e)

Figure (ii) A general parametric repeated-measures ANOVA in SPSS. (a) Dialogue box for defining the repeated-measures factor, (b) dialogue box for assigning the data to the defined repeated-measures factor, (c) multivariate test for normality, (d) result for non-normal data, (e) mean values  $\pm$  s.e. for the treatments.

variation within the columns obscures any treatment effect.

### Repeated measures ANOVA using SPSS

SPSS adopts very general approaches to analysing different experimental/observational designs, and the price of generality can be complexity. This is nowhere more evident than in the repeated-measures ANOVA, which follows a multivariate mode of analysis because it uses the values for each of the factor levels as if it were a different variable. The individuals subjected to the repeated measures are assumed by default to correspond to each line of the data.

Click on 'Analyze', 'General Linear Model' and then 'Repeated Measures' to obtain the dialogue box in Fig. (ii)a, where we can define the repeated-measures factor by naming it (here *treatmen* – since only 8 letters are allowed) and assigning the number of levels (here 3, since there are 3 levels of the factor '*treatmen*' – the three selection regimes). Then click 'Add' and 'Define', bringing up the second dialogue box in Fig. (ii)b. Here assign the columns of data to these pre-defined levels of the repeated-measures factor, either one by one, or (if they are in the correct order) all in one go (as in the figure). Clicking 'OK' then produces the output.

The first thing to check is whether the residuals conform to the assumption of normality. This is impossible unless you have very large sample sizes because of the way in which the test is done in SPSS. You have therefore to assume normality, and take extra care in interpreting the results.

The result is shown in Fig. (ii)d, and gives four different tests of the (general) null hypothesis ( $H_0$ ) that none of the selection treatments have any effect. The first line ('sphericity assumed') is for data in groups with homogeneous variances (see Box 3.1b), and the subsequent three lines are different ways of adjusting for inhomogeneity. Here, reassuringly, they all give the same answer, which is that there is a highly significant effect of treatment. The way in which they adjust for non-homogeneity is complex, and need not concern us here.

An additional point here is to notice the much greater power of the repeated-measures design in cases like this where there is large variation among individuals. The mean values of the treatments (Fig. (ii)e) show that there is large overlap in the standard errors, which would normally imply no significant differences between the means. By removing inter-individual variation, however, the impact of the treatments becomes evident over and above variation between individuals.

## Box 3.5f A non-parametric repeated-measures ANOVA

To illustrate a non-parametric approach, we shall use the same example as in the parametric case because of the significant non-normality of the residuals, demonstrated by the SPSS analysis in Box 3.5e.

### A non-parametric repeated-measures ANOVA using AQB

In this case we will make a specific prediction: we predict the 'unselected' line to have the inter-

mediate level of free radicals, so that slow < unselected < fast ( $H_1$ ). The null hypothesis ( $H_0$ ) is that the mean ranks will not follow this pattern.

In AQB (Fig. (i)) the data for the treatments are in columns, with the data for each individual on a single row. The hypothesis is placed as a predicted rank order in row 6, as shown, with the smallest predicted rank mean given the rank of '1'. Clicking 'Go' in the RESULTS table for repeated-measures produces the result. Both a

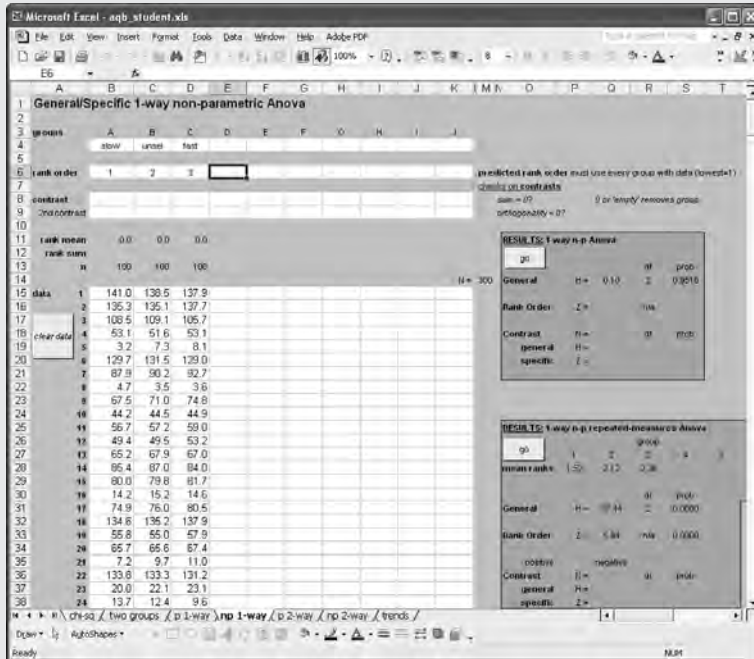
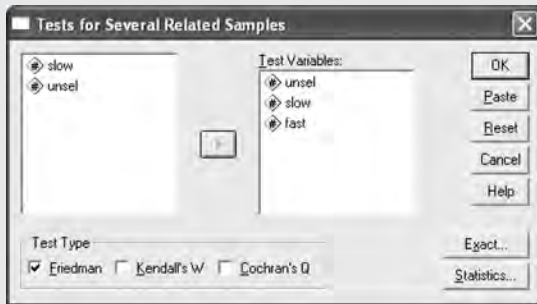


Figure (i) A specific non-parametric repeated-measures ANOVA in AQB.

general test ( $H$ ) and the specific one ( $z$ ) are produced. It is clear that the result is highly significant ( $z = 5.94$ ,  $p \ll 0.001$ ), indicating agreement with the prediction.

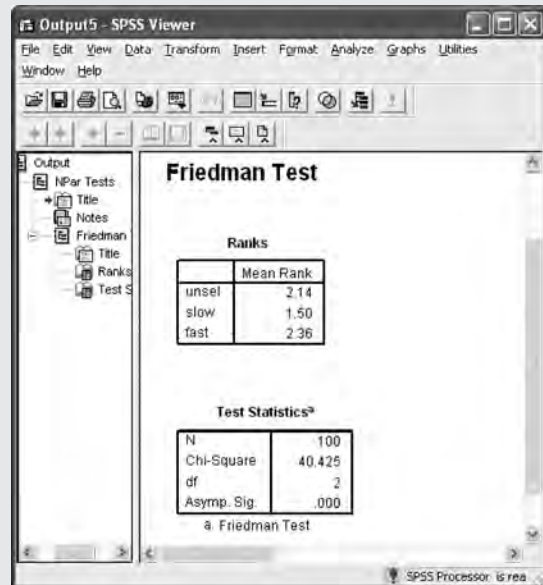
**A non-parametric repeated-measures ANOVA using SPSS (Friedman test)**

Here we can test only the general prediction that the mean ranks will not differ. Click on 'Analyze', 'Non-parametric Tests' and then 'K-related Samples'. Add the three variables of the treatments to the resulting dialogue (Fig. (ii)a), and the Output appears (Fig. (ii)b). As before, the mean ranks are quite different, and the test statistic shows that this is a highly significant difference.



(a)

Figure (ii) A specific non-parametric repeated-measures ANOVA in SPSS (Friedman test). (a) Adding the variables, (b) results.



(b)

**Covariates.** Sometimes it is useful to control for other, nuisance, factors that are measured on a constant-interval scale within an analysis of variance. For instance, an analysis of differences in mating success between territory-owning and non-territory-owning male wood mice (*Apodemus sylvaticus*) might want to control for body size to rule out an effect of territory ownership arising simply because owners tend to be bigger. In a parametric analysis of variance, body size could be incorporated as a *covariate*. The analysis would then reveal the independent effects of territory ownership and body size. Most major statistical packages allow covariates to be included in analyses of variance. Box 3.6 shows how to do this in SPSS.

### Box 3.6 Analysis of covariance – taking a covariate into account

A physiologist was interested in the impact of three foods on the basal metabolic rate (BMR) of rats, and conducted an experiment in his laboratory where rats in individual cages were randomly assigned to one of the three foods (A, B and C), and were fed this food over one week. He then measured the BMR of each rat three times, taking the average as the value for each rat. He suspected, however, that body size might also influence BMR, and might obscure the result of the experiment unless controlled for. Thus he also recorded the weight of each rat when he measured its BMR.

The dependent variable for analysis is BMR, measured on a constant-interval scale. The question concerns differences in the mean value of BMR between the three levels of the factor food type (*group* in Fig. (i)a, after having taken into account the effect of a *covariate*, body size (*size* in Fig. (i)a). After saving the residuals of BMR from the analysis of covariance described below (i.e. after fitting *group* and *size*), they were tested for normality in the manner described in Box 3.1: they conformed to the normal distribution (Shapiro-Wilk = 0.97, d.f. = 60, ns).

Frame the prediction. In this case the physiologist did not have any preconceptions about what to expect, and therefore just tested the general prediction that the mean values differed.

The analysis is called an analysis of covariance (ANCOVA), and this example is the simplest possible kind – a one-way analysis of variance

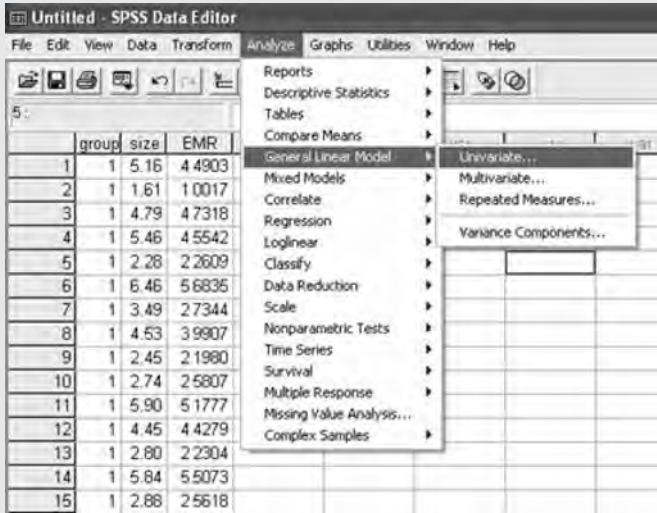
with a single covariate. ANCOVA can be very complex, involving several factors and covariates, and hence it cannot be done using the simple AQB package. Thus we shall show how to do the analysis in SPSS.

#### Analysis of covariance in SPSS

In SPSS the data are in the usual column format (Fig. (i)a), with a column for *BMR*, one for *size*, and one for the factor *group*. Click on ‘Analyze’ and ‘General Linear Model’ and ‘Univariate’ to bring up the dialogue box in Fig. (i)b. Enter the dependent variable for analysis (here, *BMR*) into the ‘Dependent Variable’ box, and the factor (here, ‘*group*’) into the ‘Fixed Factor(s)’ box.

Before doing the ANCOVA, let’s first see what the result is *without* taking the covariate of body size into account. Click the ‘Options’ button and transfer *group* into the ‘Display Means for’ box in Fig. (i)c and tick the ‘Descriptive Statistics’, and ‘Homogeneity tests’ options. When the analysis runs, we obtain the mean values of *BMR* for the three levels of *group* (Fig. (i)d) – they are clearly different, as the highly significant *F*-value of the ANOVA demonstrates (Fig. (i)e).

Now, let’s perform the ANCOVA that controls for body size in comparing the mean values of *BMR* in the different *group* treatments. Add *size* to the ‘Covariate(s)’ box (Fig. (i)f) and click ‘OK’ (leaving all other options as before). Now the result is completely different (Fig. (i)g)! There is a



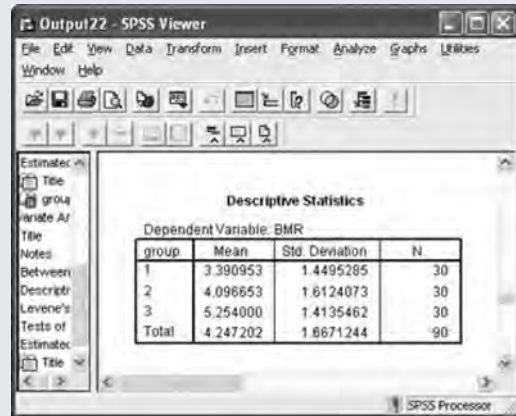
(a)



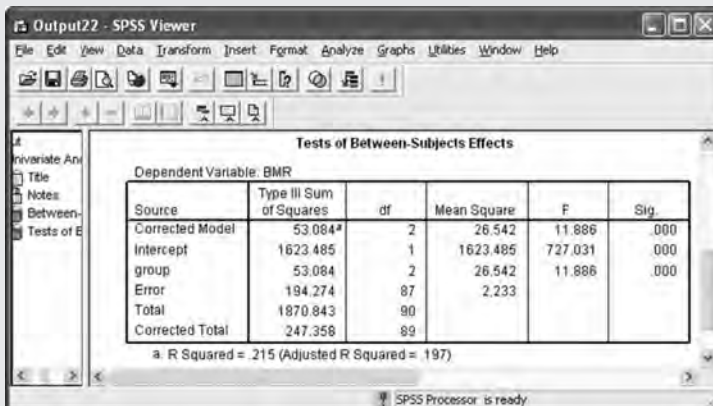
(b)



(c)



(d)



(e)

Figure (i)(a–h) Steps on the way to doing an ANCOVA. (a) Data entry and menu choice, (b) entering the variables for a simple ANOVA, (c) options to choose, (d) mean values of the groups, (e) result of the simple ANOVA, (f) entering the variables for an ANCOVA, (g) result of the ANCOVA, (h) estimated group means, allowing for the covariate.

Figure (i) cont'd



(f)

Tests of Between-Subjects Effects

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	239.780 <sup>a</sup>	3	79.927	907.095	.000
Intercept	1.472	1	1.472	16.711	.000
size	186.696	1	186.696	2118.829	.000
group	.180	2	.090	1.024	.364
Error	7.578	86	.088		
Total	1870.843	90			
Corrected Total	247.358	89			

<sup>a</sup>. R Squared = .969 (Adjusted R Squared = .968)

(g)

Estimated Marginal Means

group	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	4.242 <sup>a</sup>	.057	4.128	4.356
2	4.193 <sup>a</sup>	.054	4.085	4.300
3	4.307 <sup>a</sup>	.058	4.192	4.422

<sup>a</sup>. Covariates appearing in the model are evaluated at the following values: size = 4.7673.

(h)

very strong effect of body size on *BMR* ( $F_{1,86} = 2118.8$ ,  $p < 0.001$ ), and, once allowed for, there is *no* significant difference in *BMR* between the mean values of the food types ( $F_{2,86} = 1.02$ , ns). The 'Estimated Marginal Means' from the output (Fig. (i)h) show the mean values for the three levels of *group* estimated once size has been taken into account, and it is obvious that they hardly differ at all from one another.

Sometimes, as in the body size example here, covariates are just nuisance factors that need to be taken into account to avoid bias in the outcome of analyses. In other cases, however, we may be interested in the effect of the covariate itself. The routine in SPSS above can deal with both situations and handle a number of factors and a number of covariates at the same time.

**1 × n chi-squared.** One-way analysis of variance uses a comparison of mean (or mean rank) values of individual data values to arrive at a test statistic. Where data values are counts, however, we could, as in the two-group case (*see* Box 3.2) perform a chi-squared test on the *totals* for each group. We then have what is known as a 1 × n chi-squared analysis, where *n* is the number of groups. The



### Box 3.7 A test comparing counts classified into two or more groups ( $1 \times n \chi^2$ test)

A clinical microbiologist was assaying the effect of four antibiotics on a bacterial culture. To see whether the antibiotics differed in their ability to kill the bacterium, he counted the number of cultures on which clear plaques appeared after drop treatment with each one, and performed a  $\chi^2$  test, assuming equal expected values across the four antibiotics. The data consist of the total numbers of cultures. The nominal grouping factor is *treatment*, with four levels (the different antibiotics).

Frame the prediction, here that antibiotics differ in their ability to kill the bacterium, and hence there will be differences in the numbers of plaques between antibiotics ( $H_1$ ). The null hypothesis is that all the antibiotics are equally effective, and hence there will be no differences ( $H_0$ ).

#### Calculating a $1 \times n \chi^2$ using AQB

Under the 'chi-sq' worksheet in AQB, the counts for each treatment are entered as in Fig. (i).

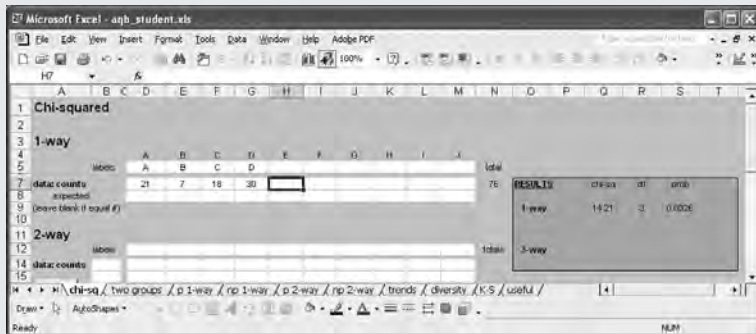


Figure (i) Performing a  $1 \times n \chi^2$  test in AQB.

Leaving the 'expected' values empty is the default assumption of equal expected numbers. The result appears automatically in the 'RESULTS' box to the right, showing, in this case, a significant difference between antibiotics ( $\chi^2 = 14.21$ , d.f. = 3,  $p < 0.01$ ).

#### Calculating a $1 \times n \chi^2$ using SPSS

To test the total counts directly in SPSS, we once again use the weighting procedure outlined in Box 3.2, but (here) for four groups instead of two. On doing this, the same result as that above will appear in the output.

two-group ( $1 \times 2$ ) chi-squared test earlier is just one form of this, and the calculation of expected values and the  $\chi^2$  test statistic are exactly the same as for the two-group case. Box 3.7 shows how to do these in AQB and SPSS.

#### Tests for differences in relation to two levels of grouping

In all the above difference tests, we were concerned with differences within a single level of grouping, e.g. between groupings based on seed colour. However many groups of seed colour we had (red, yellow, green, orange, blue, etc.) we would still be dealing only with seed colour and thus with one level of grouping. But we can easily envisage situations in which we would be interested in more than one level of grouping. For example, we might want to know not only whether chicks peck at some colours of grain more than others, but also whether pecking in males is different from that in females. More interestingly still, we might want to know whether the sex of chicks affects the *difference* in pecking at the various colours of grain. Is the difference stronger in one sex? Is it in the same direction in both sexes? Here we have *two* levels of groupings: seed colour

and sex. In the examples that follow, we shall look at analyses that cater for two levels of grouping with two groups in each.

**$2 \times 2$  chi-squared.** If we have data in the form of counts, we can again use chi-squared, but the expected values in a  $2 \times 2$  analysis (and in any other  $n \times n$  chi-squared analysis) are calculated in a different way from those in a  $1 \times n$  analysis. Instead of taking equal values, or values determined on the basis of some a priori expectation, the rows and columns of the  $2 \times 2$  (or  $n \times n$ ) table are totalled and the grand total calculated, and the respective expected values for each cell of the table are calculated as (row total  $\times$  column total)/grand total. Box 3.8 shows the procedure for performing a  $2 \times 2$  chi-squared analysis in AQB and SPSS.

**$2 \times 2$  two-way analysis of variance.** As before, if we want to compare sets of individual data values, we can use analysis of variance but this time it is a two-way rather than a one-way analysis. In a  $2 \times 2$  two-way analysis of variance, the data are cast into four cells (two  $\times$  two groups, which can be cast as two rows and two

### Box 3.8 A test comparing counts in a $2 \times 2$ classification ( $2 \times 2$ $\chi^2$ test)

A count was made of the total number of men and women with stomach biopsies showing a presence or absence of cancerous cells. We have two grouping factors each with two levels: *sex* (men/women) and *cancer treatment* (with/without), and each person's result is classified into one of the four combinations.

The prediction is that the sex of the person and whether they are treated affects their chance of showing stomach cancer ( $H_1$ ); in a  $2 \times 2$  chi-squared this would be reflected as a significant non-independence between the two factors (*see below*). The null hypothesis is that sex and treatment have no effect on the presence of stomach cancer ( $H_0$ ).

#### A $2 \times 2$ $\chi^2$ test in AQB

Under the 'chi-sq' worksheet in AQB, simply enter the group labels and four totals into the cells of the '2-way' box, as in Fig. (i), and the outcome appears automatically in the lower, two-way, line of the 'RESULTS' box.

So, in the example here, there is a significant interaction between sex and treatment ( $\chi^2 = 18.65$ , d.f. = 1,  $p < 0.001$ ).

#### A $2 \times 2$ $\chi^2$ test in SPSS

In SPSS, we either have to enter the raw data of each person, row by row, or get the package to accept the totals rather than the raw data of each person by using the 'Weight Cases' option, as in Box 3.2. Following the procedure in Box 3.2,

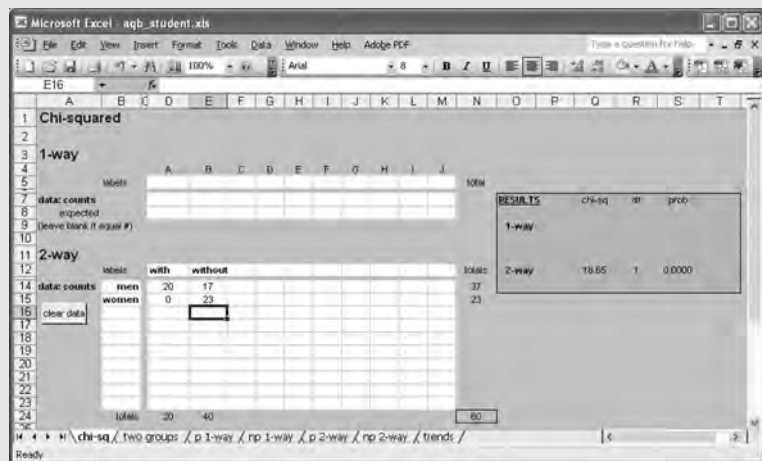
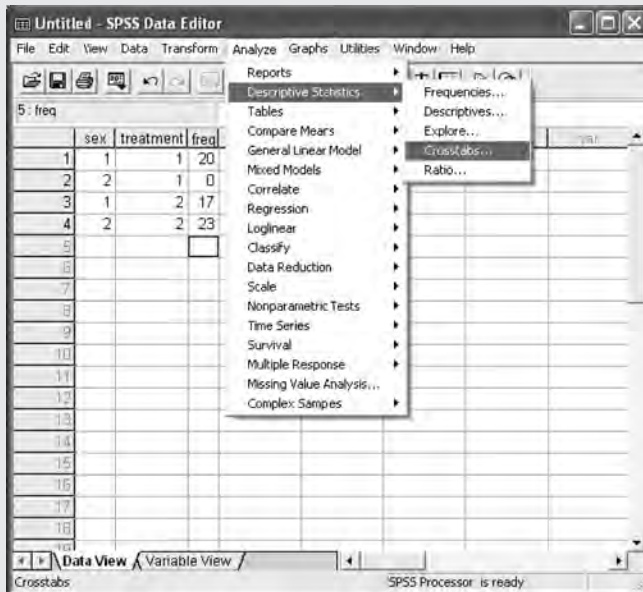
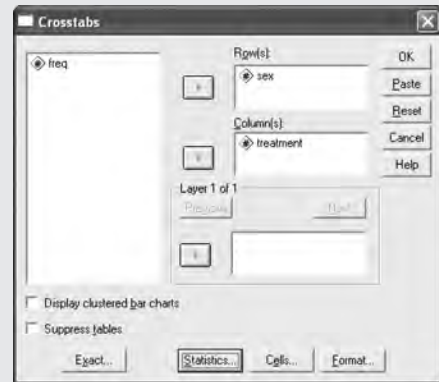


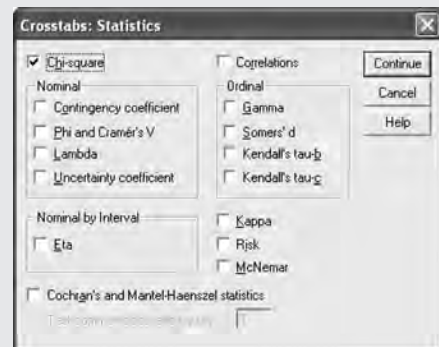
Figure (i) Data entry and results output for a  $2 \times 2$   $\chi^2$  test in AQB.



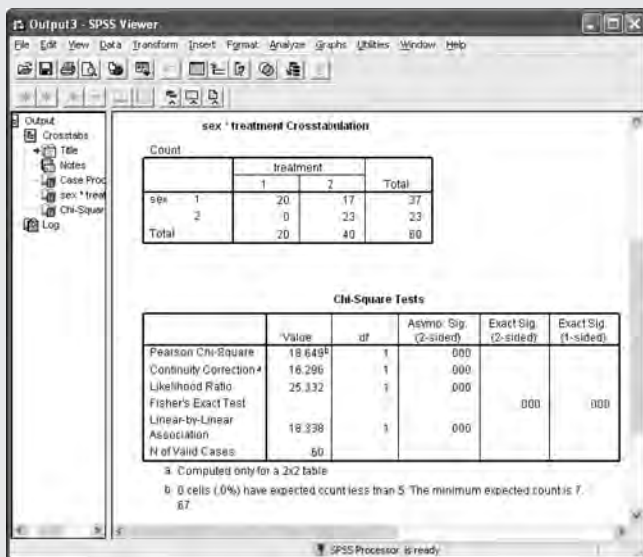
(a)



(b)



(c)

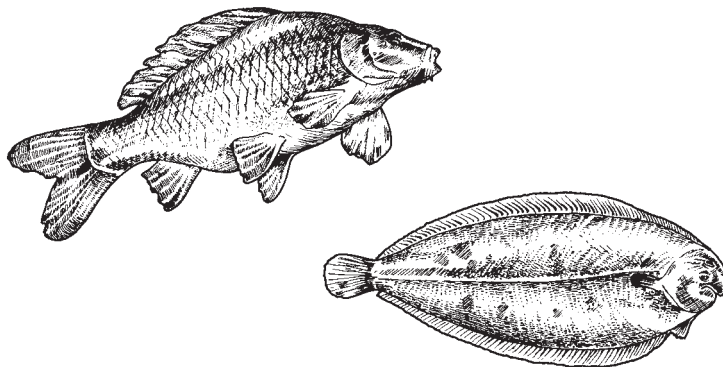


(d)

click 'Data' followed by 'Weight Cases' from the menus at the top of the screen, and select the 'Weight cases by' option, entering the column containing the data, in this case freq (see Box 3.2, Fig. (ii)c). Then click 'Analyze' followed by 'Descriptive Statistics' and 'Crosstabs' (Fig. (ii)a). Enter the two factors (*sex* and *treatment*) into the 'Row(s)' and 'Column(s)' boxes (Fig. (ii)b), click

on 'Statistics' and make sure that the 'Chi-square' box is checked (Fig. (ii)c). Click 'Continue' and 'OK' and the output appears as in Fig. (ii)d. The conclusion, once again, is that the two factors interact, i.e. the chance of getting stomach cancer differs between men and women ( $\chi^2 = 18.65$ , d.f. = 1,  $p < 0.001$ ).

columns of a table). If we wanted to do a  $3 \times 5$  two-way analysis the data would be cast into 15 cells, and so on for any combination of levels of grouping. Once again there are both parametric and non-parametric versions of the analysis (Boxes 3.9a–d). As usual, the parametric test assumes the data conform reasonably to normality. This assumption is, of course, relaxed for the non-parametric equivalent. However, *both* parametric and non-parametric tests assume that the data have the same variance within each cell (i.e. within each combination of levels of grouping) – another example of the distribution-free, but not assumption-free, nature of non-parametric tests. Both types of analysis compare the mean values of the columns within the classification and the mean values of the rows. In other words they compare means within each of the two levels of grouping. Comparisons between the column means or between the row means are known as the *main effects* and are distinguished from a second kind of comparison referred to as an *interaction*. (Note that an interaction can be calculated only if all cells contain more than one datum value.) If there is a significant interaction it means the two sets of samples at one level of grouping respond differently to differences in the second level of grouping. An example makes the distinction between main effects and interaction clear. Imagine our two levels of grouping are freshwater versus marine fish and male versus female, and the variable for comparison is growth rate. Freshwater versus marine can be the rows of the classification (*see table below*) and male versus female the columns. The analysis would be concerned with the following: (a) main effect 1: differences in row means (is there any difference in growth rate between freshwater and marine fish?), (b) main effect 2: differences in column means (is there any difference in growth rate between males and females?) and (c) any interaction between the levels of grouping (e.g. is the difference in growth rate between males and females greater in one environment than in the other?).



		Water type	
		Fresh	Marine
Sex	Male	a b c d mean = A	e f g h mean = B
	Female	i j k l mean = C	m n o p mean = D

As in the one-way analysis of variance, the non-parametric version can be used to test either general or specific predictions, but now about both the two main effects and any interactions. Making specific predictions is rather more involved in the two-way analysis because we need to be clear as to exactly what we are comparing and to calculate different coefficients for each specific comparison so predictions can be tested. Box 3.9c illustrates the procedure.

Although we have confined ourselves to a two-way analysis of variance here, the two-way model is only a particular case of a multifactor analysis of variance where there can be three, four or more levels of grouping. The principles underlying more complex analyses are the same, but as the number of levels of grouping increases it becomes more and more difficult to interpret the proliferating interactions terms. Moreover, the more levels of grouping that are included, the slimmer the chance they are all truly independent. Full discussion of these analyses can be found in Sokal and Rohlf (1995). In parametric analyses, as in the one-way case, it is also possible to incorporate covariates to control for factors on a constant interval scale.

### Box 3.9a Mean values – a *general* parametric two-way ANOVA

An experiment aimed to detect the effects of food type and parasitic infection on an enzyme (liver alcohol dehydrogenase) in rats (*Rattus norvegicus*). The variable being analysed (*enzyme*) is a measure of the enzyme's rate of reaction, and there are two factors, each with two levels: *food* (peas/beans) and *parasitism* treatment (unparasitised, parasitised). The experimental design required each rat to be in one food group (peas or beans) and one parasitism treatment group (unparasitised or parasitised). Testing the residuals for normality (see Box 3.1) showed that the data could be assumed to be normal (Shapiro-Wilk = 0.99, d.f. = 160, ns).

Because of the two-way design, we have three predictions, one about the effects of *food*, one about the effects of *parasitism* (together called the main effects), and one about the *interaction* between food type and parasitism (e.g. is the effect of food type affected by whether or not the rat is also parasitised?)

#### Doing a general parametric two-way ANOVA in AQB

In AQB, select the 'p 2-way' worksheet and enter the data as usual in columns, making sure that they are in the correct ones. The advantages of

intuitive inputting of the data in columns for one-way designs are lost for two-way designs, and you must be careful.

It is often helpful to draw a diagram of the design, as in Fig. (i)a. This shows the two levels of *food* as the columns, and the two levels of *parasitism* as the rows. In the cells of this design in Fig. (i)a, we have written the column in which the data for that cell should be placed in AQB, and indicated on Fig. (i)b by arrows where the data should go. On row 3 of the worksheet are spaces for the factor labels – in our case, Factor A is *parasitism*, and factor B is *food*. There is space on the worksheet for four levels of Factor A and six levels of Factor B – these are indexed by number on rows 4 and 5 of the worksheet, and have spaces for the level labels on rows 6 and 7, respectively. Thus columns B–G are for the six possible levels of Factor B for level 1 of Factor A; columns H–M are for the six possible levels of Factor B for level 2 of Factor A; and so on up to level 4 of Factor A. Rows 9 and 10 of the worksheet are for the contrasts, as we have met before and will meet in Box 3.8b.

Then come five rows which automatically display (as you type in the data) the summary statistics of the data: rows 12, 13 and 14 record

		Food	
		Peas	Beans
Parasitism	Parasitised	Column_B	Column_C
	Unparasitised	Column_H	Column_I

(a)

(b)

(c) Note that the width of some of the columns has been reduced to fit data and result into a single frame. Note carefully how the two-way design is put into the column format.

Figure (i) (a) Experimental design, (b) data entry, and (c) results output for a general parametric two-way ANOVA using AQB.

the mean, s.e. and sample size of the data for each column (i.e. for each cell of the two-way design). Row 15 has four boxes, which show the mean value for each of the levels of Factor A (*parasitism* in our design, and hence only two are filled in, since we have only two levels for this factor). Finally, row 16 has six boxes for the mean values of the six possible levels of Factor B (but only two are filled in here, because *food* only has two levels).

There is room for a  $4 \times 6$  two-way design in the spreadsheet, but anything larger will need to be done in SPSS or another package. The result automatically appears in the RESULTS box (Fig. (i)c).

### Doing a general parametric two-way ANOVA in SPSS

In SPSS, the data are entered in their usual format, with all the enzyme activity values in a single column, with other columns coding the parasitism (1 = parasitised, 2 = unparasitised) or food type (1 = peas, 2 = beans) treatment groups (Fig. (ii)a). To carry out a two-way ANOVA, click on 'Analyze' then 'General Linear Model' and 'Univariate' (Fig. (ii)a). General linear models (or GLMs for short) are a broad family of parametric tests to which analysis of variance belongs (*see below*). Univariate is chosen because only one variable is being analysed (here, enzyme activity). Transfer the variable being analysed (*enzyme*) into the 'Dependent Variable' box, and the two treatment groups (*food*, *parasitism*) into the 'Fixed Factor(s)' box (Fig. (ii)b).

While the dialogue box offers a choice between 'Fixed Factor(s)' and 'Random Factor(s)', the distinction between the two is not hard and fast, and the same factors can potentially be treated as 'fixed' or 'random' in different tests (Dytham, 2003). The distinction really lies in the extent to which any significant difference emerging from a test can be generalised. Thus, for example, if *food* in our example is entered as a fixed factor, and we find a significant difference between the levels in whatever we are measuring (here, enzyme activity), this implies only that there is a difference in the effect of peas and beans *per se*. If they are entered as levels of a random factor, however, it implies differences in effect

between any two randomly selected foodstuffs. In other words, using them as a random factor implies peas and beans represent foodstuffs in general. The extent to which generalisation is possible depends critically on how the levels of the factor were selected – they really do have to be a random selection in order to be able to generalise. Sometimes a factor is random because it is a blocking factor where we have no idea what the nature of the variation might be (e.g. differences among individuals, rearing cages, fields used for fertiliser trials, or physiological preparations), but it must be allowed for in the analysis. The distinction is not offered by many, more basic, statistical packages, such as AQB. Whether a factor is fixed or random makes a big difference to the way the test statistic is calculated, and hence to the results, so the choice is important.

Once the variables are entered, click the 'Options' button (Fig. (ii)c) and transfer the interaction term (*food\*parasitism*) into the 'Display Means for' box (Fig. (ii)c) so that we get the means and standard errors for each of the combinations of the factors in the eventual output. Tick the 'Observed power' and 'Homogeneity tests' boxes too. Click 'Continue' then 'OK' and the output appears as in Fig. (ii)d. First, this gives the test for homogeneity of variances among groups (one of the main assumptions of parametric ANOVA), which shows no evidence of differences in variance among groups (i.e. Levene's test is not significant). Then we get the ANOVA itself, which is identical to the one provided by AQB except for some extra information – for example, there is a line for the 'intercept', because GLM models combine both differences and trends into a single framework of regression models and fit coefficients to express the quantitative contribution of variables in the analysis. In this case, in very simple format the model is:

$$\begin{aligned} &\text{Rate of enzyme activity (enzyme)} \\ &= (\text{Overall Mean}) + a(\text{food}) + b(\text{parasitism}) \\ &\quad + c(\text{food*parasitism}) \end{aligned}$$

where a, b and c are coefficients weighting the contribution of each variable and their interaction

Untitled - SPSS Data Editor

File Edit View Data Transform Analyze Graphs Utilities Window Help

Reports  
Descriptive Statistics  
Tables  
Compare Means  
General Linear Model  
Mixed Models  
Correlate  
Regression  
Loglinear  
Classify  
Data Reduction  
Scale  
Nonparametric Tests  
Time Series  
Survival  
Multiple Response  
Missing Value Analysis...  
Complex Samples

	food	parasitism	enzyme
1	1	1	9.16
2	1	1	16.70
3	1	1	11.69
4	1	1	12.10
5	1	1	7.49
6	1	1	10.22
7	1	1	13.49
8	1	1	13.75
9	1	1	10.98
10	1	1	5.86
11	1	1	11.40
12	1	1	10.08
13	1	1	12.78
14	1	1	5.02
15	1	1	10.83
16	1	1	12.98
17	1	1	9.68

(a)

Univariate

Dependent Variable: enzyme

Fixed Factor(s): food, parasitism

Random Factor(s):

Covariate(s):

WLS Weight:

Model: Contrasts... Plots... Post Hoc... Save... Options...

OK Paste Reset Cancel Help

(b)

Univariate: Options

Estimated Marginal Means

Factor(s) and Factor Interactions: (OVERALL), food, parasitism, food\*parasitism

Display Means for: food\*parasitism

Display:  Descriptive statistics,  Estimates of effect size,  Observed power,  Parameter estimates,  Contrast coefficient matrix,  Homogeneity tests,  Spread vs. level plot,  Residual plot,  Lack of fit,  General estimable function

Significance level: .05 Confidence intervals are 95%

Continue Cancel Help

(c)

Output - SPSS Viewer

Univariate ANOVA

Levene's Test of Equality of Error Variances<sup>a</sup>

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Corrected Total
Corrected Model	50.607 <sup>a</sup>	3	16.869	4.280	.005	.12140	866
Intercept	20007.953	1	20007.953	5185.195	.000	.51851	1.000
food	4.002	1	4.002	1.009	.326	.1239	.005
parasitism	29.808	1	29.808	7.760	.006	.7760	.791
food * parasitism	15.719	1	15.719	4.082	.045	4.082	.419
Error	600.794	156	3.851				
Total	20059.353	160					
Corrected Total	651.400	159					

a. Computed using alpha = .05  
b. R Squared = .076 (Adjusted R Squared = .059)

Estimated Marginal Means

food \* parasitism

food	parasitism	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	1	10.697	.310	10.274	11.500
1	2	11.125	.310	10.512	11.737
2	1	10.614	.310	10.001	11.227
2	2	12.105	.310	11.492	12.718

(d)

Figure (ii) (a) Data entry and test selection for a parametric general two-way ANOVA in SPSS, (b) entering the variables for the analysis, (c) entering the interaction term(s) and checking the appropriate 'Display' options, (d) the results output.

and the GLM estimates the values of the Overall Mean and the coefficients  $a$ ,  $b$  and  $c$ . The details need not concern us here, except to note that the 'intercept' = Overall Mean. Thus the table is telling us that the overall mean value for *enzyme* is very highly significantly different from zero. This

is not usually very informative, since normally any dependent variable in analysis will have an overall mean value different from zero.

Another important difference compared with AQB lies in the table of 'Estimated Marginal Means' that we requested when we entered the



interaction term in Fig. (ii)c. Note that every standard error is identical. This is because the GLM takes the assumption of equal variances in ANOVA to be literally true, and estimates the common standard error from the overall within-group variance. In AQB, each group is treated separately, and hence the standard errors are all different (see Fig. (i)). A further difference lies in the 'Power' column, which tells us the power with which the actual difference in mean values has been detected (see Box 3.14).

Finally, note that the relevant degrees of freedom to cite with the  $F$ -ratio for each effect are the ones on the same line as the variable (or interaction) in the table (the between-group d.f.), and, in each case, the Error degrees of freedom (the within-groups d.f.).

From all this, we conclude that:

- there is no evidence of a difference in enzyme activity between food types (because  $F = 1.3$ , d.f. = 1,156, ns)
- there is evidence of a difference in enzyme activity between parasitism treatments (because  $F = 7.8$ , d.f. = 1,156,  $p < 0.01$ )

- there is evidence of an interaction between food type and parasitism that affects enzyme activity (because  $F = 4.1$ , d.f. = 1,156,  $p < 0.05$ ).

We can *either* present the statistical results in the text of our report, as above, *or* put them in the legends to figures that show the means and s.e.s, *or* reproduce the ANOVA table (Fig. (i) or Fig. (ii)d), but they should be given *only once*. It is good practice to give the table for all but the simplest of designs, because it is very informative to experienced scientists. When we are testing specific predictions within the ANOVA, the table becomes all but essential, as we shall see in Box 3.9b.

N.B. The 'General Linear Model'/'Univariate' option in SPSS can be used to perform a one-way ANOVA (Boxes 3.5a,b) too. Indeed, as a multi-factor ANOVA model, it can be used for ANOVAs involving, in principle, any number of factors, in which case we might be dealing with three-way, four-way etc. analyses, all working in the same basic way. One limitation on multifactor ANOVAs with large numbers of factors, however, is the near impossibility of interpreting some of the multi-way interactions they produce.

### Box 3.9b Mean values – a *specific* parametric two-way ANOVA

We have deliberately chosen a complex example here to show you how this works.

An experiment was designed to test the effect of temperature and motivation on the mathematical performance of students in a standard test. Students were tested one by one in an environmental chamber that could be set to one of three conditions of temperature (cold, normal warm room temperature or hot). Each student was pre-treated in one of three ways to alter their motivation: (a) in a very off-hand manner and generally given the impression that they were incompetent (low motivation), (b) neutrally with respect to their competence (neutral) or (c) in a praising fashion, giving every encouragement (high motivation). Each student was randomly allocated to

one temperature and one motivation treatment. Thus the target variable for analysis is mathematical performance, measured on a constant-interval scale. There are two grouping factors: *motivation*, with three levels (low, neutral, high) and *temperature*, with three levels (cold/warm/hot). We shall treat these as fixed factors (see Box below). Testing the residuals for normality (see Box 3.1) demonstrated that the assumption of normality is acceptable (Shapiro-Wilk = 0.99, d.f. = 225, ns). The specific predictions deriving from this were:

- Prediction 1 (main effect 1): lower motivation will reduce mathematical performance,
- Prediction 2 (main effect 2): higher temperatures will reduce mathematical performance,

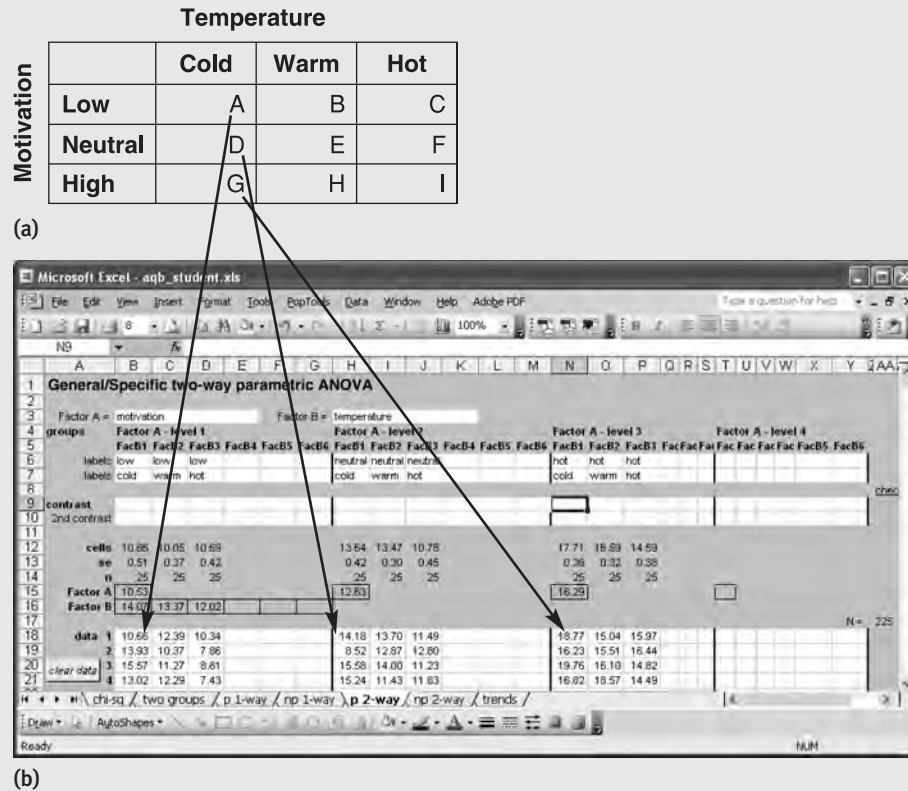


Figure (i) (a) The design, and (b) data entry in AQB of the two-way ANOVA under consideration.

- Prediction 3 (interaction): higher temperatures and lower motivation together will reduce mathematical performance more than expected from their additive effects.

In order to see how to proceed, it is best to visualise the pattern of the design (Fig. (i)) and label each of the combination of treatments (the cells of the design matrix) with letters. As with the one-way ANOVA (Box 3.5b), we should, in theory, be able to test for a specified rank order of mean values, but, as we noted in Box 3.5b, this method has not been implemented in most statistical packages. We can test for a specified rank order in a non-parametric ANOVA (see Box 3.9c), but not easily in a parametric ANOVA. Once again, therefore, we shall use contrasts to test hypotheses that are almost the same thing.

### Setting up the predictions

*Prediction 1:* We have three levels of the factor *motivation* (the rows in Fig. (i)), giving us two degrees of freedom, and hence two single-degree-of-freedom contrasts that we can use (as long as they are independent of one another). The first will contrast low motivation versus the other two, and the second will contrast neutral versus high. We proceed by casting them in terms of the cells of the design matrix (Fig. (i)):

- (a) *high motivation* > (*neutral* + *low*). The average performance for the high motivation cells of the design ( $n = 3$ ) will be greater than those for the other cells ( $n = 6$ ), i.e.:

$$(G + H + I)/3 > ((D + E + F) + (A + B + C))/6$$

which rearranges as:

$$(-1)A + (-1)B + (-1)C + (-1)D + (-1)E \\ + (-1)F + (2)G + (2)H + (2)I > 0$$

where the numbers in the brackets are the coefficients of the first contrast (cf. Box 3.5b).

- (b) *neutral > low motivation*. The average performance measured for neutral motivation ( $n = 3$ ) will be higher than that for low motivation ( $n = 3$ ), i.e.:

$$(D + E + F) > (A + B + C)$$

which rearranges as:

$$(-1)A + (-1)B + (-1)C + (1)D + (1)E \\ + (1)F + (0)G + (0)H + (0)I > 0$$

where the numbers in the brackets are the coefficients of the second contrast.

*Prediction 2*: This is done in exactly the same way, involving the *temperature* factor (the columns of Fig. (i)) rather than *motivation* (the rows of Fig. (i)). The relevant coefficients are:

- (a) *cold temperature versus other*. The average performance measured for the cold temperature cells of the design ( $n = 3$ ) will be greater than those for the other cells ( $n = 6$ ), i.e.:

$$(A + D + G)/3 > ((B + E + H) + \\ (C + F + I))/6$$

which rearranges as:

$$(+2)A + (-1)B + (-1)C + (2)D + (-1)E \\ + (-1)F + (2)G + (-1)H + (-1)I > 0$$

- (b) *warm versus high temperature*. The average performance measured for the normal temperature cells of the design ( $n = 3$ ) will be greater than those for the high temperature cells ( $n = 3$ ), i.e.:

$$(B + E + H)/3 > (C + F + I)/3$$

which rearranges as:

$$(0)A + (1)B + (-1)C + (0)D + (1)E \\ + (-1)F + (0)G + (1)H + (-1)I > 0$$

*Prediction 3*: This concerns the interaction between *temperature* and *motivation*, and it has four degrees of freedom (because the interaction has  $(a - 1)(b - 1)$  degrees for the two factors A and B which have  $a - 1$  and  $b - 1$  degrees, respectively). We shall only use two of these degrees of freedom in casting the prediction into independent contrast sets. An interaction is *non-additivity of factors*, which means that the effect on one factor (say *temperature*) is different at the different levels of the other factor (*motivation*) (see *text*). This is complicated to visualise, and it is best to set out the interaction prediction as a difference between levels in the following way:

If the hypothesis is correct, then the *motivation* differences will increase as *temperature* decreases because both cold temperature and high motivation increase the ability of students to perform. The prediction is that this ability will increase more than the additive effect of each factor, thus:

- (a) the motivation difference (*high - (neutral + low)*) will increase with temperature, so

$$cold > (warm + hot)$$

$$(G - (A + D)) / 3 > [(H - (B + E)) + \\ (I - (C + F))] / 6$$

which rearranges as:

$$(-2)A + (1)B + (1)C + (-2)D + (1)E \\ + (1)F + (2)G + (-1)H + (-1)I > 0$$

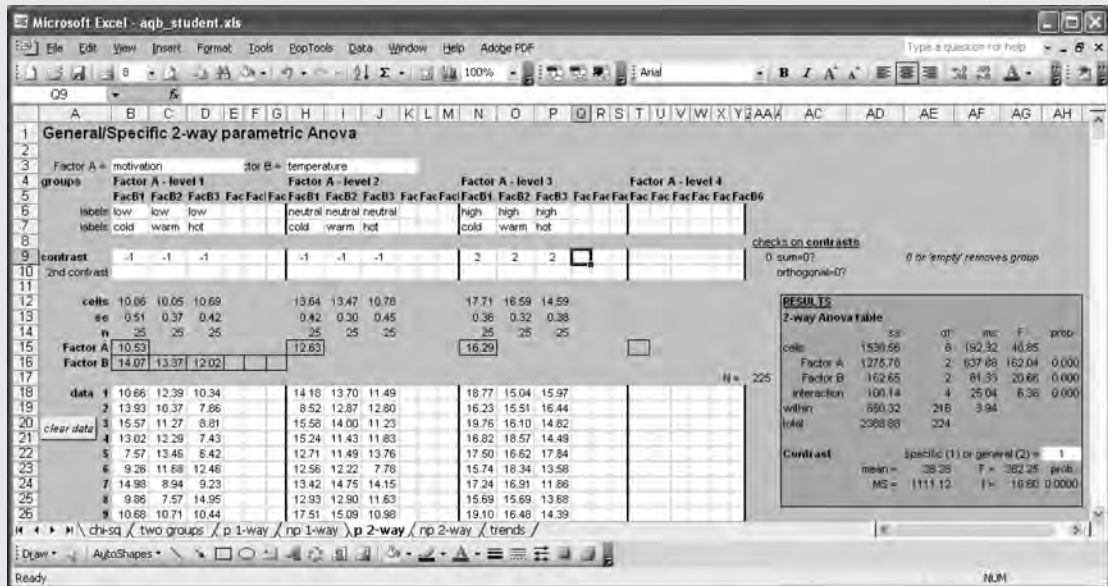
and:

- (b) the motivation difference will follow the temperature pattern *warm > hot*

$$H - (B + E) > I - (C + F)$$

which rearranges as:

$$(0)A + (-1)B + (1)C + (0)D + (-1)E \\ + (1)F + (0)G + (1)H + (-1)I > 0$$



(a) Note that the width of some of the columns has been reduced to fit data and result into a single frame.



(b)

Figure (ii) (a) Results output, and (b) checking the independence of contrast sets, for a specific parametric two-way ANOVA using contrasts in AQB.

### Testing for specific differences in a two-way classification using a parametric two-way ANOVA in AQB

In AQB (Fig. (i)b), the data are as usual put into columns as indicated in the figure (see Box 3.9a for a description). Note the labels at the top that show how the two-way design is unpacked into a single set of adjacent columns (as designs become more complex, so the SPSS method of data entry becomes more advantageous). The output for a general ANOVA appears automatically in the 'RESULTS' box (Fig. (ii)a), and you can see that the effect of Factor A (*motivation*) is significant (since  $F = 162.0$ , d.f. = 2, 216,  $p < 0.001$ ), as is Factor B (*temperature*) (since  $F = 20.7$ , d.f. = 2, 216,  $p < 0.001$ ), and there is a significant *motivation*  $\times$  *temperature* interaction (since  $F = 6.4$ , d.f. = 4, 216,  $p < 0.001$ ).

Within the general ANOVA, we are testing specific hypotheses, since we expect higher tem-

perature to *reduce* (rather than merely change) performance, and hence we enter '1' in the white box within the 'RESULTS' box (Fig. (ii)a). If the contrast coefficients sum correctly to zero, then entering them will bring up the result automatically. You can tell if they are correct because their sum is on the AQB spreadsheet in cell AA9, just to the right of the row for the coefficients. As it reminds you in the adjacent cell, the sum needs to be zero. If all is OK, then the results output appears automatically in the lower part of the 'RESULTS' box labelled 'Contrast'. They will not appear if the sum of the coefficients is not zero. The test statistic in this case is ' $t$ '.

For Prediction 1, that high motivation will increase performance, the first contrast is highly significant (since  $t = 16.8$ , d.f. = 216,  $p < 0.001$ ): the mean values do indeed fall into the expected pattern because the contrast mean is positive (28.3). For the second contrast, we need to check

	ss	df	ms	F	prob
<b>cells</b>	<b>1538.56</b>	<b>8</b>	<b>192.32</b>	<b>48.85</b>	
<b>Factor A (motivation)</b>	<b>1275.76</b>	<b>2</b>	<b>637.88</b>	<b>162.04</b>	<b>&lt;0.001</b>
contrast 1: high > (neutral + low)		1	1111.12	t = 16.8	<0.001
contrast 2: neutral > low		1	164.64	t = 6.5	<0.001
<b>Factor B (temperature)</b>	<b>162.65</b>	<b>2</b>	<b>81.33</b>	<b>20.66</b>	<b>&lt;0.001</b>
contrast 1: cold > (warm + hot)		1	94.43	t = 4.9	<0.001
contrast 2: warm > hot		1	68.22	t = 4.2	<0.001
<b>interaction (motivation x temperature)</b>	<b>100.14</b>	<b>4</b>	<b>25.04</b>	<b>6.36</b>	<b>&lt;0.001</b>
contrast 1: motivation difference, hot > (warm + cold)		1	105.47	t = 5.2	<0.001
contrast 2: motivation difference, warm > cold		1	0.01	t = -0.04	ns
<b>within</b>	<b>850.32</b>	<b>216</b>	<b>3.94</b>		
<b>total</b>	<b>2388.88</b>	<b>224</b>			

Figure (iii) Setting out the results from Fig. (ii) as a conventional ANOVA table for the purpose of a written report.

whether it is independent of the first. We can do this by entering the second set of coefficients into the line below the first set (Fig. (ii)b). Once again, the box at the end of the row (cell AA10) will show zero if the two sets of coefficients are truly independent of one another (Fig. (ii)b). If they are, then we can transfer the second set of coefficients into the contrast row proper (Row 9, where the first set are) to obtain the result. We then find that, once again, the value of '*t*' is highly significant ( $t = 6.5$ , d.f. = 216,  $p < 0.001$ ), so the mean values fall into the predicted pattern here too (again with a positive contrast mean).

Similarly for Predictions 2 and 3, we enter the coefficients and read off the result from the Contrast part of the Results box.

Because of the complexity of the predictions tested in this case, it is a good idea to set out the results in an ANOVA table, as in Fig. (iii); this contains all the details that are needed for readers to see exactly what the results are.

### Testing for specific differences in a two-way classification using a parametric two-way ANOVA in SPSS

In SPSS, as usual, the data are in one column and the codes for each of the two factors in two other columns (Fig. (iv)a). Click 'Analyze' then 'General Linear Model' and 'Univariate' and transfer the relevant data variable to the 'Dependent Variable' box, and the two groups to the 'Fixed Factor(s)' box (Fig. (iv)a). Then click the 'Contrasts' button, to bring up the relevant dialogue box (Fig. (iv)b). SPSS does not allow

us to choose our own contrasts; instead we have to choose one of a set of predetermined contrasts. It does not allow any contrasts for the interaction – only the main effects are available to be chosen. Our contrasts are actually what are called 'Helmert' contrasts (clicking the 'Help' option and searching for Contrast Types explains what the various options mean in terms of hypotheses).

Clicking 'Continue' then 'OK' gives the output (Fig. (iv)c–h). First we get Levene's test for equality of variances (Fig. (iv)c); this is not significant, and hence the normal tests are justified. Then comes the general ANOVA (Fig. (iv)d), as with AQB (except it has the 'intercept', as before [Box 3.9a]). Then we get the contrast results for each main effect, *motivation* (Fig. (iv)e–f) and *temperature* (Fig. (iv)g–h) (but not the interaction). Notice that SPSS calculates contrasts quite differently, giving a mean difference and its standard error, together with the significance. Again it is best to report the results as an ANOVA table, as in Fig. (v), to clarify exactly what has been tested and the statistical evidence.

In conclusion, there is good evidence that:

- higher temperatures reduce mathematical performance,
- lower motivation reduces mathematical performance,
- higher temperatures and lower motivation together reduce mathematical performance more than expected from their additive effects (Fig. (iii)).

Univariate

Dependent Variable: score

Fixed Factor(s): motivation, temperature

Random Factor(s):

Covariate(s):

WLS Weight:

OK Paste Reset Cancel Help

(a)

Univariate: Contrasts

Factors: motivation(Helmert), temperature(Helmert)

Change Contrast: Contrast: Helmert

Reference Category: Last First

Continue Cancel Help

(b)

### Levene's Test of Equality of Error Variances<sup>a</sup>

Dependent Variable: score

F	df1	df2	Sig.
0.890	8	216	0.526

Tests the null hypothesis that the error variance of the dependent variable is equal across groups

<sup>a</sup>. Design: Intercept + motivation + temperature + motivation \* temperature

(c)

### Tests of Between-Subjects Effects

Dependent Variable: score

Source	Type III Sum of Squares	df	Mean square	F	Sig.	Noncent. parameter	Observed Power <sup>a</sup>
Corrected model	1538.556 <sup>b</sup>	8	192.319	48.853	0.000	390.826	1.000
Intercept	38918.346	1	38918.346	9886.091	0.000	9886.091	1.000
Motivation	1275.762	2	637.881	162.035	0.000	324.071	1.000
Temperature	162.653	2	81.326	20.659	0.000	41.317	1.000
Motivation * temperature	100.141	4	25.035	6.359	0.000	25.438	0.989
Error	850.322	216	3.937				
Total	41307.224	225					
Corrected Total	2388.878	224					

<sup>a</sup>. Computed using alpha = 0.05. <sup>b</sup>. R Squared = 0.644 (adjusted R squared = 0.631)

(d)

### Contrast Results (K Matrix)

		Dependent variable
Motivation Helmert contrast		Score
Level 1 vs. Later	Contrast estimate	-3.929
	Hypothesised value	0
	Difference (estimate - Hypothesised)	-3.929
	Std. error	0.281
	Sig.	0.000
	95% confidence interval for difference	Lower bound: -4.482 Upper bound: -3.375
Level 2 vs. Level 3	Contrast estimate	-3.666
	Hypothesised value	0
	Difference (estimate - Hypothesised)	-3.666
	Std. error	0.324
	Sig.	0.000
	95% confidence interval for difference	Lower bound: -4.305 Upper bound: -3.028

(e)

**Figure (iv)** (a) Data entry and selecting variables for a specific parametric two-way ANOVA in SPSS, (b) setting up the contrasts for the analysis, (c) test for equality of variances among groups, (d) the general ANOVA output (see above), (e, f) and (g, h) the output for the contrasts for the first and second main effects, respectively (see left and next page). →

Figure (iv) cont'd

## Test Results

Dependent variable: score

Source	Sum of squares	df	Mean square	F	Sig.	Noncent. parameter	Observed power <sup>a</sup>
Contrast	1275.762	2	637.881	162.035	0.000	324.071	1.000
error	850.322	216	3.937				

<sup>a</sup>. Computed using alpha = 0.05.

(f)

## Contrast Results (K Matrix)

Temperature Helmert contrast		Dependent variable Score
Level 1 vs. Later	Contrast estimate	-1.699
	Hypothesised value	0
	Difference (estimate – Hypothesised)	-1.699
	Std. error	0.281
	Sig.	0.000
	95% confidence interval for difference	Lower bound -2.252 Upper bound -1.146
	Level 2 vs. Level 3	Contrast estimate
	Hypothesised value	0
	Difference (estimate – Hypothesised)	-0.700
	Std. error	0.324
	Sig.	0.032
	95% confidence interval for difference	Lower bound -1.338 Upper bound -0.061

(g)

## Test Results

Dependent variable: score

Source	Sum of squares	df	Mean square	F	Sig.	Noncent. parameter	Observed power <sup>a</sup>
Contrast	162.653	2	81.326	20.659	0.000	41.317	1.000
error	850.322	216	3.937				

<sup>a</sup>. Computed using alpha = 0.05.

(h)

Source	SS	df	MS	F	prob
Motivation	1275.8	2	637.9	162.0	<0.001
Helmert contrast 1: high > (neutral + low)			difference = 3.93	s.e. = 0.28	<0.001
Helmert contrast 2: neutral > low			difference = 3.67	s.e. = 0.32	<0.001
Temperature	162.7	2	81.3	20.7	<0.001
Helmert contrast 1: cold > (warm + hot)			difference = 1.70	s.e. = 0.28	<0.001
Helmert contrast 2: warm > hot			difference = 0.70	s.e. = 0.32	<0.05
Motivation × temperature	100.1	4	25.0	6.4	<0.001
Within	850.3	216	60.1		
Total	2388.9	224			

Figure (v) Setting out the results from Fig. (iv) as a conventional ANOVA table for the purpose of a written report.

### Box 3.9c Medians (mean ranks) – a *general* non-parametric two-way ANOVA

General predictions in a two-way non-parametric ANOVA can be tested only if the sample sizes are equal for all groups. If they are not, then only specific predictions are testable.

Suppose a behavioural experiment measured the reaction times of different-sized insects of two kinds to a threatening stimulus. The factor *size* had three levels (small/medium/large), and the factor *taxa* had two levels (flies/beetles). Testing the residuals showed that they were far from normal (Shapiro-Wilk = 0.87, d.f. = 84,  $p < 0.001$ ) and hence a non-parametric test is appropriate. The predictions were the usual general ones for a two-way design, involving the main effects (here, *size* and *taxa*) and their interaction:

- Prediction 1: beetles and flies will differ in reaction times,
- Prediction 2: large, medium and small insects will differ in reaction times,
- Prediction 3: there will be an interaction between size and insect type that affects reaction times.

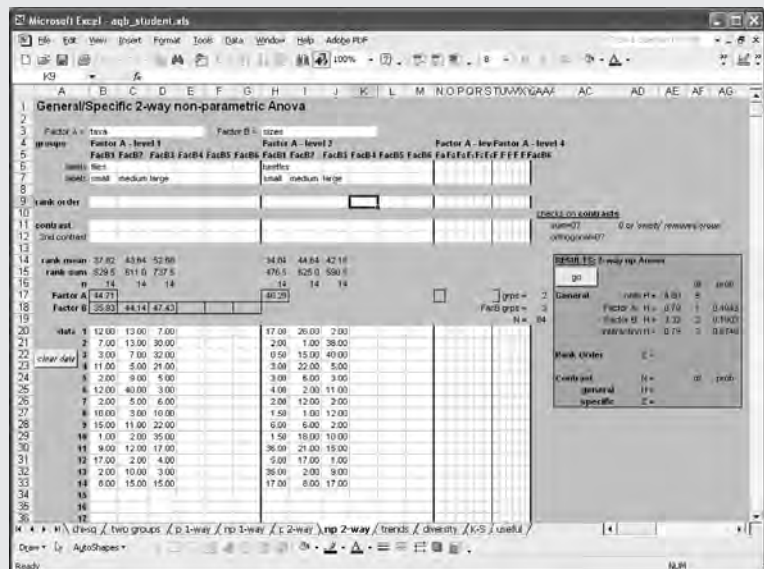
#### Testing for general differences in a two-way classification using a non-parametric two-way ANOVA in AQB

In AQB, choose the 'np 2-way' worksheet, and then put the data in columns in the usual way (Fig. (i), and *cf.* Box 3.9b, Fig. (i)b). Note the 'labels' at the top that show how the two-way design (two sets of three groups) is arranged into two sets of three adjacent columns in the 'data' box below. Clicking the 'Go' button in the 'RESULTS' box will produce the output (if sample sizes are equal – otherwise a message-box reminder will appear). Like the non-parametric one-way

ANOVA in Box 3.5c, the analysis generates the test statistic  $H$ , but in this case we get one  $H$ -value for each of the main effects and one for the interaction. You can see that the medians (mean ranks) of neither the main effects nor the interaction differ significantly, so we cannot reject the null hypothesis for any of them.

#### Testing for general differences in a two-way classification using a non-parametric two-way ANOVA in SPSS

Curiously, it is not possible to test a non-parametric two-way ANOVA in SPSS, nor in any of the major commercial statistical packages. The nearest approximation is Friedman's test, which is sometimes referred to as a two-way ANOVA, but is actually a one-way ANOVA where data come in sets of related measurements (such as those involving repeating each treatment level on the same individual). It is *not* a two-way ANOVA in the true sense considered here.



Note that the width of some of the columns has been reduced to fit data and result into a single frame.

**Figure (i)** Data entry and results output for a non-parametric two-way ANOVA in AQB.



### Box 3.9d Medians (mean ranks) – a *specific* non-parametric two-way ANOVA

A team of researchers interested in the consequences of regular exercise on various measures of health and well-being investigated the joint effects of dietary fat and exercise (running on three or more days per week) on levels of blood cholesterol. They measured cholesterol in people who had low and high fat intake, and who did or did not run. They then carried out a non-parametric two-way ANOVA. The dependent variable being analysed is blood cholesterol level. There are two factors each with two levels: *diet* (high-fat/low-fat) and *exercise* treatment (no exercise/exercise). The two-way design requires that each person being investigated has been allocated randomly to one diet treatment (low or high fat) and one exercise treatment (not running, or running). Testing the residuals for normality showed that they were far from normal (Shapiro-Wilk = 0.82, d.f. = 120,  $p < 0.001$ ), and hence a non-parametric approach is justified. Their predictions were:

- Prediction 1: running will reduce blood cholesterol,
- Prediction 2: a low-fat diet will reduce blood cholesterol,
- Prediction 3: the effect of running will be greater when on a low-fat diet.

We can test these predictions in two ways, either by predicting the *rank order* of the mean rank values, or by using single-degree-of-freedom *contrasts*. Since there are only two levels of each factor in the design, there is only one degree of freedom for each main effect, and one for the interaction. We shall see that, in a  $2 \times 2$  design like this, both methods produce exactly the same result. Any other design (e.g.  $3 \times 2$ ) will not do this.

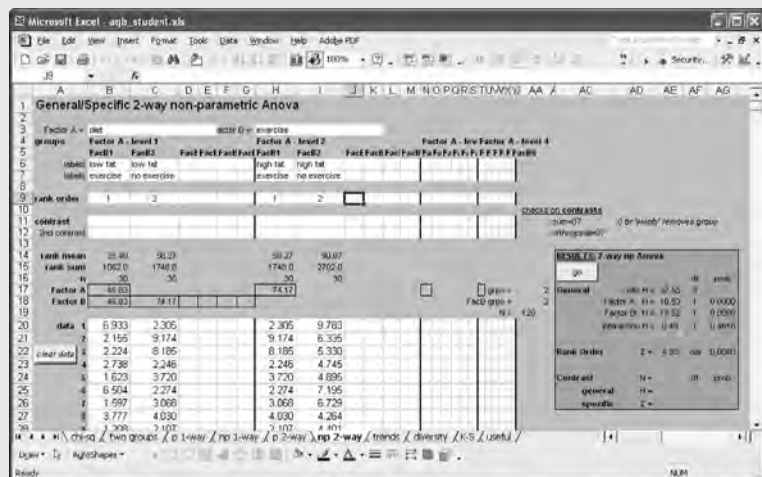
As before (Box 3.9c), it is not possible to test the specific non-parametric two-

way ANOVA in SPSS, nor in any of the major statistical packages. We can only use AQB.

#### Testing for specific differences in a two-way classification using a non-parametric two-way ANOVA in AQB

##### Testing rank order

To predict the rank order in AQB, select the 'np 2-way' worksheet, and enter the data in the usual way (see Box 3.9a). Then place the predicted rank order in the appropriate cells of the 'rank order' row (row 9) of the sheet (Fig. (i)). The example shows the test of the main effect of *exercise*. We expect cholesterol levels to be higher where no exercise is taken, and therefore place '2' in the two cells related to no exercise, and '1' in the two cells related to taking exercise. Clicking 'Go' gives us the result. You can see that a general ANOVA is produced anyway, emphasising the fact that specific tests are a component of the general ANOVA. The rank order test has  $z$  as its test statistic: it is highly significant ( $z = 4.30$ ,  $p < 0.001$ ). A similar test for the impact of diet is also significant ( $z = 2.38$ ,  $p < 0.01$ ). In the case of the interaction, there is no way to make a



Note that the width of some of the columns has been reduced to fit data and result into a single frame.

Figure (i) Data entry, setting up predicted rank orders and the resulting output for a specific two-way ANOVA in AQB.

rank-order prediction of all four cells, and therefore this prediction cannot be tested in this way.

### Using contrasts

The second method involves setting up contrasts (see Boxes 3.5b and 3.9b) for the respective predictions as follows:

		Exercise	
		Yes	No
Dietary fat	Low	A	B
	High	C	D

Figure (ii) The two-way classification in the study.

*Prediction 1:* This can be cast in the form of an inequality, following the design of the study (Fig. (ii)):

$$\text{i.e.} \quad B + D > A + C$$

and hence  $(-1)A + (+1)B + (-1)C + (+1)D > 0$

The numbers in the brackets are the coefficients to use for the contrasts, to be placed in the 'contrast' row (Row 11) of the 'np 2-way' AQB sheet (Fig. (iii)). Clicking 'Go' produces the result, which is highly significant. Comparing it with the output in Fig. (i), we can see that it is identical.

*Prediction 2:* this can be described as:

$$C + D > A + B$$

and hence

$$(-1)A + (-1)B + (+1)C + (+1)D > 0$$

Entering these coefficients produces the test result of  $z = 2.38$ ,  $p < 0.01$ . Once again, this is identical to the previous result with a predicted rank order, because of the  $2 \times 2$  design.

*Prediction 3:* this can be described in the form of differences:

$$B - A > D - C$$

and hence  $(-1)A + (+1)B + (+1)C + (-1)D > 0$

Entering these coefficients results in  $z = -0.70$ , ns. Thus there is no evidence from these data for the predicted interaction: dietary fat and exercise appear to be purely additive in their effects. If we look at the general ANOVA in the 'RESULTS' box, we can see that no interaction of *any* kind seems to be present, so the lack of a significant outcome for our particular interaction is hardly surprising.

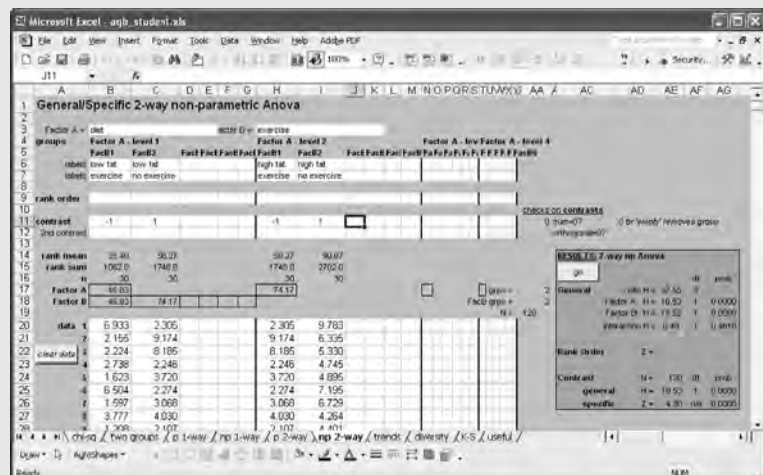


Figure (iii) Setting up the contrasts and resulting output for a specific two-way ANOVA in AQB.

### 3.3.4 Tests for a trend

As with analysis of differences, there are many tests that cater for trends. We shall introduce two simple ones here, both looking at the relationship between two sets of data. More complex versions of these tests allow multiple relationships to be tested at the same time and we shall also look at how these work.

### Correlation analysis

The first test is one of correlation. Correlation analyses calculate a test statistic known as a correlation coefficient, the two most commonly used being the parametric Pearson's product-moment correlation coefficient  $r$  and the non-parametric Spearman rank correlation coefficient  $r_s$  (Box 3.10). A correlation coefficient quantifies the extent to which there is an association between two sets of data values. A large *positive* coefficient indicates a strong tendency for high values in one set to co-occur with high values in the other and low values in one set to co-occur with low values in the other. A large *negative* coefficient indicates a strong tendency for high values in one set to co-occur with low values in the other and vice versa. Correlation coefficients take a value between +1.0 and -1.0, with values of +1 and -1 indicating, respectively, a perfect positive or negative association. 'Perfect association' means that every value in one set is predicted perfectly by values in the other set. A coefficient of 0 indicates there is no association between the two sets of values so that values in one set cannot be predicted by those in the other. If a correlation is significant, it implies that the size of the coefficient differs significantly (positively or negatively) from zero, the value expected under the null hypothesis.

As a parametric test, the Pearson product-moment correlation is subject to a number of assumptions about the data being analysed. The two sets of data must be normally distributed individually *and* jointly (a bivariate normal distribution in the jargon), both must be measured on a constant interval scale and the relationship between them must be linear. Satisfying these assumptions can be a tall order and the Pearson correlation is probably used more liberally than it should

#### Box 3.10 Analysis of trends – tests of association (correlation)

An experimenter was interested in the association between the flight speed of house sparrows (*Passer domesticus*) and their body sizes. To see whether size and speed were associated, she took a number of individual birds for which she already had some body weight data and measured their maximum flying speed in a flight chamber. Two points are important about these measures. The first is that the experimenter had not pre-determined the body weights in any way (she hadn't deliberately chosen birds of a particular weight, or fed any up to make them heavier in an experimental manipulation; she had simply caught some birds arbitrarily with respect to their weight). The second is that, should an association between weight and flight speed emerge, the cause-and-effect relationship between the two could conceivably be in either direction: heavier

birds may be stronger and so be able to fly faster (or have to spend more energy getting up speed, so fly more slowly), or birds that fly faster may burn up fat reserves and so be lighter. In other words, weight might affect flight speed, or flight speed might affect weight, or both could depend on some other (unmeasured) variable: we can't tell a priori. For these reasons, the appropriate trend analysis here is a correlation. Depending on the distribution of the data (parametric correlation requires both variables to be normally distributed, and for both together to show a bivariate-normal distribution – a three-dimensional (3D) bell shape), we could perform either a parametric or a non-parametric correlation.

This analysis is the only case where testing the raw data for normality is correct. Body *mass* seems to be normally distributed (Shapiro-

Wilk = 0.98, d.f. = 99, ns), but flight *speed* is not (Shapiro-Wilk = 0.95, d.f. = 99,  $p < 0.001$ ). We cannot test for bivariate-normal distributions without extremely large sample sizes. We shall therefore do both a parametric and a non-parametric correlation with the same data to demonstrate their similarity.

Let's suppose we have good reason for thinking heavier birds will fly faster because they are more powerful. We thus predict a positive correlation between weight and flight speed ( $H_1$ ) (the null hypothesis,  $H_0$ , is that there is not a positive correlation). This is an a priori prediction, so a specific test is appropriate.

### Correlation analysis in AQB

In AQB, select the 'trends' worksheet, and enter the weight and speed data in two columns of paired values, as in Fig. (i). The 'RESULTS' box contains output for both regression and correlation analyses, as in many statistical packages, but these are different analyses with different underlying assumptions, so, for the present case, ignore the regression output. The 'normal' correlation (called a Pearson product-moment correlation in

full) appears automatically. Since we are testing the specific prediction of a positive trend, we use the 'specific' rather than 'general' output, which gives us the one-tailed probability associated with the outcome *if* the value for the test statistic,  $r$ , has the predicted sign (positive for a positive trend, negative for a negative trend). It clearly has here, and hence the result is highly significant here.

For a non-parametric (Spearman rank) correlation, click the 'Rank' button. With large ( $n > 30$ ) sample sizes, as here, the difference in outcome between Pearson and Spearman correlations is negligible.

### Correlation analysis in SPSS

In SPSS, the data are entered in a similar paired fashion (Fig. (ii)a). Click on 'Analyze', followed by 'Correlate' and 'Bivariate'. In the resulting dialogue box (Fig. (ii)b), transfer the variables for the analysis (here *mass* and *speed*) into the 'Variables' box. For our specific test as a Pearson correlation, we check the 'Pearson' and 'one-tailed' options in the box and click 'OK'. The outcome is in the form of a table (Fig. (ii)c) because it is designed to allow any number of variables to

be entered and correlations calculated for all possible pairs. Thus the diagonal presents the correlation of each variable with itself (thus with a coefficient of 1.0, signifying, naturally enough, a perfect correlation). The correlation for each pair of different variables is given twice: once as the correlation of A on B, and once as the correlation of B on A. Thus, in Fig. (ii)c we have the correlation for *mass* on *speed* and for *speed* on *mass*. While these are exactly the same, presenting them this way underlines the fact that correlation makes no assumptions about cause and effect, i.e. about which is the

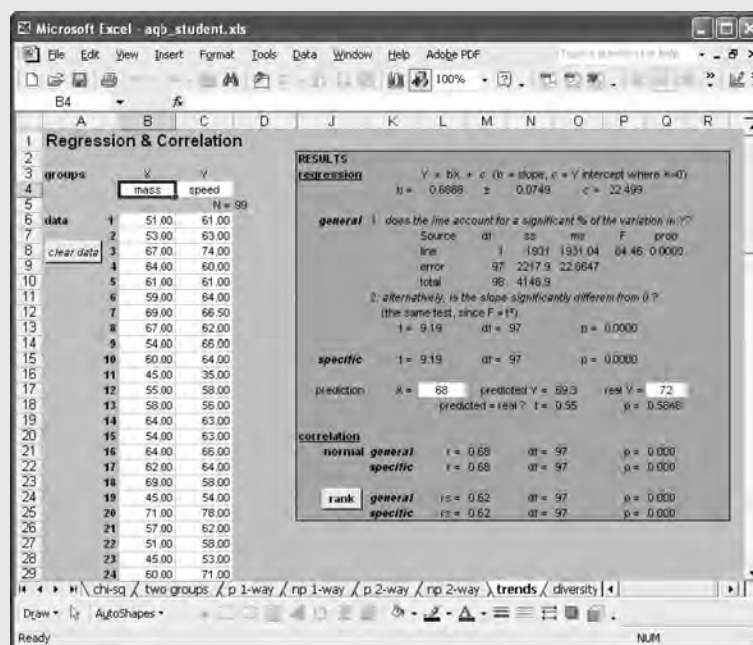
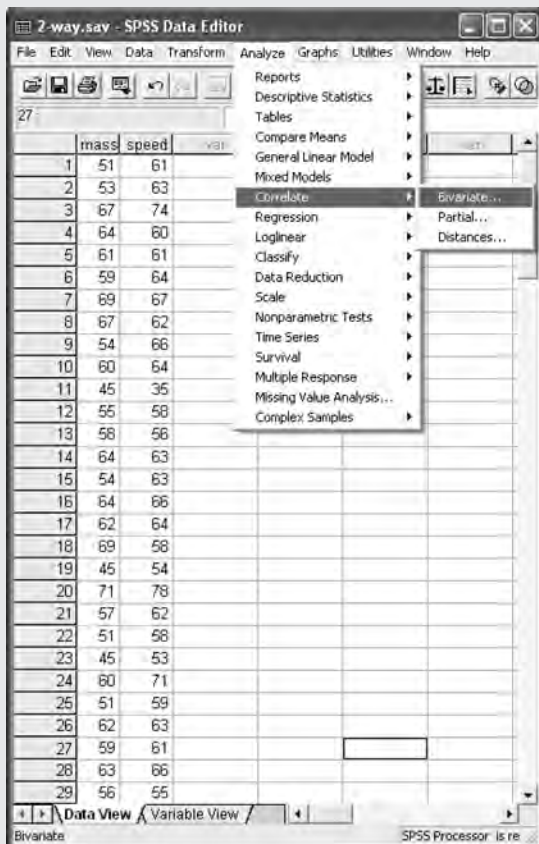
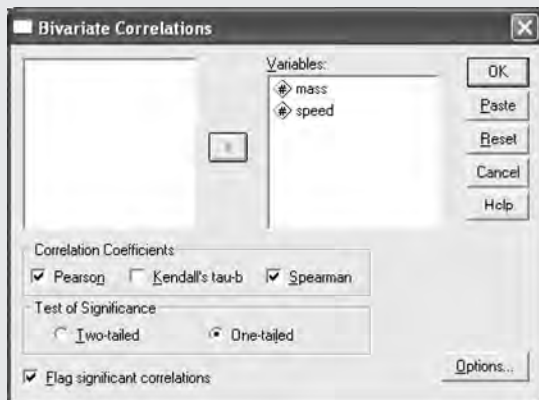


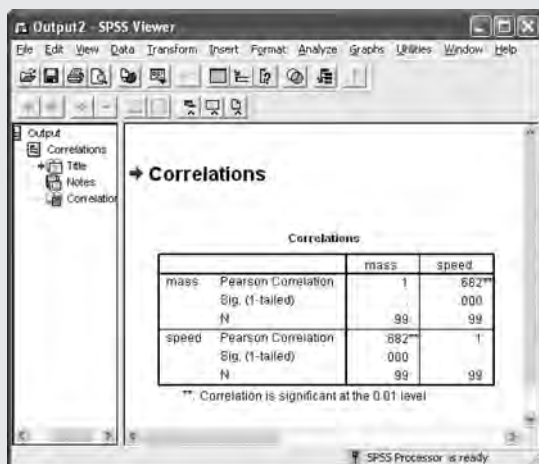
Figure (i) Data entry and results output for trend analyses in AQB.



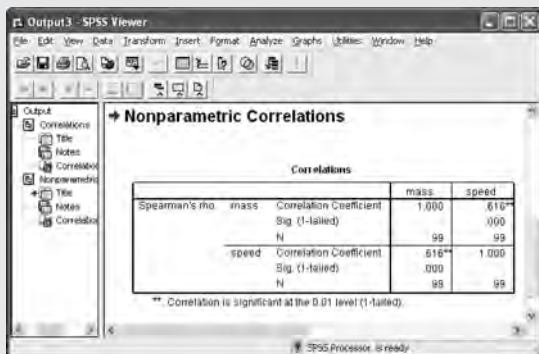
(a)



(b)



(c)



(d)

dependent and which the independent variable in the analysis. If the sign of the correlation is as predicted, we can use the one-tailed probability; if not, then the result is not significant. Here the result is clearly highly significant.

To obtain a non-parametric (Spearman rank) correlation coefficient in SPSS, we check the 'Spearman' option in Fig. (ii)b and click 'OK' to produce the table in Fig. (ii)d, which has the same layout as the Pearson table in Fig. (ii)c.


be. Being a non-parametric test, the Spearman rank correlation can be used with ordinal (ranking) or constant interval measurements and, of course, is not sensitive to departures from normality. Importantly, the relationship also need not be linear, but merely continuously increasing or decreasing (monotonic); thus  $r_s$  will test for many sorts of curved patterns of association.

While we can test for trends with correlation analyses, we must interpret them with care. Two things in particular should always be borne in mind. First, a correlation does not imply cause and effect. While we may have reasons for *supposing* that one measure influences the other rather than vice versa (*see* earlier discussion of  $x$  and  $y$  measures in trend analysis), a significant correlation cannot be used to confirm this. All it can do is demonstrate that two measures are associated. A well-known example illustrates the point. Suppose we acquired some data on the number of pairs of storks breeding in Denmark each year since 1900, and also the number of children born per family in Denmark in the same years. Plotting a scattergram, with breeding storks as the  $x$ -axis and babies as the  $y$ -axis, and calculating a correlation coefficient reveals a significant positive correlation at  $p < 0.001$  between the two measures. Do we conclude that storks bring babies? Of course not! All we can conclude is that, over the period examined, there is some association between the number of breeding storks and the human birth rate; perhaps both species simply reproduce more during long, hot summers! Second, as we have said, correlation analyses assume that associations between measures are linear (or reasonably so) if the test is parametric, or at least monotonic if the test is non-parametric. If they are neither, a lack of significant correlation cannot be taken to imply a lack of association. This is made clear in Fig. 3.4 in which the relationship between two measures is U-shaped. A correlation coefficient for this would be close to zero, but this doesn't mean there is no association. The book by Martin and Bateson (1993) contains a very useful discussion of these and other problems concerning correlation analyses.

### Linear regression

Correlation analyses allow us to judge whether two measures are associated, but that's all. Two important things that they do not allow us to establish are: (a) whether changes in the value of one measure *cause* changes in the value of the other and (b) anything quantitative about the association.

**Cause and effect.** In the case of the storks and babies above, it is fairly clear that changes in  $x$  (the number of breeding storks) do not cause changes in  $y$  (the number of human babies): the association arises either through some indirect cause-and-effect relationship (probably that habitat destruction for storks and the demographic transition to smaller family sizes in humans both happened over the course of the twentieth century) or through trivial coincidence. In many cases, though, it is not so clear whether there is or is not a direct cause-and-effect relationship between  $x$  and  $y$ . The best way to decide is to do an experiment where, instead of merely measuring  $x$  and  $y$  as in correlation analysis, we experimentally change  $x$ -values and measure what subsequently happens to  $y$ . If we see that  $y$  changes when  $x$  is changed, we can be reasonably



**Figure 3.4** Associations between two variables need not be linear. Correlation analysis would not reveal a significant trend, but an association between the variables appears to exist (see text).

happy with a cause-and-effect interpretation. This is quite different from correlation analysis.

**Quantitative relationships.** In addition, correlation analysis does not allow us to say much about the quantitative relationship between  $x$  and  $y$  (if  $x$  changes by  $n$  units, by how much does  $y$  change?) or allow us to predict values not included in the original trend analysis (e.g. can we predict the response of an insect pest to a 40 per cent concentration of pesticide when an analysis of the effect of pesticide concentration goes up to only 30 per cent?). With certain qualifications, linear regression analysis may allow us to do all the above. The qualifications arise mainly from the usual assumptions made by parametric tests, though the requirement for normality applies to one of the data sets only (the  $y$ -axis measure). However, like correlation analysis, and as its name implies, linear regression also requires the relationship to be passably linear, though there are ways of overcoming some forms of non-linearity, for instance by log-transforming the data.

Regression analysis proceeds as follows: the  $x$ -values are decided upon in advance by the investigator to cover a range of particular interest, and  $y$ -values are measured in relation to them to see how well they are predicted by  $x$ . Because  $x$ -values are selected by the investigator, they are regarded as fixed, error-free values, hence the requirement of normality only on the  $y$  measure. Linear regression then calculates the position of a line of best fit through the data points and uses the equation for this line (the *regression equation*) to predict other values for testing. The criterion of ‘best fit’ in this case is the line that minimises the magnitude of positive and negative deviations from it in the data – a more precise, mathematical way of doing what we attempt to do when drawing a straight line through a scattergram by eye. The significance of a trend can be assessed in one of two ways: the first is based on the difference in the *slope* of the best fit line from zero (a line with zero slope would be horizontal) and is indicated by the test statistic  $t$ ; the second tests whether a significant amount of variation in  $y$  is accounted for by changes in  $x$ , and is indicated by the test statistic  $F$ . Generally speaking, if the trend fails to reach significance, the best fit line should not be drawn through the points of the scattergram. Box 3.11 shows how to perform a linear regression in AQB and SPSS.

**Multiple regression.** So far we have just dealt with simple linear regression involving a single  $x$  and single  $y$  variable. It is not difficult to imagine, however, that we might have several candidate independent ( $x$ ) variables that could potentially explain variation in our dependent ( $y$ ) variable. If so, we need some way of taking them into account simultaneously so we can control for the effects of other variables when seeking effects of any particular one. Happily, there is a relatively straightforward set of techniques for doing this, which come under the umbrella term *multiple regression*. As ever, there are various ‘dos’ and ‘don’ts’ associated with the procedures, but Box 3.12 outlines a reasonably robust way of setting about the business.

### Data reduction

We conclude our summary of statistical procedures on a somewhat different note from the preceding significance tests. As we have seen in the case of analysis of variance and regression, we are often confronted with several variables that could individually or in various combinations and interactions, explain variation in a dependent variable. Multifactor and multivariate analysis of variance and multiple regression are one set of procedures that allows us to deal with this kind of situation. However, there is another way, which is not itself a significance test of any kind, but is instead a statistical technique for collapsing several interrelated variables down to one or two composite variables that can then be treated as new data (either dependent or independent variables). With the caveat that resulting composite variables can sometimes be complicated to interpret, this can be an extremely useful way of rendering multivariate data manageable for analysis. By way of example, we summarise a commonly used such procedure, known as *principal components analysis*, in Box 3.13.



### Box 3.11 Analysis of trends – linear regression

A microbiologist was interested in the effect of the concentration of a particular vitamin on the growth of the fungus *Aspergillus*. He therefore did an experiment in which the concentration of vitamin solution added to the fungal culture was varied and subsequent growth of the culture measured (as mg dry weight of fungus).

As the vitamin, and its requirement by the fungus, is known, prior information suggests that there should be a relationship between concentration and growth and that it should be positive, with higher concentrations leading to greater growth. We thus have a specific prediction of a positive relationship between the two variables ( $H_1$ ) and a null hypothesis ( $H_0$ ) of no positive relationship. Unlike the correlation analysis in Box 3.9a, however, we are here predicting a very definite cause-and-effect relationship, because we have *experimentally fixed (manipulated)* vitamin concentrations in order to see the impact of these manipulations on growth. In statistical terminology, therefore, vitamin concentration is the independent variable (the predictor), and growth the dependent variable (the response). Analysis of the residuals (see Box 3.1) shows that they can be regarded as normally distributed (Shapiro-Wilk = 0.98, d.f. = 25, ns).

#### Performing regression analysis in AQB

On the 'trends' worksheet in AQB, data for regression are entered in paired columns exactly as for correlation (Fig. (i)). The output appears automatically in the 'RESULTS' box. The slope ( $b$ ) of the regression line and its standard error (here,  $b = 0.184 \pm 0.014$ ) is given at the top, along with the equation describing the line (in which

$b$  is the coefficient determining the slope, and  $Y$  is the *growth* value predicted by the equation for any given vitamin concentration). Next the value of  $r^2$  is given, the proportion of the variation in growth accounted for by the regression relationship with vitamin concentration. There are then two ways of assessing the significance of the regression: by an ANOVA (see Boxes 3.9a–d), which calculates an  $F$ -ratio (here,  $F = 170.5$ , d.f. = 1,23,  $p < 0.0001$ ), or by a  $t$ -test (which is itself a form of ANOVA [ $t$  is just the square root of  $F$ ] – see text for the difference in what they test). Here,  $t = 13.1$ , d.f. = 23,  $p < 0.0001$ . An important point is that the ANOVA version with its  $F$ -test is a general test only, whereas the  $t$ -test can be general (two-tailed) or specific (one-tailed), depending on the nature of the prediction.

There is a further possibility in AQB. We can use the regression to make a prediction about new values of  $Y$  (here *growth*) for any given new value of  $X$  (*vitamin*). The new value of  $X$  must lie

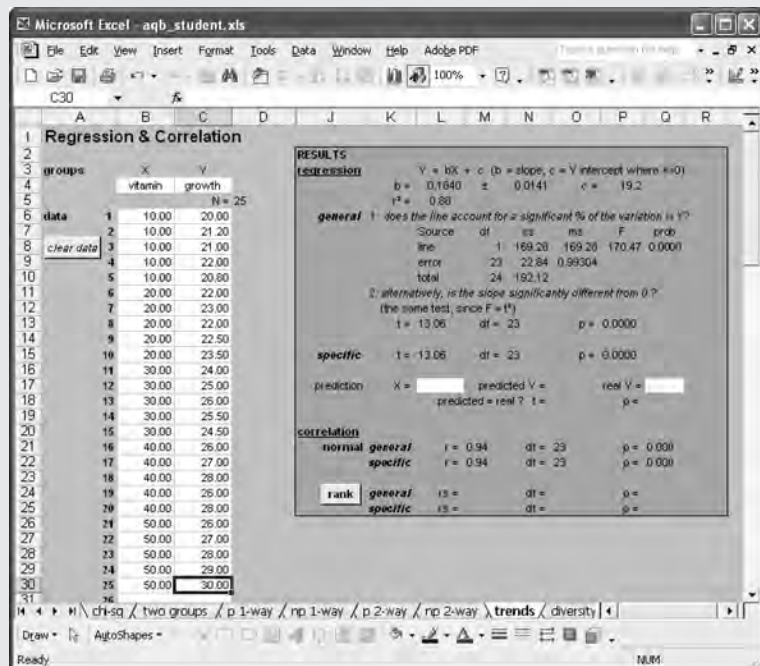


Figure (i) Data entry and results output for linear regression analysis in AQB.

within the range of X-values used to generate the regression equation, otherwise the prediction can be very inaccurate indeed. This risk of inaccuracy arises from two potential sources of error: first, there is some uncertainty in the calculated value of the slope – indicated by its standard error – which will magnify the uncertainty of the Y-value the further away we get from the middle of the line (the mean values of X and Y); second, the prediction error increases towards the ends of the regression line simply because there are usually fewer data points out there compared with the middle of the line.

Bearing this caveat in mind, we can enter a new X-value into the 'X=' box on the 'prediction' line of the AQB screen – here we've entered a new value of 45 (Fig. (ii)) – and the new value of Y predicted by it will appear as 'predicted Y' next to it. The experimenter can then carry out some new tests to see if this is upheld by the data. By entering the newly measured Y from his test into the 'real Y=' box, the experimenter can perform a *t*-test to see if it differs significantly from the value

predicted by the regression equation. The *t*-value and its associated probability appear next to 'predicted = real?' on the next line in Fig. (ii). Here we can see that the observed value of 28 is not significantly different from the predicted one of 27.5; so the regression equation seems to be a good predictor of new Y-values, at least for X-values within its range.

### Performing regression analysis in SPSS

Data for regression analysis are entered in X-Y pairs in SPSS too. Click 'Analyze', then 'Regression' and 'Linear' to obtain the dialogue box in Fig. (iii)a. Enter *vitamin* in the 'Independent(s)' box and *growth* in the 'Dependent(s)' box and click 'OK' for the output (Fig. (iii)b). This has three components. First is the correlation *r* and *r*<sup>2</sup> values, the latter of which is the proportion of variance in Y (*growth*) explained by the regression relationship with X (*vitamin*), as before. Then there is the ANOVA, as before. Finally, there is a table headed 'Coefficients', where

the estimated values for the intercept and the slope of the regression line appear, together with *t*-tests of the null hypothesis that their values are zero. Here, we can see that both differ significantly from zero (Fig. (iii)c). No facility is provided for the one-tailed, as opposed to the default two-tailed, test, so if the sign (positive or negative) of the slope is as predicted, we must halve the probability in order to get the one-tailed outcome. In the case here, the probability is so low as to be indistinguishable from zero anyway, so the one-tailed test will be highly significant.

The conclusion from both packages, then, is that increasing vitamin concentration

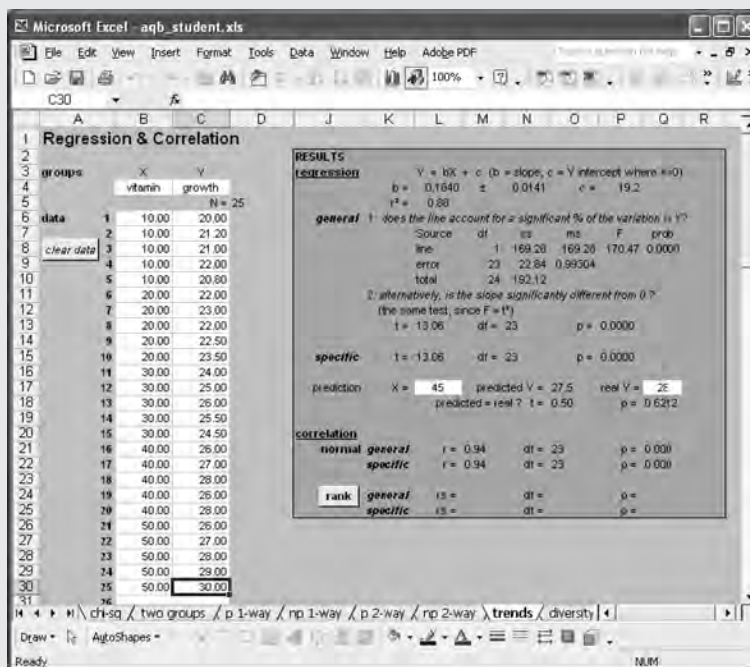
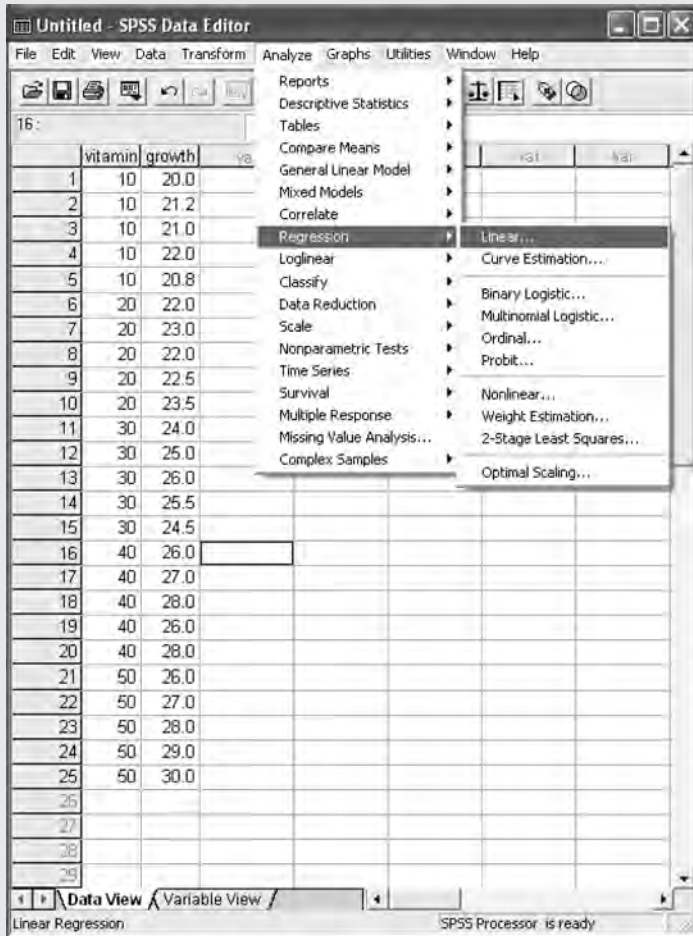
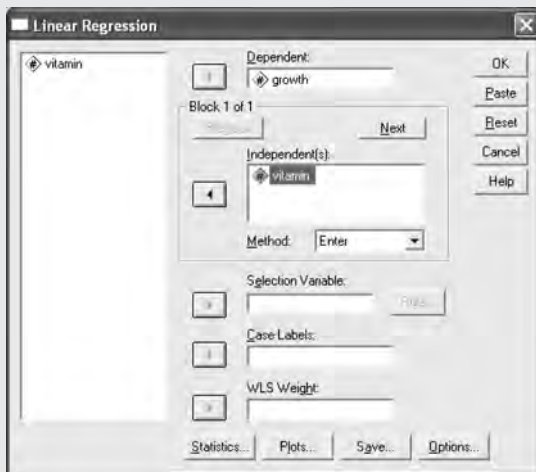


Figure (ii) Predicting and testing new values of Y for given new values of X in regression analysis using AQB.



(a)

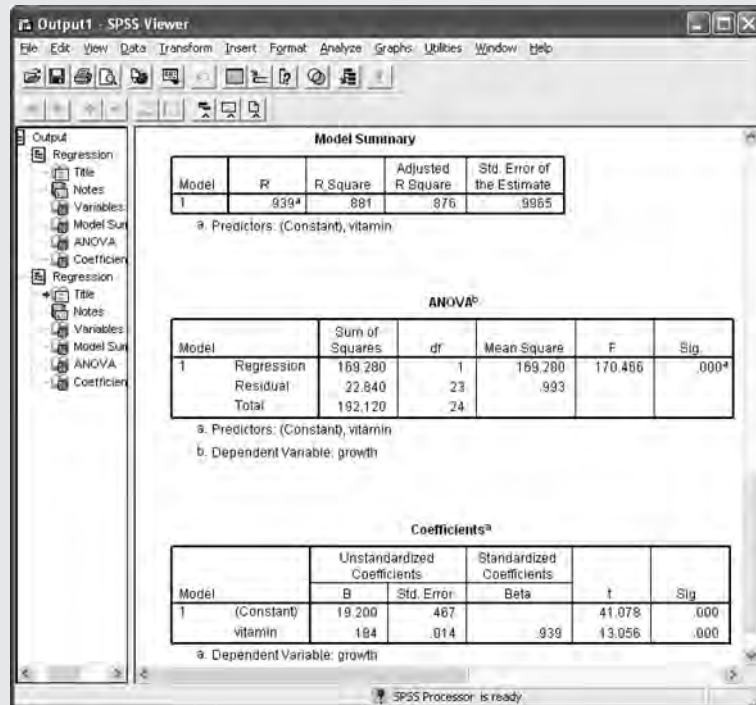


(b)

does indeed increase the growth of *Aspergillus*, and the regression equation quantifies the amount by which growth is increased by any given increase in vitamin concentration.

Figure (iii) (a) Data entry and test selection for linear regression analysis in SPSS, (b) transferring variables for the analysis, (c) the results output of the analysis (see above and below). →

Figure (iii) cont'd



(c)

### Box 3.12 Multiple regression

Often we are interested in the predictability of some measure from a number of other variables. For example, suppose we know that the body size of, and the number of eggs laid by, an insect (*Cynips divisa*) that lives inside a plant gall on oak leaves (*Quercus robur*) are very closely related to the size of gall in which it develops. The galls grow on the veins underneath the leaves. What determines gall size? We suspect that the availability of resources from the plant might be one influence, and also the extent of competition for those resources (how many galls there are per leaf or per vein?). Thus we obtain a data set taken from a large number of leaves (Fig. (i)a), where we measure the width of the gall (*gdiam*) as a measure of gall size, the length of the leaf (*llen*), the width of the leaf (*lwid*), the number of galls on the leaf (*ngalls*), the number of galls on a vein

(*vgno*), the order in which each gall developed on a vein (first gall, second, third, etc.) and the distance of each gall from the midrib of the leaf (*vdist*). Thus we might predict a priori that gall diameter (*gdiam*) should:

- increase with leaf length (*llen*) and/or leaf width (*lwid*), because bigger leaves will offer more resources,
- decrease with the extent of competition either on the whole leaf (*ngalls*) or on an individual vein (*vgno*),
- decrease with increasing *order* on a vein, and for a given order, decrease with distance from the midrib (*vdist*) (presuming that the resources for growth come down the midrib and along the veins – something that has been demonstrated in other gall systems).

Here we have one dependent variable (gall diameter, *gdiam*), and a set of independent predictor variables (*llen*, *lwid*, *ngalls*, *order*, *vngo*, *vdist*). The predictions can only be tested observationally since no manipulations are possible in this system. Even though regression is a technique developed for situations where the experimenter determines the  $x$ -values, it is very common for it to be used in other situations too, and we shall do this here. We shall check for the normality of the residuals when we have settled on the best model for predicting the dependent variable. To find the best model, we shall use an approach known as *stepwise multiple regression*, a procedure whereby the different available independent variables are added (or not, depending on their individual contribution to the model) successively to the analysis according to certain statistical rules. Experience with multiple regression shows that its results depend on how it is used, and the stepwise method of adding variables is one of the more robust ways of going about it.

What it does is to proceed in a stepwise, logical manner, adding variables to the model one by one, according to which best predicts the dependent variable at each stage. So first of all it asks which single independent variable (among all those independent variables that exceed a set threshold of good prediction) is the best predictor. Having been identified as the best, this variable is regressed against the dependent variable, and the residuals saved. Then the process is repeated, asking which of the remaining independent variables best predicts these residuals (again, as long as they exceed the threshold). This procedure carries on until either all the independent variables have been added into the model, or the threshold is no longer exceeded by any of the remaining ones. There is also a threshold to remove independent variables from the model.

### Stepwise multiple regression in SPSS

Click on 'Analyze' and 'Regression' and 'Linear' to bring up the Linear Regression dialogue box (Fig. (i)b), where *gdiam* is transferred into the 'Dependent' box, and the other variables into

the 'Independent(s)' box. Make sure the 'Method' drop-down list reads 'Stepwise'. Click on the 'Statistics' button, and in the resulting dialogue (Fig. (i)c), tick the 'Estimates', 'Model fit' and 'R-squared change' boxes. Clicking on the 'Options' button allows us to alter the thresholds of prediction used by the stepwise procedure for adding (and removing) independent variables, but this is best left to experienced statisticians.

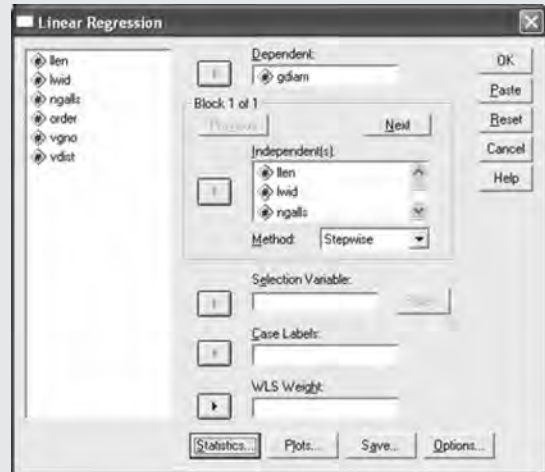
Now click 'Continue' and 'OK'. This brings up the 'Output' window (Fig. (ii)a–d). The output summarises each step of the stepwise procedure, labelled here 1–3 since there were three steps of adding independent variables. We can see (Fig. (ii)a) that the first step added leaf length (*llen*), the second added the number of galls on the leaf (*ngalls*) and the third added the order of the gall on a vein (*order*). The model summary (Fig. (ii)b) shows the change in significance of the test statistic ('*F* Change') as each independent variable is added, the correlation between the actual values of the dependent variable and those predicted by the resulting regression equation (*R*), and the amount of variation in the dependent (*gdiam*) that is explained by the independent variables at each step ( $R^2$ ). Note that in ordinary linear regression with just a single independent variable (Box 3.10), the latter is called  $r^2$ , but when there is more than one variable it becomes  $R^2$  to show it is a composite of the effects of several variables. The 'adjusted  $R^2$ ' values are the ones to look at, since these are adjusted for the correlations between the independent variables (i.e. the variables are often not truly independent of one another, but are intercorrelated, and the adjustment takes this into account). We can see that about 35 per cent (0.349) of the variation in *gdiam* is predicted by the final three-variable model. The ANOVA is provided for each step (Fig. (ii)c).

The most important part of the output is the table of coefficients (Fig. (ii)d) because this shows *how* the included independent variables affect the dependent variable. We can see that in the final model, there is a Constant (the overall mean gall diameter,  $3.3 \pm 0.5$  mm), a positive effect of leaf length (*llen*, with a slope of  $0.026 \pm 0.004$ ), a

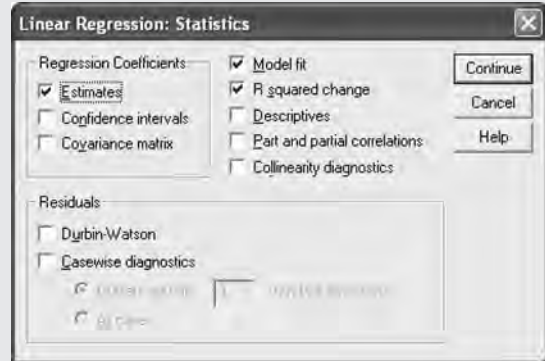
	llen	lwid	ngalls	order	vgno	vdist	gdiam
1	83	55	18	2	2	14	1.3
2	83	55	18	2	2	15	1.5
3	87	56	9	1	2	10	1.5
4	79	46	12	1	2	8	1.6
5	115	79	18	2	2	22	1.7
6	82	56	27	4	4	19	2.2
7	83	55	18	3	4	16	2.3
8	81	54	12	2	2	21	2.4
9	71	50	15	2	3	10	2.6
10	76	30	12	2	3	0	2.8
11	82	56	18	2	2	16	2.8
12	115	65	13	2	2	4	2.8
13	67	38	6	1	2	4	3.0
14	97	72	13	2	2	21	3.1
15	82	56	27	3	5	10	3.1

(a)

Figure (i) Doing a stepwise multiple regression in SPSS: (a) data input, (b) transferring the variables for analysis, (c) options to include in the output.



(b)

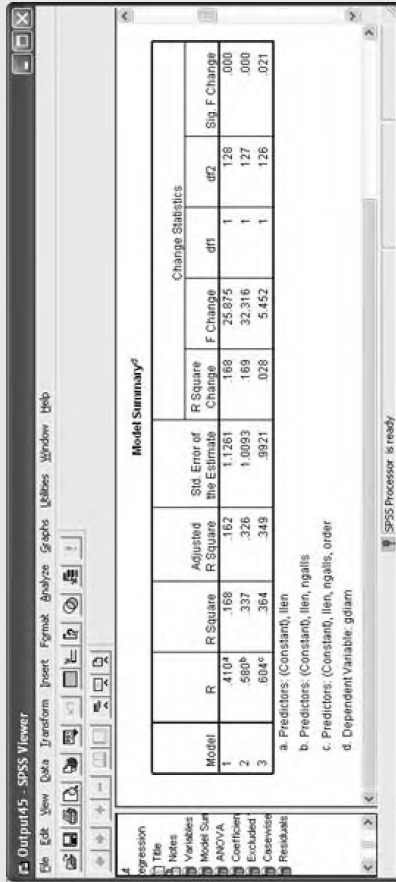


(c)

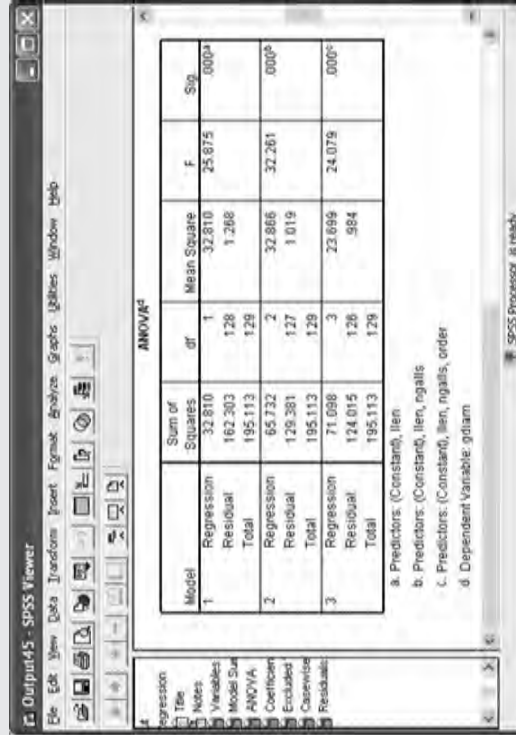
negative effect of the number of galls on a leaf (*ngalls*, with a slope of  $-0.08 \pm 0.02$ ) and a negative effect of gall order on the vein (*order*, with a slope of  $-0.3 \pm 0.1$ ).

We can now test our a priori predictions, because we have *t*-tests in the Coefficient table that test the (two-tailed, i.e. general) hypothesis that the coefficient is zero (Fig. (ii)d). We predicted the sign of the slope, and hence *if* the slope is in the predicted direction, we should halve the given probability of *t* to make the specific (one-


tailed) test. Clearly, as we predicted, there is an impact of both resources and competition, but not in all our independent variables. This is not surprising because none of them is measuring resources or competition directly, but are proxy measures intercorrelated with each other. So, for instance, once leaf length is included in the model, leaf width (*lwid*) adds nothing more and is not included. There is a table in the output (not shown) that gives details of the variable not included in the model.



(b)



(c)



(a)

Figure (ii) The output from a stepwise multiple regression analysis in SPSS. (a) Details of the independent variables entered at each step, (b) summary of the ability of the included independent variables to predict the dependent variable at each step, (c) ANOVA of the models at each step, (d) coefficients (slopes) of the effects of the predictors on the dependent variable.

Figure (ii) cont'd

Output45 - SPSS Viewer

File Edit View Data Transform Insert Format Analyze Graphs Utilities Window Help

Regression

Model Sum of Squares

ANOVA

Coefficients<sup>a</sup>

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	2.441	.476		5.126	.000
	llen	.023	.004	.410	5.087	.000
2	(Constant)	2.767	.431		6.424	.000
	llen	.028	.004	.503	6.787	.000
	ngalls	-.094	.017	-.421	-5.685	.000
3	(Constant)	3.258	.473		6.891	.000
	llen	.026	.004	.467	6.283	.000
	ngalls	-.080	.017	-.360	-4.651	.000
	order	-.294	.126	-.178	-2.335	.021

a. Dependent Variable: gdiam

SPSS Processor is ready

(d)

### Box 3.13 Principal components analysis – creating a reduced set of variables from a larger set of intercorrelated ones

A common problem in many studies is that data are collected for a large number of variables that are interrelated, so looking for differences or trends in each of them separately would result in a high degree of non-independence between tests. However, there are statistical techniques that help overcome this by producing single variables that are composites of the original separate ones. One of these techniques is known as principal components analysis (PCA).

Unlike the tests we've encountered so far, PCA isn't an hypothesis-testing method at all, but a method for reducing the number of variables to something more manageable (the so-called principal components). Thus, for us, its main attraction is that we can replace many variables with just two or three that (a) capture the main features of the data, and (b) are (by definition) statistically independent of one another. These new variables can then be used like any others in statistical tests of hypotheses, usually representing

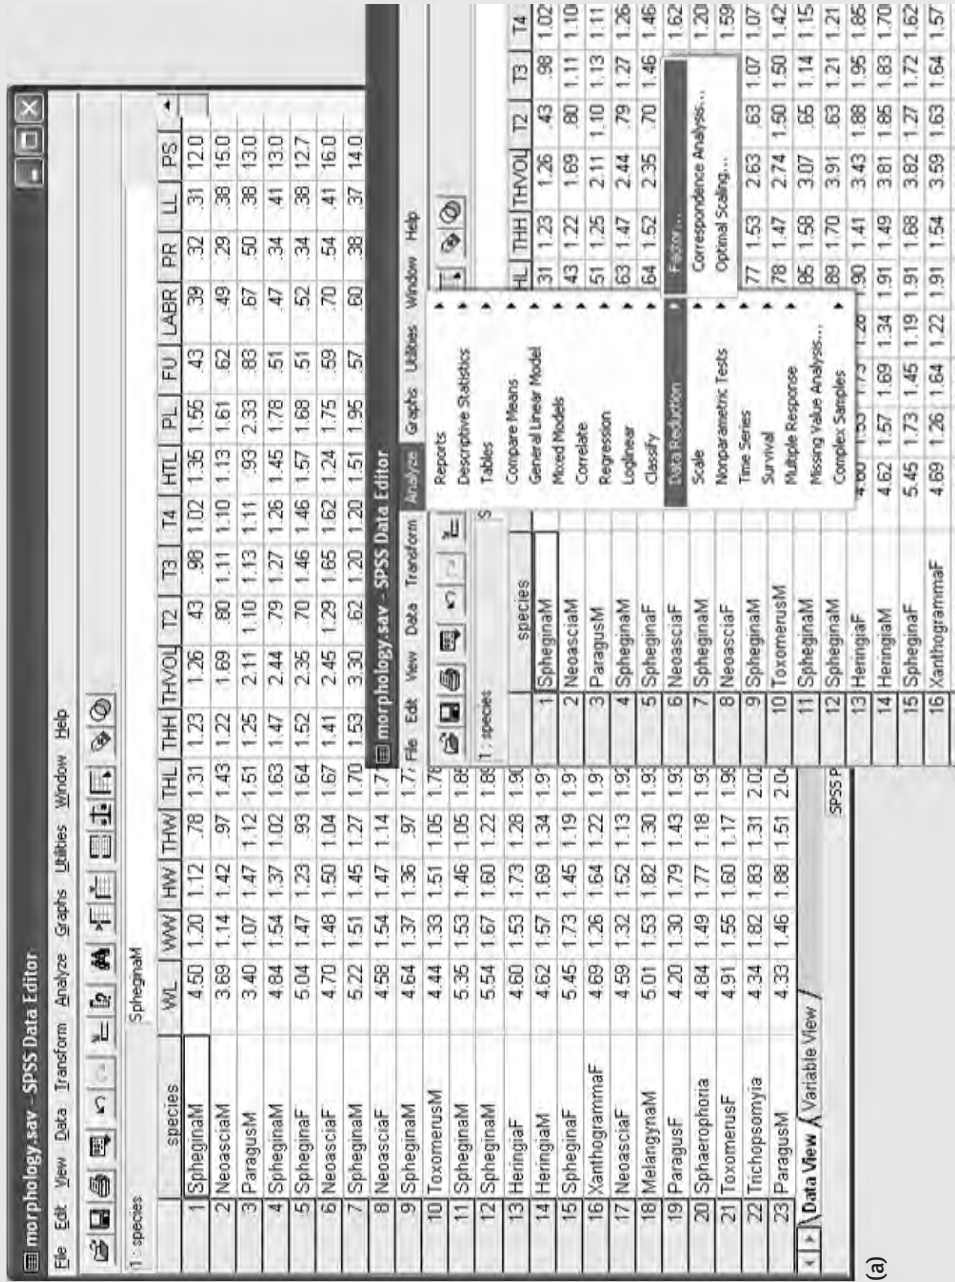
an enormous gain in clarity and efficiency. It is best learned through an example, because mathematically it is complicated!

Suppose a biologist was interested in analysing the differences in morphology among the species of flies that he worked on, and therefore he measured a set of 17 variables on a large number of individuals of many species. He then took the averages of each of the variables for each species, but kept males and females separate. This resulted in the dataset of Fig. (i). Details of the variables are in the figure legend.

#### Doing principal components analysis in SPSS

To carry out a PCA in SPSS, click 'Analyze', 'Data Reduction' and 'Factor...'. This brings up the Factor Analysis dialogue (Fig. (ii)a), where we transfer all the variables to be included into the 'Variables' box. Factor analysis covers a number of methods, but the default is PCA, so we

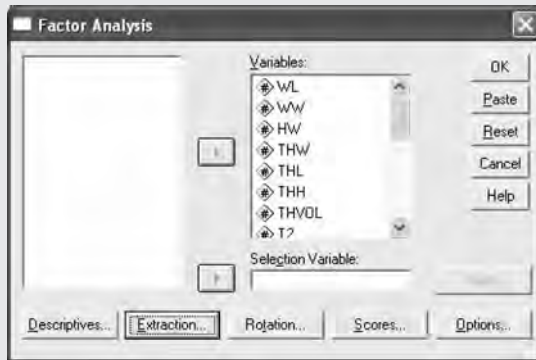




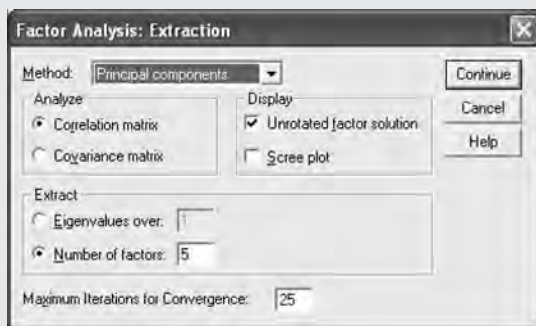
(a)

Figure (i) A data set of the morphology of a set of species of flies. The rows are labelled with the genus name only, followed by M (males), or F (females). The variables were: wing length (WL), wing width (WW), head width (HW), thorax width (THW), thorax length (THL), thorax height (THH), thorax volume (THVOL = THW\*THL\*THH), the widths of abdominal segments 2–4 (T2, T3, T4), hind tibia length (HTL), proboscis length (PL) and the lengths of the parts of the proboscis (FU, LABR, PR), width of the fleshy pad at the end of the proboscis (LL), and the number of sucking channels on the pad (PS).

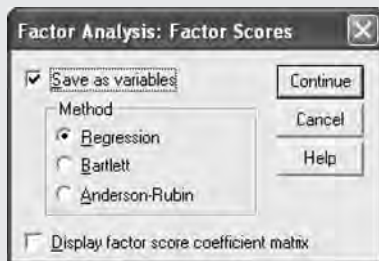
(b)



(a)



(b)



(c)

**Figure (ii)** How to do a PCA in SPSS. (a) Transferring the variables for analysis, (b) choosing the number of axes to be extracted, (c) making sure the components are saved as new composite variables in the data set.

needn't change any of the options except two. Click 'Extraction' (Fig. (ii)b) and force the PCA to extract five 'factors' by checking 'Number of factors' in Fig. (ii)b and entering '5' in the box next to it (the 'factors' are the new composite variables we are after; within PCAs they are usually referred to as 'components' or 'axes'). We do this

because often the third and subsequent components contain significant biological information, even though mathematically they don't account for much of the variation (since we had 17 variables, it is possible in principle to extract 17 new components, but beyond around five, components are likely to account for so little of the remaining unexplained variation that they can usually be ignored).

Click also on 'Scores' (Fig. (ii)a) and get the PCA to save those five components as new variables in our data set by checking 'Save as variables' in the dialogue box (Fig. (ii)c) and leaving the default 'Regression' selection under 'Method' below. This will save the components in the data file as 'fac1\_1', 'fac1\_2' to 'fac1\_5', respectively. It is thus a good idea to rename these once we have clarified what they each represent in terms of the original variables (*see below*).

Clicking 'Continue' and 'OK' brings up the Output window (Fig. (iii)a,b), with a number of tables. Unlike some statistical packages, SPSS doesn't actually provide much in the way of output for a PCA, which is curious given the reams of material it churns out for other analyses.

Looking first at the table labelled 'Total Variance Explained' (Fig. (iii)a), we can see that (as we asked) it has extracted five components (the new composite variables), but it has extracted them in order of their importance in accounting for variation in the data set. Thus the first component accounts for almost 84 per cent of the variation in the data – a huge amount. The second accounts for an additional 9 per cent, so cumulatively the first two components account for almost 93 per cent of the variation in the data. Axes 3–5 only contain small amounts of variation, but they might nevertheless be interesting biologically, so let's wait and see.

The second table of the SPSS output (Fig. (iii)b) is labelled 'Component Matrix', and represents the correlations between the new components and the original variables; it is critical to the interpretation of the new variables. Look at the first column, the correlations for the first component (headed '1' in the table). All the correlations are positive, of approximately the same magnitude,

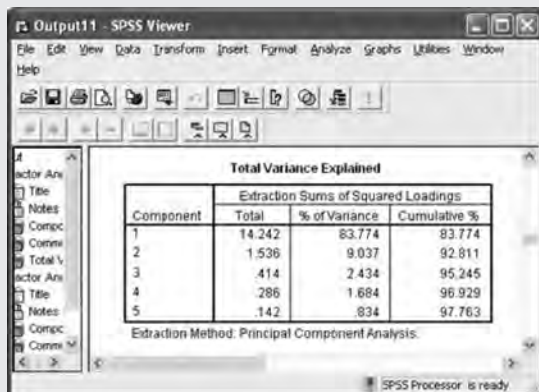


Figure (a) shows the 'Total Variance Explained' table from an SPSS PCA analysis. The table lists five components with their respective total variance, percentage of variance explained, and cumulative percentage. Component 1 explains the largest portion of the variance (83.774%).

Component	Extraction Sums of Squared Loadings		
	Total	% of Variance	Cumulative %
1	14.242	83.774	83.774
2	1.536	9.037	92.811
3	.414	2.434	95.245
4	.286	1.684	96.929
5	.142	.834	97.763

Extraction Method: Principal Component Analysis.

(a)

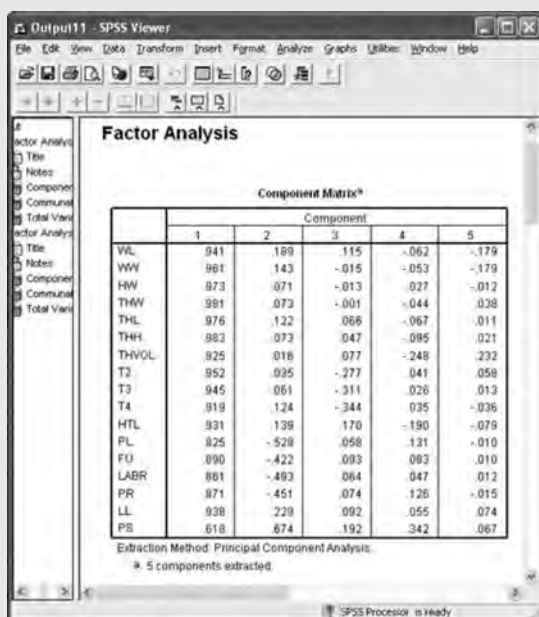


Figure (b) shows the 'Component Matrix' from an SPSS PCA analysis. This table displays the correlations between the original variables and the five extracted components. Component 1 shows strong positive correlations with variables like WL, WW, HW, THW, THL, THH, THVOL, T2, T3, T4, HTL, PL, FU, LABR, PR, LL, and PS. Component 2 shows strong negative correlations with variables like WW, HW, THW, THL, THH, THVOL, T2, T3, T4, HTL, PL, FU, LABR, PR, LL, and PS. Component 3 shows strong positive correlations with variables like THW, THL, THH, THVOL, T2, T3, T4, HTL, PL, FU, LABR, PR, LL, and PS. Component 4 shows strong negative correlations with variables like THW, THL, THH, THVOL, T2, T3, T4, HTL, PL, FU, LABR, PR, LL, and PS. Component 5 shows strong positive correlations with variables like THW, THL, THH, THVOL, T2, T3, T4, HTL, PL, FU, LABR, PR, LL, and PS.

	Component				
	1	2	3	4	5
WL	.941	.189	.115	-.062	-.179
WW	.961	.143	-.015	-.053	-.179
HW	.873	.071	-.013	-.027	-.012
THW	.981	.073	-.001	-.044	-.038
THL	.976	.122	.066	-.067	-.011
THH	.983	.073	.047	-.085	-.021
THVOL	.925	.018	.077	-.248	.232
T2	.952	.025	-.277	.041	.058
T3	.945	.061	-.311	.026	.013
T4	.919	.124	-.344	.035	-.036
HTL	.931	.138	.170	-.190	-.079
PL	.925	-.528	.058	.131	-.010
FU	.890	-.422	.093	.083	-.010
LABR	.861	-.483	.064	-.047	.012
PR	.871	-.451	.074	.126	-.015
LL	.938	.229	.092	.055	.074
PS	.918	.674	.192	.342	.067

Extraction Method: Principal Component Analysis.  
5 components extracted.

(b)

Figure (iii) Output from a PCA and its interpretation (a) details of the extracted components (the new composite variables), (b) correlations between the new axes and the original variables.

and high. Clearly this component represents all the values of the original data set relating positively together; as the value of the component increases so does the value of each of the constituent variables. We can thus characterise component 1 as a composite measure of *body size*. The second column shows that component 2 is hardly correlated with any of the variables except those that measure aspects of the proboscis. Proboscis length (PL) and the length of its parts (FU, LABR, PR) are all negative with similar magnitude, while the fleshy pad (LL) and the number of channels in it (PS) are positive. This component is therefore measuring proboscis shape (independently of body size). There are two extremes of proboscis shape: short ones with large pads, and long ones with small pads. Thus, here, as the value of the component increases, the values of variables measuring proboscis length decrease, but that of the fleshy pad increases. We must thus be aware that our new composite variables can reflect *different* directions of change among the variables making them up.

Component 3 is the only one of the remaining components with any correlation greater than 0.3. It is difficult to interpret components with only very low correlations with the original variables, and hence, in this case, we can ignore components 4 and 5. Component 3 represents the contrast between species with narrow versus those with broad abdomens (since T2, T3 and T4 are widths of the abdominal segments). Increasing scores along component 3 represent species with narrower abdomens (since the correlations are negative).

### 3.4 Testing hypotheses

The previous section has introduced an armoury of basic significance tests with which we can undertake confirmatory analyses of differences and trends. Knowing that such tests are available, however, is not much use unless we know

how and when to employ them and gear our data collection to meet their requirements. It is important to stress again, therefore, that the desired test(s) should be borne in mind from the outset when experiments and observations are being designed and the data to be collected decided upon.

### 3.4.1 Deciding what to do

Having arrived at some predictions from our hypotheses, we must decide how best to test them. This sounds straightforward in principle but involves making a lot of careful decisions. Are we looking for a difference or a trend? What are we going to measure? How are we going to measure it? How many replicates do we need? What do we need to control for? There is no general solution to any of these problems; the right decision depends entirely on the prediction in hand and the material available to test it. In a moment, we shall go back to the predictions we derived from our observational notes to see how we can test some of them. Before doing that, however, we should be aware of some important principles and pitfalls of experimental/observational design and analysis.

#### **Significance, sample sizes and statistical power**

As should be obvious from what we've said already, much hinges on the quality of the data sample we have at our disposal. It should be as representative of the population from which it derives as possible if we're to stand a chance of coming to sensible conclusions. But what does that mean? How many data values make a representative sample? What sample size do we need to make our tests for differences or trends reasonably powerful, i.e. actually capable of detecting the effects we're testing for? Sadly, there is usually no simple answer. Vague rules of thumb, such as 'at least 20 values per sample', can be and have been suggested (e.g. Dytham, 2003), but they are just that, vague rules of thumb.

What is really needed is a formal analysis of the so-called statistical *power* of a significance test in relation to the sample size available. Box 3.14 provides a summary introduction to power tests.

#### **Some dangerous traps**

**Confounding effects.** One of the commonest problems in collecting and analysing data is avoiding so-called confounding effects. Confounding effects arise when a factor of interest in an investigation is closely correlated with some other factor that may not be of interest. If such a correlation is not controlled for, either in the initial design of an investigation or by using suitable techniques during analysis, such as analysis of covariance (*see* Box 3.6), the results will inevitably be equivocal and any potential conclusions compromised. For example, suppose we wanted to know whether the burden of a particular parasitic nematode increased with the body size of the host (e.g. a mouse). We might be tempted simply to assay worm burden and measure body size for a number of arbitrarily chosen host individuals, and then perform a correlation analysis. If we got a significant positive correlation coefficient, we might conclude that burden increased with body size. An unwelcome possibility, however, is that host body

### Box 3.14 Power tests

When we carry out an experiment, we are trying to discern a pattern (a difference or a trend) in the face of variation from a number of sources. We use a statistical test to guide us, which in essence boils down all the data into a single number, the test statistic for that analysis, the distribution of which is known *under the null hypothesis*. Remember that the  $p$ -value we get from the test statistic is the probability of obtaining the data we did *if the null hypothesis is true*.

We can think about statistical testing in terms of discriminating a signal from the noise of natural variability. This noise can cause us to make errors in the conclusions of our experiment, since we don't know what the true situation is, i.e. whether the null hypothesis is true or not. This sets up four possibilities (Fig. (i)).

We can make two correct decisions (reject the null hypothesis correctly because it is actually false, or not reject the null hypothesis correctly because it is in fact true). We can also make two mistakes. One (a Type I error) is important because we are drawing a conclusion – we reject (erroneously) the null hypothesis in favour of its alternative,  $H_1$ . All statistical tests of hypotheses are designed to fix the probability (denoted  $\alpha$ ) with which such mistakes are made – it's the familiar threshold for significance, usually 0.05. We're prepared to be wrong 5 per cent of the time, but that's low enough to be reasonable.

The other mistake (a Type II error) is when we erroneously fail to reject the null hypothesis even though it is false. This is less serious since we are not drawing any conclusions about the data (unless we would like to make the much stronger claim that the null hypothesis is *actually correct*, an error of logic in scientific procedure). The rate at which this mistake happens,  $\beta$ , is a function of a number of factors, including  $\alpha$ , the design of the experiment, the test used, the sample sizes, the true (unknown) magnitude of the difference or trend, and the true (unknown) variation. The rate at which we get it right,  $(1 - \beta)$  is called the *power* of the test. Clearly, there is a trade-off between the two error rates: if we decrease  $\alpha$ ,

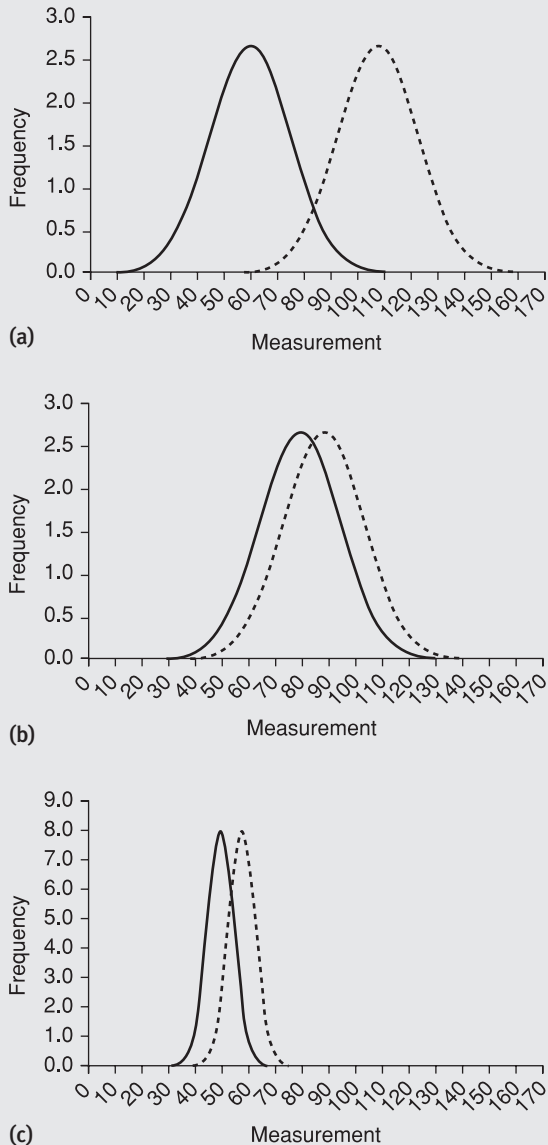
		The null hypothesis is (unknown to us)	
		TRUE	FALSE
Statistical decision	Do not reject $H_0$	✓	Type II error, at a rate $\beta$ , in part a function of the experimental design
	Reject $H_0$	Type I error, at a rate $\alpha$ set at 0.05	✓

Figure (i) The four possible outcomes of a significance test in terms of accepting or rejecting the null hypothesis (see text).

then we increase  $\beta$  (all else being equal) and the power decreases.

Suppose we have two groups we are trying to discriminate with a  $t$ -test, each with the same normal distribution of measurements but with different mean values. If the means are far apart with scarcely overlapping distributions, then they are easy to distinguish by eye, but, more importantly, are likely to be different statistically (Fig. (ii)a). Move the mean values closer so that the distributions overlap substantially, and they become much harder to tell apart by eye, and are less likely to differ statistically (Fig. (ii)b). Reduce the variance so that once more the distributions hardly overlap, and once again they are easy to tell apart (Fig. (ii)c). These situations emphasise that the power of a test depends in part on the true (unknown) situation of the data, in this case the true difference between the means, and the true variance. How easy it is to tell the two groups apart is measured by a parameter called the *effect size*, which represents the 'distance' between  $H_0$  (no difference) and  $H_1$  (a significant difference).

The accuracy of our estimates of means and variances depend on the sample size we have chosen to take, and hence is under our control. Larger sample sizes provide more accurate estimates, and hence a given difference between two



**Figure (ii)** Three different degrees of spread and separation in the distribution of two variables, leading to different degrees of discriminability (*see text*).

mean values is detected with greater power. Thus for a given statistical test, if we know a set of parameters (the Type I error rate  $\alpha$ , the effect size, and the sample size), then we can calculate the power.

How much power is enough? What is an acceptable error rate for  $\beta$ ? Like the significance level,  $\alpha$ , this is essentially an arbitrary decision. The usual levels quoted in the literature are 80 per cent (i.e.  $\beta = 0.20$ ), or (for those who believe that both Type I and Type II error rates should be the same) 95 per cent.

So what is a power test *for*? There are two main uses:

- (i) a priori power analysis. This is to ensure, *in advance*, that we are going to test our hypothesis adequately, that the experimental design we have chosen is actually capable of doing the job. We specify the power we want to have, and two of the three other parameters of  $\alpha$ , effect size and sample size: given those, we can then calculate the third parameter. For a given design, the most obvious aspect under our control is the sample size, and hence a priori power tests are usually used to calculate how large the sample sizes should be in order to achieve a particular power. We specify  $\alpha$  to be the conventional 0.05, and choose an effect size that indicates just how far away from  $H_0$  we consider important enough to warrant attention. We can choose to pay attention to ‘small’, ‘medium’ or ‘large’ effects, depending on what theory, data or cost–benefit considerations tell us. The recommended values vary with the kind of test: for a  $t$ -test they are 0.2 (a small effect size), 0.5 (medium) and 0.80 (large); for ANOVA they are 0.1, 0.25 and 0.4, respectively; and for correlation and regression, 0.02, 0.15 and 0.35, respectively (what is being calculated varies among these tests, so they are not directly equivalent). Given a particular value of effect size, we can then calculate the required sample size using an appropriate software program (such as the free-to-download G\*Power – SPSS cannot perform a priori power tests). G\*Power can be downloaded from <http://www.psych.uni-duesseldorf.de/aap/projects/gpower/>. When we do this (Fig. (iii)), the results are often a bit of a blow – for example, using a one-tailed



(a)



(b)

Figure (iii) Power calculation using G\*Power. (a) Opening screen, (b) example of the sample size required for a one-tailed  $t$ -test to detect a medium-sized effect with  $\alpha = 0.05$ .

$t$ -test to detect a medium-sized effect with  $\alpha = 0.05$  needs a sample size of 102!

- (ii) *post hoc* power testing. Here, researchers retrospectively calculate the power of the test they have done. This is the kind of power test implemented in SPSS, and demonstrated in Box 3.9a. An unfortunate tradition has built up that suggests that this kind of power test should routinely be used after a non-significant test and can be used to interpret the non-significant result – *it cannot* (see Hoenig and Heisey, 2001). Indeed, *post hoc* power tests are virtually useless since they tell us nothing more than the  $p$ -value itself gives.

size correlates with age so that bigger hosts also tend to be older. If there is some age-related change in immune competence (e.g. older hosts are less able to resist infection), a positive correlation with size could arise that actually has nothing to do with host body size. In this case size is confounded with age and simple correlation analysis cannot disengage the two. The best solution here would be to select different-sized hosts from a given age group so that the confounding effect is controlled for from the outset. Order effects are common confounding factors in many undergraduate projects. Testing animals in, say, treatment 1 first, then in treatment 2, then in 3, etc., confounds treatment with time. Animals may simply be tired by the time they get to treatment 3, so any difference in their response to the treatment could be due to that.

**Floor and ceiling effects.** Floor and ceiling effects arise when observational or experimental procedures are either too exacting or too undemanding in some

way to allow a desired discrimination to be made. For example, looking for differences in mathematical ability among people by asking them the solution to  $5 + 3$  is unlikely to be very fruitful because the problem is too easy; everyone will get the right answer straight away. A ceiling effect (everyone performs to a high standard) will thus prevent any differences there may be in ability becoming apparent. Conversely, if the same people were asked to solve a problem in catastrophe theory the odds are that no one would be able to do it. In this case, a floor effect (everyone does badly) is likely to prevent discrimination of ability. Floor and ceiling effects are not limited to performance-related tasks. Similar limitations could arise in, for instance, histological staining. If a particular tissue requires just the right amount of staining to become discriminable from other tissues, the application of too little stain would result in everything appearing similarly pale (a floor effect) while too much stain would result in everything appearing similarly dark (a ceiling effect). Real differences in tissue type would thus not show up at the extremes of stain application. Floor and ceiling effects are clearly a hazard to be avoided and are well worth testing for in preliminary investigations.

**Non-independence.** Another common source of error in data collection and analysis arises from non-independence of data values. In many circumstances, there is a temptation to treat repeated measures taken from the same subject material as independent values during statistical analysis. As Martin and Bateson (1993) point out, this error arises from the misconception that the aim of a scientific observation or experiment is to obtain large numbers of measurements rather than measurements from a large number of subjects. The point is, of course, that obtaining additional measures from the same subject is not the same as increasing the number of subjects in the sample. An example of such an error would be as follows. Suppose an investigator wished to assess the average rate of nutrient flow in the phloem of a particular plant species. Setting up a preparation might be involved and time-consuming. To save effort, the investigator decides to take as many measurements as possible from each preparation before discarding it. As a result, there are 15 measurements from one preparation, 10 from another and 12, 16 and 5 from three more. To calculate the average, the investigator totals the measurements and divides by  $n = 58$  (i.e.  $15 + 10 + 12 + 16 + 5$ ). Of course, the measurements from each preparation are not independent; there may be something about the plant in each case that gives it an unusually high or low rate of nutrient flow relative to most of the plants in the population. Incorporating each measurement taken from it as an independent example of flow rate in the population as a whole is clearly going to bias the average upwards or downwards. The true  $n$ -size in the above example is five (the number of preparations), not 58. Measurements from each preparation should thus be averaged, or collapsed in some other way, to provide a single value for use in analysis. The fallacy of this kind of approach becomes obvious if we consider estimates of average plant height rather than nutrient flow rate. Few people would seriously measure the height of the same plant 16 times and regard these as independent samples of the height of the species concerned. The principle, however, is exactly the same in the flow rate example and is referred to as *pseudoreplication*.



The problem with non-independence is that it can operate at several different levels and sneak insidiously into analyses unless careful attempts are made to exclude it (*see*, for example, Boxes 3.4 and 3.5e). We have discussed it only at the level of the individual subject. However, depending on what is being measured, it could arise if, for example, related individuals are used as independent subjects or if plants grown on the same seed tray or animals kept in the same cage are used. We should thus be on our guard against it at all times.

## 3.5 Testing predictions

Having highlighted some potential pitfalls, we must now bear them in mind as we return to our main observational examples and design experiments to test some of their predictions. We shall take one prediction from each.


### Example 1

#### Plants and herbivores

**E.g. Prediction 1A(ii)** *Leaf damage by slugs will decrease the further up a plant that samples are taken.*

This predicts a negative trend between leaf damage and height up the plant. The assumption is that the height of a leaf off the ground influences its vulnerability to slugs. Two approaches immediately suggest themselves. We could conduct a survey in the field, measuring the height of leaves above the ground and scoring the amount of slug damage on each, or we could carry out an experiment, in the field or the laboratory, exposing leaves at different heights to slugs in a controlled environment. Either way, there is a formidable number of factors to take into account if we're to get a sensible outcome.

The most obvious is that our exploratory samples came from a wide range of plant species. At the very least there are likely to be confounding effects of species-specific attributes such as the presence of distasteful toxins or other deterrents. A first consideration in a field survey, therefore, might be to select plants of a similar range of heights within each of several species. This would ensure that the confounding effects of height and species were removed, but on its own it would still not be enough for a robust comparison. A major uncontrolled factor remains: the prevalence of slugs. Different species of plant are likely to occupy different habitats, some of which are more suitable for slugs than others. Another potential confounding effect may therefore need to be removed. To check, we could carry out a simple census of slug populations around our subject plants. If numbers did not differ significantly, we could happily ignore them. If they did differ, however, we should need some way of



taking them into account. One way would be to weight the recorded damage by the observed prevalence of slugs before analysis. We could analyse the relationship between weighted damage and height as a trend, using a correlation or regression analysis, but an alternative approach, that would allow us to look in more detail at the effect of plant species, would be to use two-way analysis of variance. A two-way analysis of variance, with height of leaves above the ground (low, medium, high) and plant species as the two levels of grouping, would reveal the separate main effects of height and plant species, but also the (very likely) interaction between them. The interaction term would probably be very important here because structural and developmental differences between species will almost certainly influence the effects of position up the plant. An alternative way to take slug prevalence into account in a parametric analysis of variance would be to include it as a covariate, a constant interval measurement whose effect on the data can be controlled for within the analysis of main effects and interaction.

If height did emerge as a significant predictor of slug damage, it would, of course, lead to further questions. Is the height effect due to slugs being unwilling to climb beyond a certain height? Is it due to leaves further up being tougher or more noxious? Is it due to taller plants having greater gaps between successive leaves so discouraging further ascent by slugs? Any height  $\times$  species interaction might offer a clue to some of these (and other) possibilities by highlighting species characteristics that increase or decrease the height effect. An easier way to get at them, however, might be to do some laboratory experiments.

A laboratory study would attempt to control things more tightly at the outset. One approach might be to cultivate individual plants of some of the species sampled so that they were of similar height and the important morphological characteristics, such as the number and spacing of leaves, were, as far as possible with different species, standardised. Plants of different species could then be arranged randomly or in a regular, alternating pattern on a bench, so that any systematic confounding of position and species was avoided, and each plant was exposed to the same number of similarly sized slugs for a set period, say overnight, and then scored for leaf damage. We might be tempted simply to catch a few slugs in the field and use those. However, this would be unwise. Freshly caught slugs would enter the experiment with an unknown feeding history. We would know neither their level of hunger, nor what they had recently been feeding on. Both factors could introduce unwelcome bias into the experiment or even cause a floor effect. The best thing to do would be to bring slugs into the laboratory well before the experiment (or culture them in the laboratory), feed them all on the same material (a combination of the plant species to be tested) and deprive them of food for a short time (e.g. 12 hours) prior to testing. All slugs would then be standardised for feeding experience and hunger.

The design above would allow us to assess the effect of leaf height on damage and, if we chose to observe slug activity (directly or using, for example, time-lapse photography), we might be able to conclude something about how any effect came about. Suppose we found that higher leaves did indeed sustain less damage, but that this was due not to slugs failing to get up to them but to slugs feeding for a shorter time when they did get there. Two possible explanations

might be: (a) slugs were nearly satiated by the time they reached the higher leaves or (b) higher leaves are less palatable. An easy way to test for the latter would be to present slugs with standard-sized discs of material cut from leaves at different heights. We could choose three heights (high, medium and low) on different, but standard-sized, plants. If higher leaves are less palatable, discs cut from them would sustain less damage within the test time.

## Example 2

### Hosts and parasites

**E.g. Prediction 2A** *Parasite burdens will increase with host testosterone levels.*

This prediction derives from the hypothesis that reproductive hormones might influence susceptibility to infection. Both sex and stress hormones are known to affect the immune system, often in concert, though their effects on resistance to parasites are very variable. Prediction 2A is based on the observation that adult voles in the samples generally had greater parasite burdens than juveniles, and males greater burdens than females (Fig. 2.1). Since the difference between age classes is much more pronounced in males, testosterone becomes a plausible candidate for driving the effects of age and sex on parasite burdens. As with the seemingly simple prediction about leaf height and herbivore damage, however, much needs to be thought about in testing whether there is a connection.

We could start by taking some animals from the field and assaying their parasite burdens and testosterone levels. This could be done non-destructively by taking blood samples and faeces for blood and gut parasites, and inspecting the fur for ticks, fleas and other ectoparasites. Circulating testosterone concentrations could be assayed from either the blood or the faecal samples. Since testosterone secretion is highly pulsatile, and we are interested in chronic effects, faecal samples might provide the more appropriate measure since they accumulate testosterone metabolites over a period. Depending on the degree of discontinuity of testosterone concentration across age and sex classes, and the extent to which it can be normalised, we could test for testosterone as a predictor of age and sex differences in parasite burden by including it as a covariate in a two-way analysis of variance. The variable of interest would be our measure of parasite burden and the levels of grouping sex and age class. If we first ran the analysis without testosterone and found significant age and sex effects (as Fig. 2.1 suggests we might), but then found that the effects disappeared and were replaced by a significant covariate effect of testosterone when the latter was included, we should have some evidence that testosterone was important in gen-



erating our initial age and sex differences in parasite burden. If the distribution of testosterone values did not permit this, we could instead test for differences in testosterone between age and sex classes using non-parametric analysis of variance and seek correlations between testosterone levels and parasite burden within classes.


Such analyses might tell us something, but it would be limited. One of the main problems is that simply taking animals from the field and testing for associations between hormone levels and parasites does not allow us to say anything about cause and effect. Testosterone may well influence parasite burden, but equally both may correlate with something else, such as level of social activity, which affects exposure to infection. The association would then be an artefact of testosterone levels and parasite burden being linked to exposure. A further problem is that testosterone may covary with circulating levels of corticosterone, a glucocorticoid 'stress' hormone that also influences the immune system. The best way to control for these potentially confounding effects is either to manipulate levels of testosterone and corticosterone directly, by injection or slow-release implants, and then see what happens when animals are given a controlled infection with a known parasite, or to monitor spontaneous levels of the hormones over a period and challenge in the same way. In order to use either approach, however, animals would need to be cleaned of existing parasites and given a period of acclimation before any experiment. Corticosterone plays a role in the immune response to many infections, particularly gut helminths (worms), and any residual infection is likely to compromise investigations of hormonal effects on resistance. The best approach to start with would probably be to monitor spontaneous levels of testosterone followed by challenge. While manipulating levels experimentally allows selective control of individual hormones, it can also cause unwanted side effects and disrupt the delicate interactions between physiological systems that underpin hormonal effects on resistance. It may also be necessary to remove relevant endocrine tissue surgically or by chemical ablation to prevent spontaneous secretion affecting the control of circulating levels. Such manipulations might thus be better as a follow-up to test conclusions arising from the more observational approach. The assumption in the latter, of course, is that differences in spontaneous circulating levels of hormone will predispose individuals to correspondingly different degrees of resistance. We should then look for an association between hormone levels prior to infection and the subsequent severity of the infection. Regression analysis would be the obvious candidate, perhaps using a suitable multivariate model to take the effects of both testosterone and corticosterone levels into account simultaneously.

### Example 3

#### Nematodes and pollutants

**E.g. Prediction 3B** *Species present at unpolluted sites but missing from polluted sites will show greater mortality when exposed to pollutants.*

This prediction assumes that pollution is causally responsible for the absence of certain species from polluted samples. While this seems simple enough to test, we might want to explore its basis a little more before embarking on a set of experiments.



The most obvious point is that, with samples from only three sites, pollution status and site are confounded. Particular species may be present or absent at a given site purely by chance, or because sites happened to differ in some other important respect (e.g. interstitial water content of the soil) that affected their viability for different species. Ideally, therefore, we should first replicate our samples within site categories by choosing a number of sites, say four to six, of each kind (unpolluted, polluted with heavy metals, polluted with organophosphate) across which there is some variation in other environmental features. Species that are consistently absent from polluted sites in these samples would provide a better basis for further investigation.

One way forward might then be to collect or culture representative nematode species that are present only at unpolluted sites and expose them, say in standardised Petri dish cultures, to representative concentrations of heavy metal or organophosphate pollutant, not forgetting a suitable control (e.g. distilled water). Each treatment might be replicated half a dozen times. One-way analysis of variance of mortality by treatment would then reveal any significant effect due to the experimental pollutants. Significantly greater mortality in the two polluted treatments would be evidence in favour of a direct impact of pollution on species survival and thus presence/absence at particular sites. However, a lack of any effect would not necessarily rule out pollution as being responsible for the absence of certain species from polluted sites. Simply bathing adult worms in solutions and seeing if they die is a crude approach to say the least. Pollutants may work at any of a number of points in the worms' life cycle, perhaps reducing fecundity (the number of eggs produced) or the survival of larvae. Similar experiments could be performed to test these possibilities. More subtly, the effects may depend on particular environmental conditions, for example interactions between pollutants and other chemicals in the soil. This may be the case even though the simple experiment above showed a mortality effect; bathing worms in raw pollutant may kill them, but this may not be the way they are killed by pollution in the field. More complex experimental treatments, simulating patterns of exposure in the soil, might thus be called for.


**Example 4**  
**Crickets**

**E.g. Prediction 4C** *Encounters will progress further when opponents are more similar in size and it is more difficult to judge which will win.*

This prediction derived from observing that encounters between male crickets followed an apparently escalating pattern from chirping and antenna-tapping to out-and-out fighting and that, on the whole, bigger crickets tended to win. A possibility, therefore, is that progressive escalation reflects information-gathering about the relative size of an opponent and the likelihood of winning if the encounter is continued. If relative size is difficult to judge, as when two opponents are closely matched, the likelihood of winning cannot be judged in advance and the only way to decide the outcome is to fight. The prediction is thus of a negative trend between degree of escalation and the relative size of opponents: degree of escalation should increase with decreasing difference in size.

At first sight, this seems easy enough to test using a Spearman rank correlation or regression analysis. However, we first need some way of measuring degree of escalation. So far all we have are behavioural descriptions – chirping, antennating, fighting, etc. – from which we have inferred levels of escalation. Somehow we must put numbers to these. It is clear that we cannot put the behaviours on some common constant interval scale; we cannot, for instance, say that antennating is twice as escalated as chirping and fighting ten times as escalated. The easiest thing is simply to rank them. Thus what we assume to be the lowest level of escalation, say chirping, takes a rank of 1 and the highest level a rank of  $n$ , where  $n$  is the number of levels we decide to identify. We can then use the ranks of 1 to  $n$  as the  $y$ -values in our trend analysis.

To obtain our  $x$ -values, we must decide on a suitable measure of size. Ideally, the measure should be reliable and repeatable within and between individuals; measuring the size of the flexible abdomen, for instance, might not be a good idea because this could vary with food and water intake and thus vary from one encounter to another. It would be better to measure some component of the hard exoskeleton, e.g. the length of the long hind leg or the width of the thorax, which will not vary over the time course of observations. Of course, in our analysis, we are interested in a measure of the *relative* size of opponents, so our  $x$ -values must be some measure of relative size. The most obvious might be the *difference* in size between opponents. However, it is not hard to see why this would be inadequate. Suppose we observed two crickets of thorax widths 7.5



and 8.5 mm respectively. Suppose we observed another pair of thorax widths 6.5 and 5.5 mm. In both cases, opponents differ by 1 mm and would score the same on a simple difference measure. In the first case, however, 1 mm is only 6 per cent of the combined width measures; in the second it is 8 per cent. A 1 mm difference may thus create a greater asymmetry in the likelihood of winning in the second case than in the first. As a result it would be better to use a ratio rather than a difference scale on the  $x$ -axis, e.g. size of bigger opponent/size of smaller opponent.

Having decided on our measures, we can now plan observations. Since we are looking for a trend, we want to end up with pairs of  $x$ - and  $y$ -values. One way we might proceed is to put a number of individually marked males into a sand-filled arena and record all encounters over, say, 20 minutes, noting the males involved and the highest level (on our scale of 1 to  $n$ ) to which each encounter progressed. One problem with this approach, however, is that some males would interact more than once. The pairs of  $x$ - and  $y$ -values arising from each repeat encounter could not be used independently because body size ratio would be confounded with pair of opponents and escalation levels might be influenced by the males' past experience of each other. It would therefore be better to arrange encounters between different pairs of males to provide independent replicates of a range of size ratios. Since we are using ordinal (rank) measures of  $y$  and selected ratios as  $x$ , and we have no reason to expect a linear relationship, we should test for significance in our trend using a Spearman rank correlation rather than regression.

## 3.6 Refining hypotheses and predictions

The discussions above do two things. First they give some simple indications as to how to set about testing particular predictions. Second, they show that the outcome of such tests differ in the extent to which they increase or decrease our confidence in the hypothesis from which the prediction derives. Thus, for example, a failure to find an association between leaf damage and height up a plant in a field survey (Prediction 1A(ii)) would not greatly undermine our confidence in the hypothesis (1A) that leaf damage reflected availability to slugs. A host of factors besides height up the plant is likely to affect attack by slugs. Prediction 1A(ii) is thus a very restrictive test of Hypothesis 1A. The important point, however, is that the process of testing Hypothesis 1A does not stop there. It isn't abandoned just because one rather simplistic prediction did not work out. Instead the predictions are refined, gradually ruling out confounding factors. We saw this in the plants and herbivores example, as the suggested investigations used laboratory experiments to test for effects of height by controlling for changes in the size, spacing and palatability of leaves further up the plant.

Refinement also takes place in the opposite direction. A prediction that *is* borne out does not necessarily offer direct support for its parent hypothesis. The test of Prediction 3B (nematodes and pollutants) is a good example. This predicts that the species of nematode absent from polluted sites will die when


exposed to pollutants in the laboratory. It derives from the general hypothesis (3A) that pollution reduces species diversity. The finding that pollutants kill these species when they are bathed in them in Petri dishes does not necessarily mean that direct susceptibility to pollutants is the reason for reduced species diversity at polluted sites, or even the absence of those particular species from such sites. Their susceptibility in the laboratory may reflect a more general susceptibility to stressors. A variety of chemicals, quite unrelated to the environmental pollutants in question, might have a similar effect and this should certainly be tested. The absence of a particular species from polluted sites may reflect an interaction between effects of pollutants and competitive ability, with absence ultimately being due to competitive exclusion by other species rather than mortality through pollution. A more complex experimental design comparing effects on putatively robust (present at polluted sites) as well as putatively susceptible species would increase confidence in a mortality effect if differential mortality of susceptible, but not robust, species emerged.

Although we have presented hypotheses and their predictions in a rather cut-and-dried fashion through this book, it is clear that there is really considerable fluidity in both. The relationship between hypothesis, prediction and test is a dynamic one and it is through the modifying effects of each on the others that science proceeds.

### 3.7 Summary

1. Predictions derived from hypotheses dictate the experiments and observations needed to test hypotheses. As a result of testing, hypotheses may be rejected, provisionally accepted or modified to generate further testable hypotheses.
2. Decisions about experimental/observational measurement and the confirmatory analysis of such measurements are interdependent. The intended analysis determines very largely what should be measured and how, and should thus be clear from the outset of an investigation.
3. Some kind of yardstick is needed in confirmatory analysis to allow us to decide whether there is a convincing difference or trend in our measurements (i.e. whether we can reject the null hypothesis of no difference or trend). The arbitrary, but generally accepted, yardstick is that of statistical significance. Significance tests allow us to determine the probability that a difference or trend as extreme as the one we have obtained could have occurred purely by chance. If this probability is less than an arbitrarily chosen threshold, usually 5 per cent but sometimes 1 or 10 per cent, the difference or trend is regarded as significant and the null hypothesis is rejected.
4. Different significance tests may demand different attributes of the data. Parametric and non-parametric tests differ in the assumptions they make about the distribution of data values within samples and the kinds of measurement they can cope with. Tests can also be used in specific/one-tailed or general/two-tailed forms depending on prior expectations about the direction of differences or trends.






5. Basic tests for a difference include chi-squared, the  $t$ -test, the Mann–Whitney  $U$ -test and one- and two-way analysis of variance. Each has a number of requirements that must be taken into account and care is needed not to make multiple use of two-group difference tests in comparing more than two groups of data.
6. Basic tests for a trend include Pearson product-moment and Spearman rank correlations and regression analysis. Correlation is used when we merely wish to test for an association between two variables. Regression is used when we change the values of one variable experimentally and observe the effect on another to test for a cause-and-effect relationship between two variables. Regression analysis yields more quantitative information about trends than correlation but makes more stringent demands on the data.
7. Multifactor and multivariate forms of difference and trend analyses allow the effects of several different variables to be analysed together within tests.
8. Testing predictions requires careful thought about methods of measurement, experimental/observational procedure, replication and controlling for confounding factors and floor and ceiling effects.
9. Whether or not a test supports or fails to support a prediction, the implications for the parent hypothesis need careful consideration. A hypothesis does not necessarily stand or fall on the outcome of one prediction; everything depends on how discriminating the prediction really is.

## References

- Bailey, N. T. J. (1981) *Statistical methods in biology*, 2nd edition. Hodder and Stoughton, London.
- Day, R. W. and Quinn, G. P. (1989) Comparisons of treatments after an analysis of variance in ecology. *Ecological Monographs* **59**, 433–463.

- Dytham, C. (2003) *Choosing and using statistics: a biologist's guide*. 2nd edition. Blackwell, Oxford.
- Gaines, S. D. and Rice, W. R. (1990) Analysis of biological data when there are ordered expectations. *American Naturalist* **135**, 310–317.
- Grafen, A. and Hails, R. (2002) *Modern statistics for the life sciences*. Oxford University Press, Oxford.
- Hawkins, D. (2005) *Biomeasurement: understanding, analysing, and communicating data in the biosciences*. Oxford University Press, Oxford.
- Hoenig, J. M. and Heisey, D. M. (2001) The abuse of power: the pervasive fallacy of power calculation for data analysis. *American Statistician* **55**, 19–24.
- Martin, P. and Bateson, P. (1993) *Measuring behaviour*, 2nd edition. Cambridge University Press, Cambridge.
- Meddis, R. (1984) *Statistics using ranks: a unified approach*. Blackwell, Oxford.
- Pentecost, A. (1999) *Analysing environmental data*. Longman, Harlow.
- Ridley, J., Kolm, N., Freckleton, R. P. and Gage, M. J. G. (2007) An unexpected influence of widely used significance thresholds on the distribution of *P*-values. *Journal of Evolutionary Ecology* (in press).
- Ruxton, G. D. and Colegrave, N. (2006) *Experimental design for the life sciences*, 2nd edition. Oxford University Press, Oxford.
- Schluter, D. (1988) Estimating the form of natural selection on a quantitative trait. *Evolution* **42**, 849–861.
- Siegel, S. and Castellan, N. J. (1988) *Nonparametric statistics for the behavioral sciences*. McGraw-Hill, New York.
- Sokal, R. R. and Rohlf, F. J. (1995) *Biometry*, 3rd edition. Freeman, San Francisco.
- Tomas, L. and Juanes, F. (1996) The importance of statistical power analysis: an example from animal behaviour. *Animal Behaviour* **52**, 856–859.
- Underwood, A. J. (1997) *Experiments in ecology*. Cambridge University Press, Cambridge.



# 4

## Presenting information

### How to communicate outcomes and conclusions

If your investigation has gone to plan (and possibly even if it hasn't), you will have generated a pile of data that somehow needs to be presented as a critical test of your hypothesis. Performing appropriate significance tests is only one step on the way. While significance tests will help you decide whether a difference or trend is interesting, this information still has to be put across so that other people can evaluate it for themselves. There are two reasons why we should take care how we present our results. The first is to ensure we get our message over; there is little point making a startling discovery if we can't communicate it to anyone. The purpose of our investigation was to test a hypothesis. We might conclude that the results support the hypothesis or that they undermine it. Whichever conclusion we reach we must sell it if we wish it to be taken seriously. Since scientists are by training sceptical, selling our conclusion may demand some skilful presentation and marshalling of arguments. The second reason is that we must give other people a fair chance to judge our conclusions. As we saw earlier, simply saying that some difference or trend is significant doesn't tell us how strong the effect is. It is important to present results in such a way that others can make up their own minds about how well they support our conclusions. In this chapter, we shall look at some conventions in presenting information that help satisfy both these requirements. We begin with some simple points about figures and tables.

#### 4.1 Presenting figures and tables

We stressed earlier that it is usually not helpful to present raw data. Raw data are often too numerous and the information in them too difficult to assimilate for useful presentation. Instead, we summarise them in some way and present the

summary form. We have already dealt with summary statistics in a general way when we discussed exploratory analysis; here we discuss their use in presenting the results of confirmatory analysis.

Although it is usually obvious that raw data require summarising, there can be a temptation to summarise everything, as if summary statistics or plots were of value in themselves. In confirmatory analyses, they are of value only to the extent that they help us evaluate tests of hypotheses. We thus need to be selective in distilling our results. Naturally, the summaries and forms of presentation that are most appropriate will depend on the type of confirmatory analysis. It is therefore easiest to deal with different cases in turn.

#### 4.1.1 Presenting analyses of difference

Where we are dealing with analyses of difference, the important information to get across is a summary of the group values being compared. The form this takes will vary with the number of groups and levels of grouping involved. There are two basic ways of presenting a summary of differences: figures and tables. As we have argued previously, figures tend to be easier to assimilate than tables, even when the latter comprise summary statistics. However, tables may be more economical when large numbers of comparisons are required, or where comparisons are subsidiary to the main point being argued but helpful to have at hand. If tables are used, it is important that they present all the key summary information necessary to judge the claims they make. This usually means (a) summary statistics (e.g. means  $\pm$  standard errors\*, medians  $\pm$  confidence limits) for each of the groups being compared, (b) the sample size ( $n$ ) for each group, (c) test statistic values, (d) the probabilities ( $p$ -values) associated with the test statistic values and (e) an explanatory legend detailing what the table tells us. The test statistics and  $p$ -values can be presented either in the table itself or in the legend. The same information, of course, should be presented in figures except that the summary statistics are represented graphically (e.g. as bar charts) instead of as numbers, and information about sample sizes, test statistics and probability levels more conventionally goes in the legend (now usually called the *figure caption*) rather than in the figure itself. (Nevertheless, as long as it doesn't clutter the figure and detract from its impact, it can be very helpful to include statistical information within the figure and we shall do this later where appropriate.)

##### Differences between two or more groups (with one level of grouping)

Here we are presenting the kinds of result that might emerge from a Mann–Whitney  $U$ -test, or a one-way analysis of variance. Suppose we have tested for a difference in growth rate (general prediction) between two groups of plants, one given a gibberellin (growth-promoting) hormone treatment, the other acting as

\* While there are various forms of summary statistic, means  $\pm$  standard errors are widely used because the mean of a set of values is an easy concept to grasp and because the standard error estimates the distribution of *means* within the population which approaches a normal distribution as the sample size increases. It is thus usually legitimate to quote means  $\pm$  standard errors as summary statistics even when the distribution of data *values* demands a non-parametric significance test.

an untreated control. The treated group contained 12 plants and the untreated group eight. Using a one-way analysis of variance for two groups, we discover a significant difference at the 0.1 per cent ( $p < 0.001$ ) level between the groups, with treated plants growing to a mean ( $\pm$  standard error) height of  $14.75 \pm 0.88$  cm during the experimental period, and controls growing to a mean height of  $9.01 \pm 0.63$  cm. We could present these results as in Table 4.1a. Note the legend explaining exactly what is in the table.

**Table 4.1a** The mean height to which plants grew during the experimental period when treated with gibberellin or left untreated.

	Experimental groups		Significance
	Treated	Untreated	
Mean ( $\pm$ s.e.) height (cm)	$14.75 \pm 0.88$	$9.01 \pm 0.63$	$H = 12.22,$ $p < 0.001$
<i>n</i>	12	8	

The significance column could be omitted from the table, in which case the test statistic and probability level should be given in the legend. The legend would now read:

**Table 4.1a** The mean height to which plants grew during the experimental period when treated with gibberellin or left untreated.  $H$  comparing the two groups = 12.22,  $p < 0.001$ .



**Table 4.1b** The mean height to which plants grew during the experimental period when treated with gibberellin or left untreated. \*\*\*,  $H = 12.22$ ,  $p < 0.001$ .

	Experimental groups		Significance
	Treated	Untreated	
Mean ( $\pm$ s.e.) height (cm)	14.75 $\pm$ 0.88	9.01 $\pm$ 0.63	***
<i>n</i>	12	8	

An alternative and frequently adopted convention in presenting significance levels is to use asterisks instead of the test statistic and probability numbers. In this case, different levels of probability are indicated by different numbers of asterisks. Usually \* denotes  $p < 0.05$ , \*\*  $p < 0.01$  and \*\*\*  $p < 0.001$ , but this can vary between investigations so it is important to declare your convention when you first use it. Using the asterisks convention, the table would now read as in Table 4.1b. Exactly the same forms of presentation, of course, could be used for comparisons of more than two groups. However, if we had used a *U*-test to test for a difference between two groups, we should now have to change to a one-way analysis of variance to avoid abuse of a two-group difference test (*see* Chapter 3).

The presentation of means and standard errors (or medians and confidence limits) is appropriate whenever we are dealing with analyses that take account of the variability within data samples. In chi-squared analyses, however, where we are comparing simple counts, there is obviously no variability to represent. If we were presenting a chi-squared analysis of the number of plants surviving each of three different herbicide treatments and one control treatment, therefore, the table would be as shown in Table 4.2a. Table 4.2b shows an alternative presentation. If we wish to present our results as figures rather than tables, we can convey the same information using simple bar charts. Thus Table 4.1a can be recast as Fig. 4.1a. Similarly, Table 4.1b could be recast as Fig. 4.1b. For a comparison of three groups, say comparing the effectiveness of the lambda bacteriophage in killing three strains of *Escherichia coli* suspected of differing in susceptibility, the figure might be as shown in Fig. 4.2a.


In Figs 4.1a, b and 4.2a, we have assumed a one-way analysis of variance was used to test a general prediction (hence the test statistic *H*). If we had instead tested a specific prediction because we had an a priori reason for expecting a rank order of effect (e.g. gibberellin-treated plants would grow taller than untreated plants (Fig. 4.1), or strain B of *E. coli* would be most resistant and

**Table 4.2a** The number of plants surviving treatment with different herbicides.

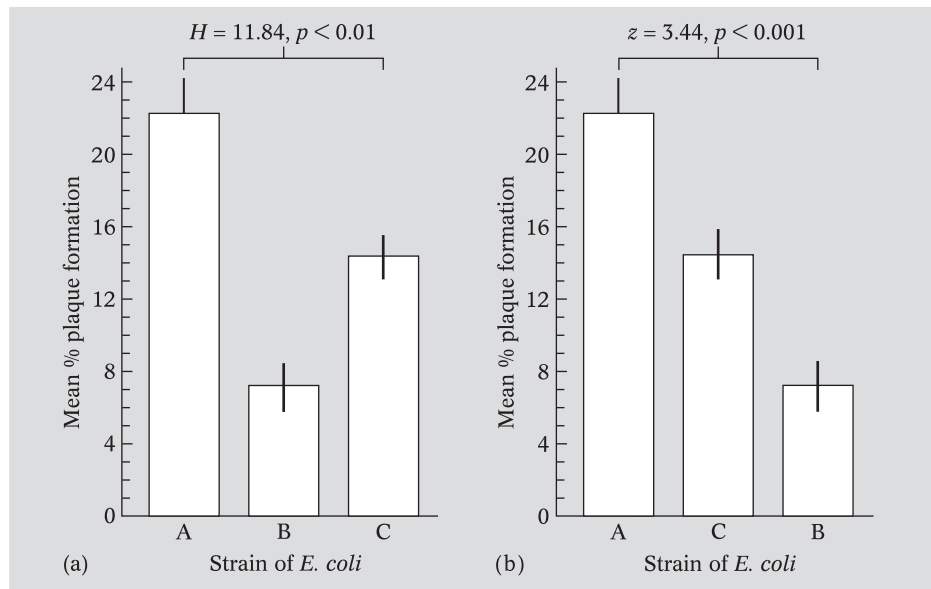
	Experimental group				Significance
	Herbicide 1	Herbicide 2	Herbicide 3	Control	
Number of plants surviving	15	8	6	27	$\chi^2 = 19.3$ , $p < 0.001$

**Table 4.2b** The number of plants surviving treatment with different herbicides. \*\*\*,  $\chi^2 = 19.3$ ,  $p < 0.001$ .

	Experimental group				Significance
	Herbicide	Herbicide	Herbicide	Control	
	1	2	3		
Number of plants surviving	15	8	6	27	***



**Figure 4.1** (a, b) The mean height to which plants grew during the experimental period when treated with gibberellin (+G,  $n = 12$ ) or left untreated (U,  $n = 8$ ). \*\*\*,  $H = 12.2$ ,  $p < 0.001$ . Bars represent standard errors. (c) Figure 4.1b with a different scale. The mean height to which plants grew during the experimental period when treated with gibberellin (+G,  $n = 12$ ) or left untreated (U,  $n = 8$ ). \*\*\*,  $H = 12.2$ ,  $p < 0.001$ . Bars represent standard errors.




**Figure 4.2** The mean percentage area of plaque (= bacterial death) formation by lambda bacteriophage on three strains (A–C) of *E. coli*. Bars represent standard errors. (a) Non-parametric analysis of variance testing a general prediction of difference between strains.  $N = 8$  cultures in each case. (b) Non-parametric analysis of variance testing a specific prediction of difference between strains (A > C > B).  $N = 8$  cultures in each case.

strain A least resistant to attack by the phage (Fig. 4.2)), we should recast the figures with groups in the predicted order and quote the test statistic  $z$  rather than  $H$ . Thus Fig. 4.2a could be recast as Fig. 4.2b.

For most people, bar charts like these convey the important differences between groups more clearly and immediately than equivalent tables of numbers. However, it is worth stressing some key points that help to maximise the effectiveness of a figure.

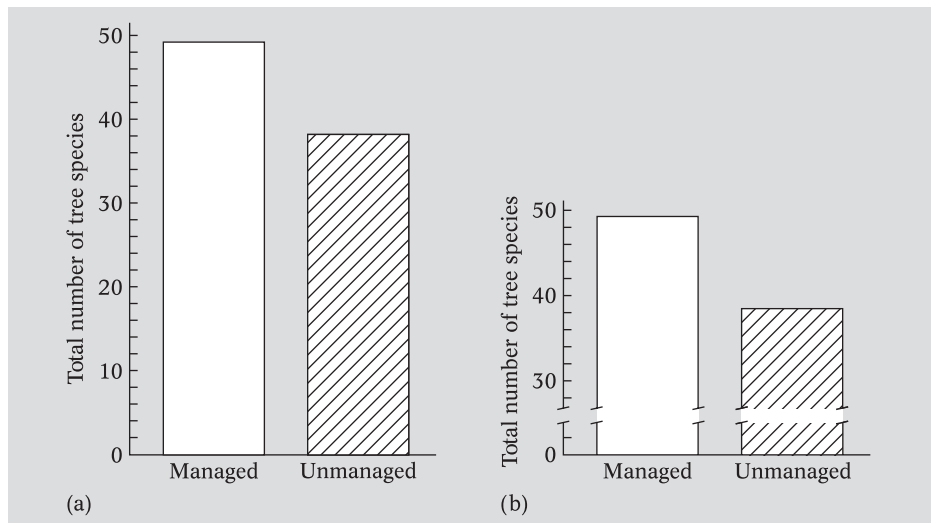
1. Make sure the scaling of numerical axes is appropriate for the difference you are trying to show. For instance, the impact of Fig. 4.1b is much reduced by choosing too large a scale (see Fig. 4.1c).
2. Always use the *same* scaling on figures that are to be compared with one another. Thus, Fig. 4.3a, b is misleading because the different scaling makes the magnitude of the bars look the same in (a) and (b). Using the same scale, as in Fig. 4.3c, shows that there is in fact a big difference between (a) and (b).
3. Make sure axes are numbered and labelled properly and that labels are easy to understand and indicate the units used. Avoid obscure abbreviations in axis labels: these can easily be ambiguous and misleading or unnecessarily difficult to interpret.
4. Axes do not have to start at zero. Presentation may be more economical if an axis is broken and starts at some other value. Thus, Fig. 4.4a could be recast as Fig. 4.4b with the break in the vertical axis indicated by a double slash.






**Figure 4.3** (a) The total number of T-helper cells in experimental samples from laboratory mice following administration of two different cytotoxic drugs (A and B).  $N = 40$  samples for each drug treatment. (b) As Fig. 4.3a, but for experimental samples from humans. (c) The total number of T-helper cells in experimental samples from laboratory mice (open bars) and humans (shaded bars) following administration of two different cytotoxic drugs (A and B).  $N = 40$  samples for each drug treatment and species.

5. Include indications of variability (standard errors, etc.) where appropriate. Also include sample sizes and  $p$ -values as long as these don't clutter the figure. If they do, put them in the legend.
6. Always provide a full, explanatory legend. The phrasing of the legend should be based on the prediction being tested and the legend should include



**Figure 4.4** (a, b) The total number of species of trees bearing epiphytes in 2 km<sup>2</sup> study areas of rainforest in Bolivia where forests are managed economically (open bar) and unmanaged (shaded bar).  $N = 1 \times 2$  km<sup>2</sup> area in each case.



any statistical information (sample sizes, test statistics, etc.) not included in the figure. Do not repeat information in both figure and legend, though. The legend should allow a reader to assess the information in the figure without having to plough through accompanying text to find more detailed discussion.


**Table 4.3** The mean ( $\pm$  s.e.) percentage number of eggs broken during the experimental period by three strains of battery hen (1–3) under three different housing conditions (A–C). Parametric two-way analysis of variance shows a significant effect of both strain ( $F = 145.09$ , d.f. = 2,27,  $p < 0.001$ ) and housing ( $F = 103.29$ , d.f. = 2,27,  $p < 0.001$ ) and a significant interaction between the two ( $F = 58.76$ , d.f. = 4,27,  $p < 0.001$ ).  $N = 4$  in each combination of strain and housing condition.

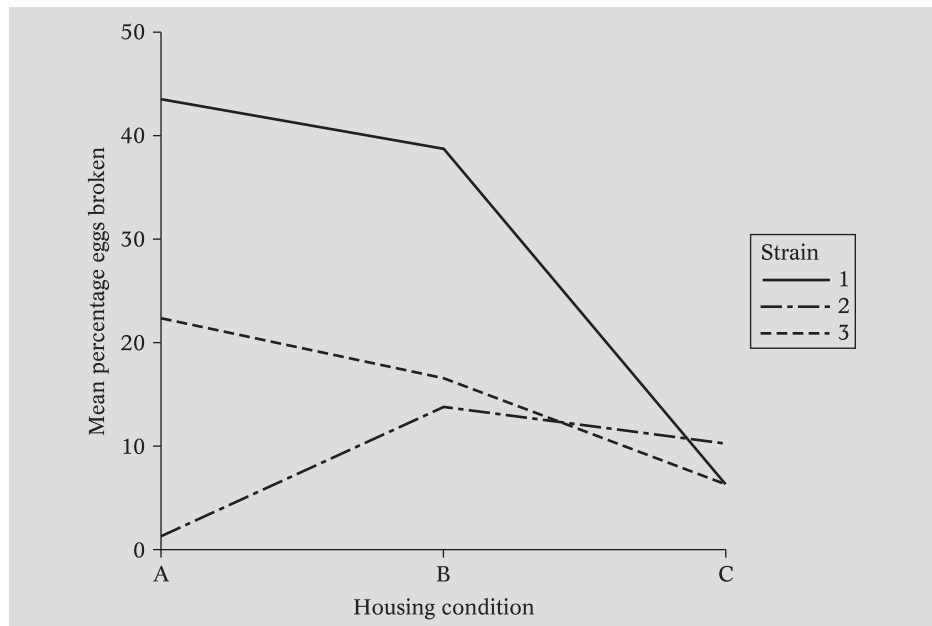
		Strain		
		1	2	3
Housing condition	A	43.50 $\pm$ 2.32	1.25 $\pm$ 0.75	22.25 $\pm$ 2.14
	B	38.75 $\pm$ 1.09	13.75 $\pm$ 1.38	16.50 $\pm$ 1.71
	C	6.25 $\pm$ 0.85	10.25 $\pm$ 1.80	6.25 $\pm$ 1.80

### Differences between two or more groups (with more than one level of grouping)

Here we are concerned with the sort of results that might arise from a two-way analysis of variance or  $n \times n$  chi-squared analysis. Presentation is now a little trickier because of the number of comparisons we need to take into account. A table is probably the simplest solution. For instance, suppose we had carried out a two-way analysis of variance looking at the difference in the frequency of accidental egg damage between three strains of battery hen maintained in three different housing conditions. Here we have two levels of grouping (strain and housing condition) with three groups at each level. The analysis tests for a difference between strains (controlling for housing condition), a difference between housing conditions (controlling for strain) and any interaction between the two levels of grouping (*see* Chapter 3). The best way to present the differences between groups within levels is to tabulate the summary statistics for each of the nine ( $3 \times 3$  groups) cells and include the test statistics in the legend. Thus testing for any difference between groups (i.e. not predicting a difference in any particular direction) might give the results shown in Table 4.3.

This analysis reveals significant effects of both strain and housing conditions on egg breakage. These are obvious from the summary statistics in the table: breakage in strains 1 and 3 is relatively high under housing conditions A and B but drops sharply in condition C. In contrast, breakage in strain 2 is highest in conditions B and C and lowest in A. Damage tends to be greater in types A and B housing than in type C. In addition to these main effects, however, there is also a significant interaction between strain and housing condition (*see* legend to Table 4.3) with the effect of housing differing between strains. Although interaction effects can also be gleaned from a table of summary statistics like Table 4.3, they can be presented more effectively as a figure; one of the levels of grouping constitutes the  $x$ -axis and the measure being analysed is the  $y$ -axis. The relationship between the measure and the  $x$ -axis grouping can then be plotted for each group in the second level. Figure 4.5 shows such a plot for the interaction in Table 4.3. The lines in the figure, of course, simply indicate the groups of data: they are in no way comparable with statistically fitted lines. Full details of the analysis are given in the legend because such a figure would not normally be presented as well as the summary table since it repeats information already given





**Figure 4.5** The mean percentage number of eggs broken by three strains of battery hen (solid, strain 1; dash/dot, strain 2; dotted, strain 3) in three different housing conditions (A–C). Parametric two-way analysis of variance showed a significant effect of both strain ( $F = 145.09$ , d.f. = 2,27,  $p < 0.001$ ) and housing condition ( $F = 103.29$ , d.f. = 2,27,  $p < 0.001$ ) and a significant interaction between the two ( $F = 58.76$ , d.f. = 4,27,  $p < 0.001$ ).  $N = 4$  in each combination of strain and housing condition.

in the table. The graph should have the s.e.'s plotted with the mean values, but it does tend to become very cluttered. Judgement is required as to whether a table or a figure is the best presentation. From Fig. 4.5 it is clear that, while all three strains show differences in egg damage across housing conditions, the direction and degree of decline are different in different strains. This implies that the effect of housing condition varies with strain, which is what is meant by an interaction between housing condition and strain.

With an  $n \times n$  chi-squared analysis, we are just dealing with total counts in each cell so there are no summary statistics to calculate and present. The simplest presentation is thus an  $n \times n$  table with each cell containing the observed and expected values for the particular combination of groups (the expected value in each cell usually goes in brackets). Table 4.4 shows such a presentation for a

**Table 4.4** The number of seeds germinating in a tray in relation to temperature (low, 5 °C; high, 25 °C) and soil type. Expected values in brackets.  $\chi^2 = 14.38$ , d.f. = 1,  $p < 0.001$ .

	Number of seeds germinating	
	In clay soil	In sandy soil
Low temperature	40 (57.97)	100 (82.03)
High temperature	131 (113.03)	142 (159.97)



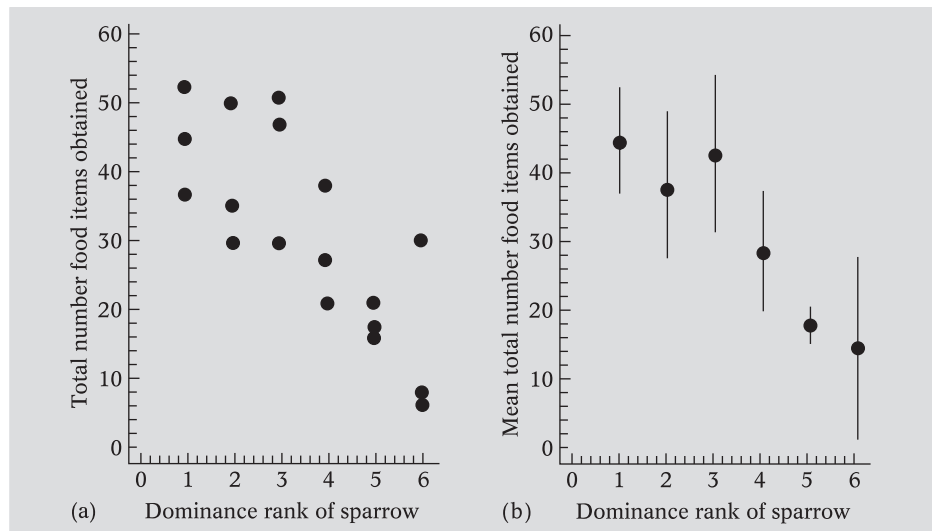
chi-squared analysis of the effects of temperature and soil type on the number of seeds out of 150 germinating in a seed tray.

#### 4.1.2 Presenting analyses of trends

Presenting trend analyses is rather simpler because, in most cases, a scattergram with or without a fitted line is the obvious format. When it comes to more complicated trend analyses that deal with lots of different measures at the same time, it is usually possible to present the various significant relationships as so-called *partial* correlation or regression plots. These can be selected within the regression analysis procedures of SPSS, for example. Alternatively, summary tables rather than figures may be necessary.

##### Presenting a correlation analysis

Since correlation analysis does not fit a line to data points, presentation consists simply of a scattergram, though depending on how we have replicated observations this may include some summary statistics (*see below*). Information about test statistics, sample sizes and significance could be given in the figure, but it is more usual to include it in the legend. Thus Fig. 4.6a shows a plot of the number of food items obtained by male house sparrows (*Passer domesticus*) in relation to their dominance ranking with other males in captive flocks of six (rank 1 is the most dominant male that tends to beat all the others in aggressive disputes and rank 6 is the least dominant that usually loses against everyone else). In this case, observations were repeated for three sets of males so there are three separate points (*y*-values) for each *x*-value in the figure.

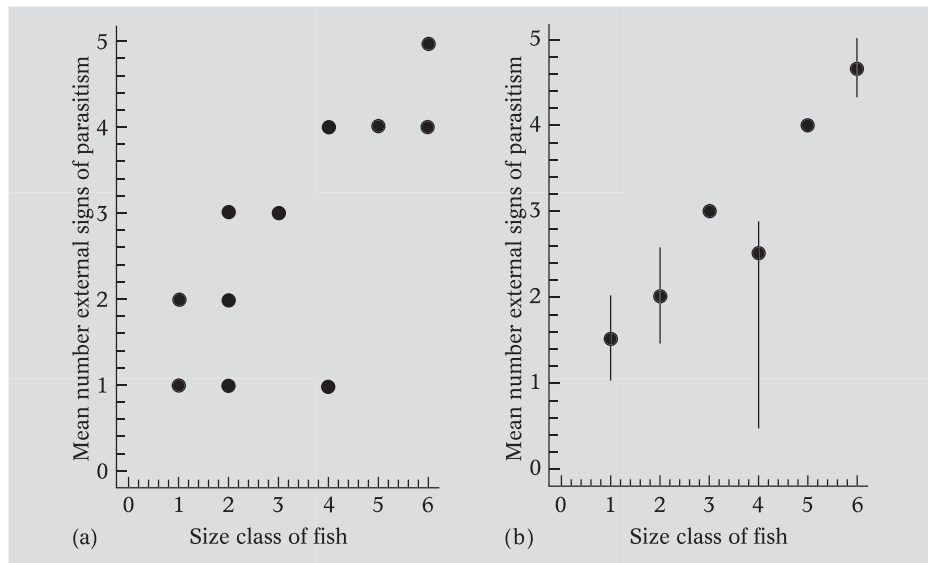


**Figure 4.6** (a) The number of food items obtained during the period of observation by male house sparrows of different dominance status in groups of six (rank 1, most dominant; rank 6, least dominant, data for three groups at each rank).  $r_s = -0.77$ ,  $n = 18$ ,  $p < 0.001$ . (b) The mean number of food items obtained during the period of observation by male house sparrows of different dominance status in groups of six (rank 1, most dominant; rank 6, least dominant).  $r_s = -0.77$ ,  $n = 18$ ,  $p < 0.001$ . Bars represent standard errors.

Although there is a significant trend towards dominant males getting more food, the correlation is negative because we chose to use a rank of 1 for the most dominant male and a rank of 6 for the least dominant. Rankings are frequently ordered in this way, leading to the slightly odd situation of concluding a positive trend (e.g. dominants get more food) from what looks like a negative trend (the number of food items decreases with increasing rank number). There is no reason, of course, why dominance shouldn't be ranked the other way round (6 = most dominant, 1 = least dominant) so that a positive slope actually appears in the figure.

Sometimes when replicated observations are presented in a scattergram, they are presented as a single mean or median with appropriate standard error or confidence limit bars. Thus an alternative presentation of Fig. 4.6a is shown in Fig. 4.6b. Note that a different explanatory legend is now required because the figure contains different information.

In some cases, replication may not occur throughout the data set. Say we decided to sample a population of minnows (*Phoxinus phoxinus*) in a stream to see whether big fish tended to have more parasites. To avoid the difficulties of making accurate measurements of fish size in the field and possibly injuring the fish, we visually assess those we catch as belonging to one of six size classes. We then count the signs of parasitism on them and return them to the water. Because we have no control over the number of each size class we catch, we end up with more samples for some classes than for others. When we come to present the data, we could present them as individual data points for each size class (Fig. 4.7a) or condense replicated data for classes to means or medians (Fig. 4.7b).



**Figure 4.7** (a) The relationship between the size of minnows (arbitrary size classes) and the number of signs of parasitic infection observed on them.  $r_s = 0.74$ ,  $n = 11$ ,  $p < 0.02$ . (b) The relationship between the size of minnows (arbitrary size classes) and the mean number of signs of parasitic infection observed on them.  $r_s = 0.74$ ,  $n = 11$ ,  $p < 0.02$ . Bars represent standard errors.

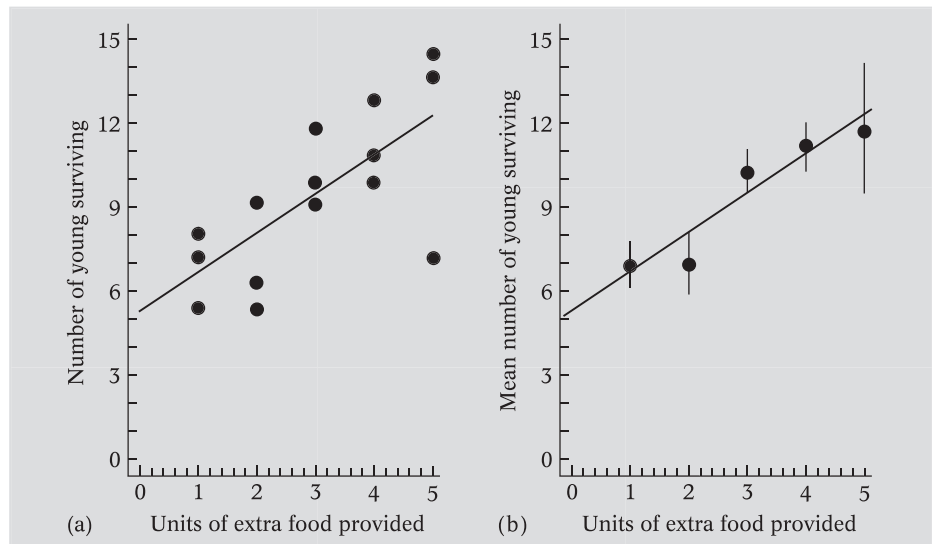
In the latter case, only some points might have error or confidence limit bars attached to them because only some classes are replicated (*see* Fig. 4.7a). Data for those classes that are not replicated are still presented as single points.

It is, of course, important to remember that even where correlations are presented as mean or median values rather than independent data points, the correlation analysis itself (i.e. the calculation of the correlation coefficient) is still performed on the independent data points, not on the means or medians. Values of  $n$  are thus the same in Figs 4.6a and 4.6b and in Figs 4.7a and 4.7b. Correlations *can* be performed on summary statistic values, but obviously a lot of information is lost from the data and  $n$ -sizes are correspondingly smaller.

### Presenting a regression analysis

Presenting a regression analysis is essentially similar to presenting a correlation except that a line needs to be fitted through the data points. If the trend isn't significant, so that a line should not be fitted, a figure probably isn't necessary in the first place. The details of calculating a regression line have been given earlier. You may sometimes come across regression plots that show confidence limits as curved lines above and below the regression line itself. However, we shall not be dealing with these here. For further information, see Sokal and Rohlf (1995).

As with correlation, data can be presented as independent points or, where replicated for particular  $x$ -values, as means or medians. Once again, where means or medians are presented, significance testing and the fitting of the line are still done using the individual data points, not the summary statistics. Figure 4.8 presents a regression of the effect of additional food during the breeding season on



**Figure 4.8** (a) The number of chicks surviving to their first winter in relation to the number of units of extra food provided during the breeding season in three populations of moorhen.  $F = 13.27$ , d.f. = 1,13,  $p < 0.01$ . (b) The mean number of chicks surviving to their first winter in relation to the number of units of extra food provided during the breeding season in three populations of moorhen.  $F = 13.27$ , d.f. = 1,13,  $p < 0.01$ . Bars represent standard errors.

the number of young moorhens (*Gallinula chloropus*) surviving into their first winter in three study populations. Five different quantities of food were used and the three populations received them in a different order over a five-year experimental period. In Fig. 4.8a, the numbers for each population are presented separately; in Fig. 4.8b they are presented as means ( $\pm$  s.e.) across the three study populations.

## 4.2 Presenting results in the text

So far in this section, we have assumed that results will be presented as figures or tables. Figures and tables, however, take up a lot of space in a report and may not be justified if the result is relatively minor or there is a strict limit on the length of the report. In such cases, analyses can instead be summarised in parentheses in the text of the 'Results' section (*see later*). The usual form for a difference analysis is to quote the summary statistics, test statistic, sample size or degrees of freedom and  $p$ -value. Thus the information in Table 4.1a could easily be presented in the text as:

Treatment with gibberellin resulted in a significant increase in growth compared with non-treated controls (mean ( $\pm$  s.e.) height of treated plants =  $14.75 \pm 0.88$  cm,  $n = 12$ ; mean height of controls =  $9.01 \pm 0.63$  cm,  $n = 8$ ;  $H = 12.22$ ,  $p < 0.001$ ).



For a trend, it is usual to quote just the test statistic, sample size or degrees of freedom and the  $p$ -value. For instance, the information in Fig. 4.8a could be summarised as follows:

The number of chicks hatching during a breeding season that survived into their first winter increased significantly with the amount of extra food provided within the population ( $F = 13.27$ , d.f. = 1,13,  $p < 0.01$ ).

It is impossible to generalise about when an analysis could be presented in the text rather than in separate figures or tables. Sometimes, as we have said, it is simply a matter of limited space. However, rough guidelines might include the following: (a) difference analyses between only two or three groups; (b) corroborative analysis, supporting a main analysis already presented as a figure or table (for instance, if a main analysis showed a significant correlation between body size and fighting ability, a corroborative analysis might check that body size was not confounded with age and that the correlation could not have arisen because bigger individuals had more experience of fighting); (c) analyses providing background information (e.g. showing a significant sex difference in body size where this is germane to, say, an analysis of the diet preferences of the two sexes).

### Units

Whether we are dealing with data in tables, figures or text, it is *essential* that appropriate units of measurement are included and cited with their conventional abbreviations (Box 4.1). Summary statistics are meaningless without them.

## 4.3 Writing reports

Just as figures and tables of data should be presented properly to ensure they are effective, so care must be taken in the text of a report. In the scientific community, reports of experiments and observations are usually published in the form of papers in professional journals, or sometimes as chapters in specialist books. In all cases, however, the aim is both to communicate the findings of a piece of research and provide the information necessary for someone else to repeat the work and check out the results for themselves. Both these elements are crucial and, as a result, scientific papers are usually refereed by other people in the same field to make sure they come up to scratch before being published. Not surprisingly, a more or less standard format for reports has emerged which divides the textual information into well-recognised sections that researchers expect to see and know how to refer to to find out about different aspects of the work. Learning to use this format properly is one of the most important goals of any basic scientific training. We shall therefore now discuss the general structure of a report and what should and should not go in each of its sections; then we shall develop a full report, incorporating our various points about text and data presentation, from some of our main example observations.

**Box 4.1**
**Some common units of measurement and their conventional abbreviations**
**Length**

Kilometre	km
Metre	m
Centimetre	cm
Millimetre	mm

**Area**

Square kilometre	km <sup>2</sup>
Hectare	ha
Square metre	m <sup>2</sup>
Square centimetre	cm <sup>2</sup>
Square millimetre	mm <sup>2</sup>

**Volume**

Cubic decimetre (≡ Litre)	dm <sup>3</sup> (l)
Cubic centimetre (≡ Millilitre)	cm <sup>3</sup> (ml)
Cubic millimetre (≡ Microlitre)	mm <sup>3</sup> (μl)

**Weight**

Kilogram	kg
Gram	g
Milligram	mg
Nanogram	ng

**Time**

Million years	My
Years	y
Hours	h
Minutes	min
Seconds	s

### 4.3.1 The sections of a report

There are five principal sections in a report of experimental or observational work: *Introduction*, *Methods*, *Results*, *Discussion* and *References*. Sometimes it is helpful to have some small additional sections such as *Abstract*, *Conclusions* and *Appendices*, but we shall deal with these later.

### Introduction

The Introduction should set the scene for all that follows. Its principal objective is to set out: (a) the *background* to the study, which means any theoretical or previous experimental/observational work that led to the hypotheses under test. Background information is thus likely to include references to previously published work and sometimes a critical review of competing ideas or interpretations. It might also include discussion about the timing (seasonal, diel, etc.) of experiments/observations and, in the case of fieldwork, the reasons for choosing a particular study site; (b) a clear statement of the hypotheses and predictions that are being tested; and (c) the rationale of the study, i.e. how its design allows the specified predictions to be tested and alternatives to be excluded. The Introduction should thus give the reader a clear idea as to why the study was carried out and what it aimed to investigate. The following is a brief example:

Reptiles are ectotherms and thus obtain most of the heat used to maintain body temperature from the external environment (e.g. Davies, 1979). Rattlesnakes (*Crotalus* spp.) do this by basking in the sun or seeking warm surfaces on which to lie (Bush, 1971). An increased incidence of snake bites in the state over the past two years has been attributed to a number of construction projects that have incidentally provided rattlesnakes with concrete or tarmac surfaces on which to bask (North, 2006). The aim of this investigation was to study the effect of the construction projects on basking patterns among rattlesnakes to see whether these might increase the exposure of people to snakes and thus their risk of being bitten. The study tests two hypotheses: (a) concrete and tarmac surfaces are preferred basking substrates for rattlesnakes and (b) such surfaces result in a higher than average density of snakes near humans.

When the reader moves on to the Methods and Results sections, they will then appreciate why things were done the way they were. Reading Methods or Results sections without adequate introductory information can be a frustrating and often fruitless business since the design of an experiment or observation usually makes sense only in the context of its rationale. As we shall see below, it is sometimes appropriate to include background material in the Discussion section, but in this case it should be to help develop an interpretation or conclusion, not an afterthought about information relevant to the investigation as a whole; if it is the latter, it should be in the Introduction.

### Methods (or Materials and Methods)

The Methods section is perhaps the most straightforward. Nevertheless, there are some important points to bear in mind. Chief among them is providing enough detail for someone else to be able to repeat what you did *exactly*. Clearly, the precise detail in each case will depend on the investigation, but points that need attention are likely to include the following.

***Experimental/observational organisms or preparations.*** The species, strain, number of individuals used, growth or housing conditions and husbandry, age and

sex, etc. for organisms; the derivation and preparation and maintenance techniques, etc. for preparations (e.g. cell cultures, histological preparations, pathogen inoculations).

**Specialised equipment.** Details (make, model, relevant technical specifications, etc.) of any special equipment used. This usually means things like tape or video recorders, specialist computing equipment, spectrometers, oscilloscopes, automatic data-loggers, optical equipment such as telescopes, binoculars or specialised microscopes, centrifuges, respirometers, specially constructed equipment such as partitioned aquaria, choice chambers, etc. Run-of-the-mill laboratory equipment like glassware, balances, hotplates and so on don't usually require details, though the dimensions of things like aquaria or other containers used for observation and the running temperature of heating devices, etc. should be given.

**Study site (field work).** Where an investigation has taken place in the field, full details of the study site should normally be given. These should include its location (e.g. grid reference) and a description of its relevant features (e.g. size, habitat structure, use by man) and how these were used in the investigation. The date or time of year of the study may also be relevant.

**Data collection.** This should include details of all the important decisions that were made about collecting data. Again, it is impossible to generalise, but the following are likely to be important in many investigations: any pretreatment of material before experiments/observations (e.g. isolation of animals, drug treatment, surgical operations, preparation of cell cultures, staining); details of experimental/observational treatments *and controls*; sample sizes and replication; methods of measurement and timing; methods of recording (e.g. check sheets, tape recording, tally counters, etc.); duration and sequencing of experimental/observational periods; details of any computer software used in data collection. Of course, it is important not to go overboard. For instance, it isn't necessary to relate that a check sheet was ticked with a red ball-point pen rather than a black one, but if the pen was used to stimulate aggression in male sticklebacks (which often attack red objects) then it would be relevant to state that a ball-point pen was used and that it was red.

## Results

The Results section is in some ways the most difficult to get right. Many students regard it as little more than a dumping ground for all manner of summary and, worse, raw data. Explanation, where it exists at all in such cases, frequently consists of an introductory 'The results are shown in the following figures . . .' and a terminal 'Thus it can be seen . ..'. A glance at any paper in a journal will show that a Results section is much more than this. At the other extreme, explanation within the Results often drifts into speculative interpretation, which is more properly the province of the Discussion (*see below*).

A Results section should do two things and *only* two things: first, it should present the data (almost always in some summarised form, of course) necessary to answer the questions posed; and second, it should explain and justify the

analytical approach taken so that the reasons for choice of test and modes of data presentation are clear. The section should thus include a substantial amount of explanatory text, but explanation should be geared solely to the analyses and presentation of data and not the interpretations or conclusions that might be inferred from them. An example might be as follows:

Figure 1 shows that rattlesnakes are significantly more likely to be found on concrete (Fig. 1a) and tarmac (Fig. 1b) surfaces around dawn and dusk than around midday. Since many of the construction projects in the survey of snake bite incidence have involved highways (Greenbaum *et al.*, 1984), this temporal pattern of basking may result in highest snake/human encounter at times when public conveniences are closed and motorists are forced to relieve themselves at the roadside. Indeed Table 1 shows a strong association for three highways between time of day and number of motorists stopping by the roadside.

It is also important that all the analyses and presentations of data involved in the report appear in the Results section (as figures, tables or in the text) and only in the Results section; no analysis should appear in any other section.


### Discussion

The Discussion is the place to comment on whether the results support or refute the hypotheses under test and how they relate to the findings of other studies. The Discussion thus involves interpretation and reasonable speculation, with further details about the material investigated and any corroborative/contradictory/background information as appropriate. As we have said, however, while the Discussion may flesh out, comment, compare and conclude, it should not bring in new analysis. Neither should it develop background information that is more appropriate to the Introduction (*see earlier*). The kind of thing we'd expect might be as follows:

The results suggest that concrete and tarmac surfaces are not favoured for basking by rattlesnakes in comparison with broadly equivalent natural surfaces when relative area is taken into account. One reason for this might be the greater proximity and greater density of cover close to the natural surfaces sampled. Many snakes (Jones, 1981), including rattlesnakes (Wilson, 1998), prefer basking areas within a short escape distance of thick cover. Despite not being preferred by snakes, the greater incidence of bites on concrete and tarmac surfaces can be explained in terms of the greater intensity of use of these surfaces by humans. However, Wilson (1998) has noted that the probability of attack when a snake is encountered increases significantly if there is little surrounding cover. The paucity of cover around the concrete and tarmac samples may thus add to the risk of attack in these environments.

### References

Your report should be referenced fully throughout, with references listed chronologically in the text and alphabetically in a headed Reference section at the end.



References styles vary enormously between different kinds of report so there is no one accepted format. However, a style used very widely is illustrated below and we suggest using it except where you are explicitly asked to adopt a different style. In this style, references in the text should take the form:

... Smith (1979, 1980) and Grant *et al.* (1989) claim that, during a storm, a tree 10 m in height can break wind for over 100 m (*but see* Jones and Green, 2002; Nidley, 1999, 2001) ...

In the References list at the end, journal references take the form:

Grant, A. J., Wormhole, P. and Pigwhistle, E. G. (1989) Tree lines and the control of soil erosion. *Int. J. Arbor.* **121**, 42-78.

Jones, A. B. and Green, C. D. (2002) Soil erosion: a critical review of the effect of tree lines. *J. Plant Ecol.* **97**, 101-107.

Smith, E. F. (1979) Planting density and canopy size among deciduous trees. *Arbor. Ecol.* **19**, 27-50.

Smith, E. F. (1980) Planting density and growth rate among deciduous trees. *Arbor. Ecol.* **20**, 38-52.

author      year      title of paper      full or abbreviated journal title      volume no.      inclusive page nos.

for books they take the form:

Nidley, R. (1999) *Deforestation and its impact on national economies.* Hacker Press, London.

title of book in italics      publisher      place of publication

and for chapters in edited volumes the form:

Nidley, R. (2001) Economic growth and deforestation. In *Sustainable economics and world resources*, eds A. B. Jones and C. D. Green, pp. 64–78. Hacker Press, London.

Where more than one source by a particular author (or set of authors) in a particular year is referred to, the sources can be distinguished by using lower case letter suffixes, e.g. (Smith, 1976a, b) indicates that you are referring to two reports by Smith in the year 1976. The order in which you attribute a, b, c, etc. is determined by the order in which you happen to refer to the publications in your report, not the order in which they were published in the relevant year.

**Personal observations and personal communications.** Although most of the references you will make will be to work by other people, or yourself, that has been published in some form, it is occasionally appropriate to refer to unpublished observations. This usually arises where some previous, but unpublished, observation is germane to an assumption, fact, technique, etc. that you are relying on in your own report. If such observations are your own, they can be referred to in the text as '(personal observation)' or '(pers. obs.)'. If they have been reported to you by someone else, then they can be referred to as, for example, '(P. Smith, personal communication)' or '(P. Smith, pers. comm.)' – note that the name of the person providing the information is given as well.

**Abstract.** A small, but important, section of many scientific reports, certainly published ones like papers in learned journals, is the *abstract*. This is a short (often strictly word-limited) summary of the aims and main findings of the investigation. The idea is to provide the reader with a quick overview of what was done and what was interesting about it, so that the reader can decide whether they want to read the report in more detail. Abstracts are particularly important in the case of published reports because they are often made available online to people browsing the various searchable scientific literature databases (*see* Box 1.1). They are thus a useful 'shop window' for available studies on the chosen topic. Increasingly, as part of the general push for greater public awareness of science, abstracts are now also being made available in the form of 'lay summaries', meaning that they are redrafted in simple, everyday language that people without formal scientific training can understand (*see* 4.4 below); often they are sent out to the media. Whether or not you intend to try to publish your report, however, producing an abstract for it is good practice because it makes you think clearly about the important messages in your work and express them succinctly. An abstract is included in the example report in Box 4.2 to illustrate the point.

### Other sections of a report

In some cases, there may be additional sections to a report.

**Conclusions.** Sometimes, especially where analyses and interpretations are long and involved, it is helpful to highlight the main conclusions in a tail-end section

so that the reader finishes with a reminder of the ‘take-home’ message of the investigation.

**Appendix.** Occasionally, certain kinds of information may be incorporated into an Appendix. Such information might include the details of mathematical models or calculations, detailed background arguments, selective raw data or other aspects of the study that potentially might be of importance to readers but which would clutter up and disrupt the main report were they to be included there. Appendices are thus for informative asides that might help some readers but perhaps distract others. It follows, therefore, that appendices should be used selectively, sparingly and for a clear purpose, not as a dumping ground for odds and ends on the grounds that they might just turn out to be useful.

**Use of abbreviations.** It is also worth saying something about the use of abbreviations. Many long-winded technical and jargon terms are often abbreviated in reports, papers and books. This is common practice and perfectly acceptable, as long as abbreviations are defined at their first point of use and conventions are adhered to where they exist (some acronyms, for example, are so well-established that people are hard put to recall the full terminology). Thus:

The hyperstriatum ventrale pars caudalis (HVC) in the forebrain of birds is associated with the production of song. The volume of the HVC also varies with the complexity of song in different species.

and

To see whether there was any effect of site on the frequency of calling, we carried out a one-way analysis of variance (ANOVA). The results of the ANOVA suggested that site had a profound effect.

present no problem, whereas:

The HVC in the forebrain of birds is associated with the production of song.

or

To see whether there was any effect on the frequency of calling, we carried out an ANOVA.

leaves the uninitiated little the wiser.

While abbreviations and acronyms are acceptable, however, they should be used judiciously. Littering text with them is a sure way to destroy its readability and confuse the reader.

### 4.3.2 Example of a report

Having outlined the general principles of structuring a report, we can finish off by illustrating them more fully in a complete report (see Box 4.2). The report is



**Box 4.2** Example report**The effect of body size on the escalation of aggressive encounters between male field crickets (*Gryllus bimaculatus*)****Abstract**

Fighting is a costly activity; it takes time and energy and risks injury or even death. One way animals may be able to reduce the cost of aggressive competition is by assessing their chances of winning before becoming involved in a fight. Various attributes of opponents might provide information about the likelihood of winning, an obvious one being body size. When male field crickets (*Gryllus bimaculatus*) were allowed to interact in a sand-filled arena, encounters were more likely to become aggressive as the difference in body size between opponents declined, suggesting relative body size was important in assessing whether or not to escalate into a fight. However, the results also suggested that experience of winning or losing a fight itself affected the tendency to initiate and win subsequent fights. Aggressive encounters between male crickets may thus depend on both assessment of opponents and the degree of confidence of competing individuals at the time of encounter.

**Introduction**

Fighting is likely to be costly in terms of time and energy expenditure and risk of injury to the individuals involved. We might thus expect natural selection to have favoured mechanisms for reducing the likelihood of costly fights. One way animals could reduce the chance of becoming involved in an escalated fight is to assess their chances of winning or losing against a given opponent before the encounter escalates into all-out fighting. There is now a substantial body of theory (e.g. Parker, 1974; Maynard Smith and Parker, 1976; Enquist *et al.*, 1985) suggesting how assessment mechanisms might evolve and much empirical evidence that animals assess each other during aggressive encounters (e.g. Davies and Halliday, 1978; Clutton-Brock *et al.*, 1979; Austad, 1983). Since the outcome of a fight is likely to be determined by some kind of difference in physical superiority between opponents, features relating to physical superiority might be expected to form the basis for assessment.

Male field crickets compete aggressively for ownership of shelters and access to females (see Simmons, 1986). Casual observation of male crickets in a sand-filled arena suggested that body size might be an important determinant of success in fights, with larger males winning more often (pers. obs.). This is borne out by Simmons (1986), who found a similar effect of body size in male *G. bimaculatus*. Observations also showed that aggressive interactions progressed through a well-defined series of escalating stages (see also Simmons, 1986, and e.g. Brown *et al.*, 2006 and Nasil, 2002 for other cricket species) before a fight ensued. One possibility, therefore, is that these escalating stages reflect the acquisition of information about relative body size and interactions progress to the later, more aggressive, stages only when opponents are closely matched in size and the outcome is difficult to predict. This study therefore tests two predictions arising from this hypothesis:

1. large size will confer an advantage in aggressive interactions among male crickets, and
2. interactions will escalate further when opponents are more closely matched in size.

### Methods

Four groups of six virgin male crickets were used in the experiment. All males were derived from separate, unrelated stock colonies a week after adult eclosion so each group comprised arbitrarily selected, unfamiliar males on establishment. Crickets were maintained on a 12 h : 12 h light : dark cycle that was shifted by 4 h to allow observation at periods of peak activity (Simmons, 1986). Before establishing a group, the width of each male's pronotum (thorax) was measured at its widest point using Vernier calipers and recorded as an index of the male's body size (the pronotum was chosen because it consists of relatively inflexible cuticle that is unlikely to vary between observations or with handling; adult body size is determined at eclosion so does not change with age). The dorsal surface of the pronotum of each male was then marked with a small spot of coloured enamel paint to allow the observer to identify individuals.

Groups were established in glass arenas (60 × 60 × 30 cm) with 2-cm deep silver sand substrate. Each arena was provided with water-soaked cotton wool in a Petri dish and two to three rodent pellets. No shelters or other defensible objects were provided to avoid bias in the outcome of interactions due to positional advantages. Arenas were maintained under even 60 W white illumination in an ambient room temperature of 25 °C throughout the experiment.

The six males in a group were introduced into their arena simultaneously and allowed to settle for 5 min. They were then observed for 30 min, during which time all encounters between males were dictated on to magnetic tape, noting: (a) the individuals involved, (b) the individual initiating the encounter (the first to perform any of the components of aggressive behaviour – *see below*), (c) the individual that won (decided when one opponent first attempted to retreat) and (d) the components of aggressive behaviour used by each opponent during the encounter. Following Simmons (1986), the aggressive behaviours recognised here are shown in Table 1.

### Results

*Do larger males tend to win aggressive encounters?* To see whether larger males tended to win more often, the percentage of encounters won by each male in the four groups was


**Table 1** Degree of escalation increases from Aggressive stridulation to Flip. Each behaviour can thus be ascribed a rank escalation value ranging from 1 (low escalation) to 6 (high escalation).

Behaviour	Description	Escalation ranking
Aggressive stridulation	One or both males stridulate aggressively. This may occur on its own or in conjunction with other aggressive behaviours	1
Antennal lashing	One male whips his opponent with his antennae	2
Mandible spreading	One male spreads his mandibles and displays them to his opponent	3
Lunge	A male rears up and pushes forward, butting the opponent and pushing him backwards	4
Grapple	Males lock mandibles and wrestle	5
Flip	One male throws his opponent aside or onto his back. Re-engagement was rare following a Flip	6


plotted against pronotum width (Fig. 1). A significant positive trend emerged. Figure 1, however, combined data from all four groups. Did the relationship hold for each group separately? Spearman rank correlation showed a significant relationship in three of the four groups ( $r_s = 0.94, 0.99, 0.97$  ( $p < 0.05$  in all cases) and  $0.66$  (ns),  $n = 6$  in all groups, one-tailed test).

If there is a size advantage as suggested by Fig. 1, we might expect larger males to initiate more encounters than smaller males since they have more to gain. Figure 2 shows a significant positive correlation between pronotum width and the percentage of the recorded encounters for each male that was initiated (see Methods) by that male. As expected, therefore, larger males tended to be the initiator in more of their encounters.

One possibility that arises from Figs 1 and 2 is that the apparent effect of body size was an incidental consequence of the tendency to initiate. There may be an advantage to initiating itself, perhaps because an individual initiates only when its opponent's ability to retaliate is compromised (e.g. it is facing away from its attacker). If the males doing most of the initiating in the groups just happened to be the bigger ones, the initiation could underlie the apparent effect of body size on the chances of winning. To test this, the percentage encounters won by each male when he was the initiator was compared with the percentage won when he was not. The



**Figure 1** The relationship between the size (pronotum width) of a male and the percentage of encounters won by the male in a group of six.  $r_s = 0.81$ ,  $n = 24$ ,  $p < 0.0001$ .

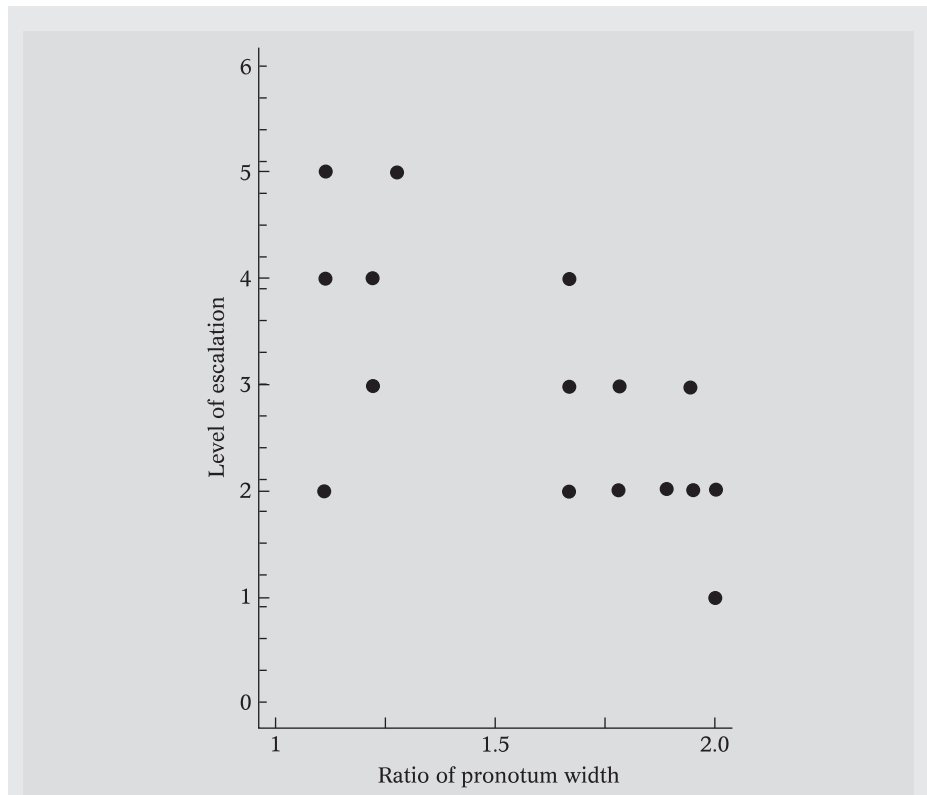


**Figure 2** As Figure 4.9 but for the relationship between male size and the percentage number of encounters initiated by the male.  $r_s = 0.63$ ,  $n = 24$ ,  $p < 0.03$ .

analysis showed no significant difference ( $U = 82$ ,  $n_1 = 19$ ,  $n_2 = 19$ , ns\*) between the two conditions.

*Does difference in body size affect the degree of escalation in encounters?* Figure 3 shows the relationship between the ratio of pronotum width for pairs of opponents and the degree of escalation of their encounters. A ratio of one indicates equal size and ratios greater than one increasing departure from equality. Degree of escalation is measured as the maximum rank value (1–6, see Methods, Table 1) recorded during an encounter. As predicted, the figure shows a significant negative correlation between size ratio and degree of escalation so that escalated encounters were more likely between opponents that were closely matched in size. The trend in Fig. 3 is for data across all groups. Does the trend hold within individual males? Correlation analysis for those males (four) that were involved in five or more encounters with opponents of different relative size suggests that it did, although the trends were significant in only two cases ( $r_s = -0.96$ ,  $n = 7$ ,  $p < 0.05$ ;  $r_s = -0.72$ ,  $n = 5$ , ns;  $r_s = -0.76$ ,  $n = 6$ , ns;  $r_s = -0.99$ ,  $n = 6$ ,  $p < 0.05$ , one-tailed

\* Although it is perfectly legitimate to use a Mann–Whitney  $U$ -test here, the fact that we are actually comparing data for the two conditions (initiated versus non-initiated encounters) *within* males means we could have used a different sort of two-group difference test (e.g. a Wilcoxon matched-pairs signed ranks test) which takes this into account.



**Figure 3** The relationship between ratio of pronotum widths of fighting males and the maximum level of escalation (1–6, see Methods) reached in fights.  $r_s = -0.71$ ,  $n = 20$ ,  $p < 0.002$ .

test). A similar analysis of the relationship between size ratio and the duration of aggressive encounters also showed a significant tendency for closely matched males to fight for longer ( $r_s = 0.62$ ,  $n = 20$ ,  $p < 0.002$ ).

*Does the experience of winning or losing affect subsequent interactions?* The results so far are consistent with males assessing each other on the basis of relative size. However, it is possible that the experience of winning or losing a fight might itself influence a male's approach to subsequent encounters. A male that has just won a fight, for example, might assess his chances of winning the next one as being higher than if he had just lost (a 'confidence' effect). To see whether this was the case, the outcome of fights for males that were recorded as having ten or more encounters during the observation period were analysed according to whether the male won or lost his first recorded encounter. A Mann–Whitney  $U$ -test showed that males winning their first encounter won a significantly greater proportion of their subsequent encounters than those losing their first encounter ( $U = 8$ ,  $n_1 = 6$ ,  $n_2 = 9$ ,  $p < 0.05$ ). They also initiated a significantly greater proportion of the encounters ( $U = 9$ ,  $n_1 = 6$ ,  $n_2 = 9$ ,  $p < 0.05$ ). Interestingly, there was no significant difference in body size between males winning versus losing their first encounter ( $U = 16$ ,  $n_1 = 6$ ,  $n_2 = 9$ , ns), so the 'confidence' effect seemed to be to that extent independent of the effect of size.

## Discussion

The results bore out both predictions about the effects of body size on the outcome of aggressive interactions between male field crickets: larger males were more likely to win and escalation was more likely between closely matched opponents. This is consistent with the outcome of fights being largely a matter of physical superiority and with the structuring of interactions into a well-defined series of escalating stages reflecting assessment.

The fact that larger males were more likely to initiate an interaction could mean that the relative size of a potential opponent is assessable in advance of physical interaction. However, it could also reflect a general confidence effect arising from previous wins by larger males (males may initiate according to the simple decision rule 'if I won in the past, I'll probably win this time, so it is worth initiating'). Indeed, Simmons (1986) presents evidence that the number of past wins has a positive influence on the tendency for males to initiate, a result consistent with the apparent confidence effect in the present study (but see e.g. Brown *et al.*, 2006 and Nosil, 2002 for more equivocal results from other species). Alternatively, initiation could reflect individual recognition, with males picking on those individuals against whom they have won in the past. Since this study did not record encounters independently of the performance of one of the categories of aggressive behaviour, it is not possible to say whether initiations against particular opponents occurred more or less often than expected by chance. Whatever the basis for deciding to initiate, however, there was no evidence that initiation itself conferred an advantage in terms of the outcome.

Although no resources (shelters and females) were available in the arenas, the size advantage in the aggressive interactions recorded here is in keeping with the tendency for larger males to take over shelters and mate successfully with females (Simmons, 1986). While females prefer to mate with males in or near shelters (because these provide good oviposition sites and protection from predators), they will mate with males encountered in open areas (Simmons, 1986). Aggression between males in the absence of shelters or females may thus reflect an advantage to reducing competition should a female happen to be encountered.

## References

- Austad, S. N. (1983) A game theoretical interpretation of male combat in the bowl and doily spider, *Frontinella pyramitela*. *Anim. Behav.* **19**, 59–73.
- Brown, W. D., Smith, A. T., Moshalik, B. and Gabriel, J. (2006) Aggressive contests in house crickets: size, motivation and information content of aggressive songs. *Anim. Behav.* **72**, 225–233.
- Clutton-Brock, T. H., Albon, S. D., Gibson, R. M. and Guinness, F. E. (1979) The logical stag: adaptive aspects of fighting in red deer (*Cervus elephus* L). *Anim. Behav.* **27**, 211–225.
- Davies, N. B. and Halliday, T. R. (1978) Deep croaks and fighting assessment in toads, *Bufo bufo*. *Nature* **274**, 683–685.
- Enquist, M., Plane, E. and Roed, J. (1985) Aggressive communication in fulmars (*Fulmarus glacialis*) competing for food. *Anim. Behav.* **33**, 1107–1120.
- Maynard Smith, J. and Parker, G. A. (1976) The logic of asymmetric contests. *Anim. Behav.* **24**, 159–175.
- Nosil, P. (2002) Food fights in house crickets, *Acheta domesticus*, and the effects of body size and hunger level. *Can. J. Zool.* **80**, 409–417.
- Parker, G. A. (1974) Assessment strategy and the evolution of animal conflicts. *J. Theor. Biol.* **47**, 223–243.
- Simmons, L. W. (1986) Inter-male competition and mating success in the field cricket, *Gryllus bimaculatus* (de Geer). *Anim. Behav.* **34**, 567–579.

one that might arise from some of the experiments we proposed earlier in the main examples, in this case aggression in crickets.

## 4.4 Writing for a more general readership

So far, we've looked at writing reports on the assumption that the target readership is other biologists, whether these are lecturers or tutors in an educational setting, or fellow researchers. The key concerns with this kind of writing are: (a) adhering to the appropriate conventions of scientific reporting, and (b) the technical clarity with which the hypotheses and predictions under test, the manner in which the study was carried out, the statistical analyses and the background and conclusions to the study are presented for critical professional scrutiny. This can make for a rather dry and jargon-strewn read for anyone coming to it out of general interest, and perhaps with little or no formal scientific education (the proverbial 'intelligent layman'). Nevertheless, there are many good reasons why such a person should be able to appreciate what has been done, and why, in scientific research (not least because they are probably helping to pay for it out of their taxes!), and equally good reasons why scientists themselves should make the effort to render their work accessible rather than leaving it to other people, such as journalists, who often have a poor grasp of the work or try to sensationalise it into a newsworthy story.

In the UK, the art of explaining science to the 'intelligent layman' now has its own buzz phrase, the 'Public Understanding of Science' (or PUS, to use its slightly unfortunate acronym), one enthusiastically promoted by government and honoured with an eponymous chair at Oxford. In its wake, scientists are actively encouraged to engage with the media to promote and explain their work, and many universities and colleges offer courses on popular scientific writing and other public communication skills. Not surprisingly, therefore, science students increasingly see career opportunities in the wider communication of their subject. So what exactly does 'wider communication' entail? Well, of course, it depends to some extent on what you're trying to talk about and to whom, but there are some general points that will help keep you on the right track.

### Avoid jargon

Probably the first golden rule is to avoid, as far as possible, using any technical jargon or unfamiliar scientific terminology. For example, the following passage in a scientific article on immune response to infection:

In a study of European bank voles (*Clethrionomys glareolus*) high burden populations were characterised by high levels of plasma corticosterone and testosterone and high fluctuating asymmetry (FA) in hind foot length (a putative measure of developmental instability; e.g. Palmer and Strobeck, 1986; Møller and Swaddle, 1997), but reduced aggressiveness among males (Barnard *et al.*, 2002, in review).

might be rendered as:

Bank voles are small woodland rodents, rather like mice, but with shorter ears and tails. They are common across Europe where, like other wild mammals, they are often infected with various parasites, such as tapeworms and fleas. In areas where infections are very severe, voles often show evidence of stress and poor physical development, and males are less aggressive than their counterparts elsewhere.

This has managed to recast the whole thing in easily understood, everyday language that most people would be able to follow. It tells them what voles are (many may not know), and what the piece is talking about when it refers to parasites. It also leaves out all the clutter of scientific referencing, essential for the professional scientist, but a distraction for the general reader.

Sometimes, however, the use of a certain amount of technical jargon may be inescapable, or even helpful in giving some idea of the scientific approach behind a report. If so, *explain it* in simple everyday terms. For instance, a different version of the passage above might have gone as follows:

Voles from areas with very high levels of parasite infection show poor physical development. This can be judged by something scientists call ‘fluctuating asymmetry’, which is a measure of the difference in size between the same parts of the body, say the front or hind legs, on the left and right sides. The bigger the difference, the greater the hardship the animal is assumed to have experienced during its development.

This includes the technical term ‘fluctuating asymmetry’, because the reader might wonder exactly how poor physical development was measured, but immediately counters its potential to confuse with a simple description of what it means.

### **Use catchy analogies**

Most people scanning a newspaper or magazine for articles that might interest them need something to draw their attention to a particular piece in the first place, and then to maintain their attention as they read it. A snappy title might achieve the first aim, but ensuring the second can be a bit trickier. A favourite device is to draw parallels with familiar aspects of our own experience. The old adage that ‘sex sells newspapers’, for example, is as true now as it ever was, and anything that reflects on, or can be compared with, our own sexual behaviour, however tangentially, can usually be assured of at least a passing glance. Health, war, culture and intelligence are other good ‘hooks’ for getting attention, because they touch on our preoccupations or views on what makes us uniquely human. Some popular science writers use characters or events from literature or works of art to draw analogies and create striking imagery, Stephen Jay Gould being an enthusiast of the style. Indeed, cleverly appropriate metaphors are an effective way in general to get points and ideas across, as illustrated by Richard Dawkins’s famous book *The Selfish Gene* (Dawkins, 1976), a successful popular account of the principle of natural selection that uses the ‘selfish’ metaphor to characterise the effects of



selection on how organisms behave. As authors like Dawkins have discovered, however, it is often wise to emphasise that you are using metaphors as just that, and not in a literal sense (for instance, the ‘selfish’ gene metaphor simply captures the idea that genes code for characteristics that enhance their chances of being passed on to the next generation; it doesn’t imply they are thoughtlessly self-centred in the sense we might use ‘selfish’ to describe other people). Bearing these points in mind, we might lead into our vole story with something like the following:

The stressed office worker falling prey to every passing infection is a familiar cliché of our pressured, industrialised lives, and good medical research suggests it has a sound basis in fact. But is it something unique to us and our artificially hectic lifestyle? Evidence from the leafy glades of Europe’s woodlands suggests not. A recent study of voles in the Polish forests of Mazury, has revealed an association between stress and infection with parasites that has more than an echo of sick office worker syndrome.

So here we’ve used a familiar everyday analogy to get the reader’s attention and cue them in to the message of what follows; the reader knows to expect something about how stress relates to disease and is thus primed to follow the piece. Box 4.3 presents a general readership piece along these lines on the cricket study reported in formal scientific style in Box 4.2.

## 4.5 Presenting in person: spoken papers and poster presentations

The outcomes and conclusions of your research can also be presented in person in the form of a talk, or in a fashion somewhat intermediate between a written and spoken presentation known as a poster presentation. As with the written word, the audience for a talk or poster can be diverse – from a group of fellow students reporting their research at the end of a field course, to the serried ranks of a thousand or so delegates at an international conference. The major difference from the traditional written word is that you are presenting the information live (poster presentations usually involve sessions where you explain and defend your findings personally to interested readers), and this difference can be enough to test the composure of even the most confident individual. At one level there are as many styles of spoken and poster presentations as there are presenters and, to a lesser extent, each person in the audience will have a slightly different reaction to any given style. However, there are features that are shared by good presentations and we discuss some of these here.

### 4.5.1 A presentation is not a written report

As we explained in section 4.3, writing for a general readership needs a different approach from writing a scientific paper. The same is true of spoken and poster presentations; indeed, they have more in common with writing for a general audience than writing a paper. And while presenting a spoken paper is sometimes

### Box 4.3 A general readership account of fighting in crickets (from Box 4.2)

#### CANNY CRICKETS SIZE UP THE OPPOSITION

As any military commander worth his salt will tell you, ‘know thine enemy’ is a wise rule when it comes to a showdown against a powerful opponent. Rush in on impulse and it probably won’t take long to regret it. This turns out to be as true for other animals as it is for ourselves. And, while they may not be hot on handy aphorisms, our fellow planetary inhabitants can be as wily as any military strategist when it comes to fisticuffs. Take the humble field cricket (*Gryllus bimaculatus*), for example.

Field crickets are distributed widely throughout southern Europe, where males contribute to the ringing chorus of insects that fills the evening air in summer. The reason males put all this effort into calling, however, is not to provide musical accompaniment for human bystanders, but to set about the serious business of competing for females. And here, the males’ tuneful performance is just for starters.

As well as providing a sound beacon by which females can home in on a male, calls are picked up by other, rival, males. Since calling is costly – it takes time and effort, and can attract predators – these males can avoid paying the cost themselves by keeping quiet and intercepting the females attracted by their rival’s efforts. Thus, callers can rapidly find themselves having to compete for their hard-won female with a bunch of freeloaders. Calling aggressively at them may have some effect, but often not much, so it’s not long before things start to get physical. This is where careful choices have to be made.

Observations of crickets in the field, and under laboratory conditions mimicking those in the field, show that males don’t just get stuck in indiscriminately. Interactions follow a well-defined sequence of behaviours that gradually escalate in their level of aggressiveness. If an opponent doesn’t give way when lashed with the aggressor’s antennae, for instance, the encounter might progress to locking jaws and wrestling. Ultimately, the larger of the contesting males tends to win, but the way encounters progressively up the ante strongly suggests males are weighing each other up before going all out. If so, an obvious prediction is that fights should last longer and be more intense as the difference in size between opponents decreases, because it would be harder for each to tell who was the biggest. Some recent experiments by scientists at the University of Nottingham, UK have shown that this is exactly what happens.

When pairs of males were carefully measured and then allowed to compete with each other, they were much more aggressive, and fought for longer, when closely matched in size. Thus, if reliable information about relative size and strength was not forthcoming when opponents prospectively probed each other, there appeared to be nothing for it but to set to. But things weren’t quite as simple as this. While assessing the opposition certainly seemed to be part of the process, it turned out that males were also affected by their previous experience of winning or losing. If they had just won a fight, they were more likely to enter the next one with gusto and win that too. If they’d lost, however, they were more circumspect the next time and tended to lose again. This suggests that the degree of confidence with which a male approaches a fight has an important role to play, and that, in the world of crickets, just as in our own, he who hesitates is lost.

termed *reading a paper*, the result, if taken literally, usually makes for deadly listening for the simple reason that written and spoken language are generally very different in style. Written prose is usually more formal in structure and use of language than the spoken word, which abounds with illustrative figures of speech, amusing asides and informal phraseology that is easy on the ear. For instance, in a published paper we might write:

In summary, therefore: (a) female thargs are on average 50 per cent larger than males, (b) only males over 4 kg in body weight initiate courtship, and (c) only males over 5 kg mate successfully.

But we'd probably *say* something more like:

OK, so what can we conclude from all this? Well for one thing it's clear that females are the larger sex in thargs, which may mean that only relatively big males stand a chance of success in courtship. This may be why it's only 4 kg plus males that attempt courtship and only 5 kg plus males that actually get anywhere with it.

Very few people can write successfully in spoken language, so a paper read out verbatim usually sounds stilted and wooden. One past master of the art of writing in spoken English (in his case) is the evolutionary biologist Richard Dawkins, who customarily reads his talks from a prepared script, but he is a real exception.

The same is really true of poster presentations. One thing a poster absolutely should *not* be is a written paper pinned to a board, however attractively the author may feel the pages are set out. The art of good poster presentation owes more to the skills of the advertising agency than the professional writer: it's all about getting a message over clearly and immediately (people browsing anything from 20–30 upwards to 500 or more posters are not going to hang around reading sheets of dense prose; they want headlines and images they can take in at a glance). We have occasionally seen a scientific paper stapled to a poster board at a conference, but this is more likely a consequence of a catastrophe befalling the original, intended poster than a serious attempt to use a written report as a poster presentation.

## 4.5.2 General considerations

### Know your audience

It should probably go without saying that the content of a presentation will depend on the audience. A more general audience, such as a local natural history society or schoolchildren, will require more background to the study and less technically demanding language than a group of specialists, such as we should find in a research group. As a rule, the more general the audience, the more a presentation should deal in simple (but nonetheless accurate) take-home messages rather than the detailed mechanics of how the messages were arrived at (*see below*). We are not advocating skipping important qualifying information in general presentations, just recommending that care is taken not to lose the

interesting facts and findings in it all. The key is therefore that the pitch and style needs to be geared appropriately for different audiences.

### **Know your allocation**

As well as choosing an appropriate style for your presentation, there is almost always another challenge to be met: it has to fit into a predetermined allocation of time or space. Where talks at conferences are concerned, this can be extremely short, 10 minutes or so sometimes, but may be up to 50 minutes or an hour if you're a key speaker. People often underestimate how long it takes to speak to a slide in a talk, and therefore tend to have too many slides for the time. As a general rule you should allow about two minutes per slide, which means around 10 slides for a 20-minute talk. Remember to talk the audience through each slide, pointing to the relevant pieces of information to which you are referring; don't just put the slide up and talk as if it wasn't there. Poster sizes are usually determined by the size and shape of available boards, and, of course, the number of posters that have to be squeezed onto them. Either way, it is an essential skill of giving spoken or poster presentations that you can tailor them to a required slot. It is thus vital to know the time allotted for your spoken presentation or the dimensions of the poster space. If in doubt (though you shouldn't be), err on the short side – few in the audience will mind if the talk finishes a few minutes early or if your poster doesn't completely fill the poster board. However, taking more than your allocation of time or space will usually elicit negative reactions or result in you being cut off mid-presentation by an irritated session chairperson or not being able to put your poster up at all.

### **Know your aim**

A presentation is sometimes seen as an opportunity (or temptation) to report every last detail of your study. This is a fatal mistake! There is never enough time or space to do it, and nobody's attention span would take it even if there was. Carefully, but ruthlessly, choose the aspects of your study that will be of most interest to the audience. Thus, if you are reporting interesting mating behaviour in your study species, focus on what makes it interesting and perhaps sets it apart from mating behaviour in other, similar species rather than on exactly how many dreary wet days you spent freezing in a remote hide collecting the data, or how you finally arrived at the clever statistical analysis you used. Make sure also that your points can be readily conveyed in the presentation – a personal demonstration of your study species' mating display on a table top could be the highlight of a spoken presentation, but you will need to think of another way of conveying the information if presenting a mostly unattended (by you) poster – a carefully chosen, clear photograph being an obvious option.

### **Less is more**

Although a cliché, the spirit of this little aphorism has probably contributed more to good presentations than almost anything else. Whether the text will appear as a slide in a talk or as a panel on a poster, it is often surprising just how few words are needed to convey the meaning. For example, a first attempt at an introductory slide for a talk about fighting in crickets might set the agenda as follows:

### Introduction

- Fighting is likely to be expensive in terms of time and energy and the risk of injury to individuals
- We might thus expect selection to have favoured mechanisms to reduce the likelihood of costly fights
- One way this could be achieved is by individuals assessing their chances of winning against an opponent before getting involved in a fight
- There is now a substantial body of theory suggesting how assessment mechanisms might evolve, and much evidence that animals do assess each other during aggressive encounters
- Male field crickets (*Gryllus bimaculatus*) compete aggressively for ownership of shelters and access to females, and casual observation suggests that body size might be an important determinant of success in fights

However, this is a lot of text, which will probably take more time than the slide is on the screen to read properly. The point is, you don't need to use full sentences to convey the essential message, and less text means you can use larger fonts and maybe some associated images to make the whole thing more digestible. The following box might be an effective distillation of the full-blown text in the box above.

### INTRODUCTION

- Fighting can be costly
- Selection may thus favour reduced fighting
- Perhaps through assessing opponents first
- We can model the evolution of assessment
- And test models using fighting in crickets

This style is often referred to as *telegraphese*, because messages sent by telegram were charged by the word and shorter messages saved money. (An extreme example was an exchange of telegrams between the nineteenth-century French author Victor Hugo and his publisher: Hugo sent a telegram containing only a question mark, correctly assuming his publisher would understand he was asking how the sales of his recent novel were going – the publisher replied with an exclamation mark, indicating they were extremely good).

Less is also more when it comes to images (like photographs or video windows) and embellishments (like borders, cartoons and coloured text boxes), so the advice is the same as for words: unless an image or embellishment is really central to the message you are trying to convey, or essential for clarity, carefully consider whether it merits inclusion (*see also* Box 4.4). The aim is to design the presentation around the core message, not to produce a kaleidoscope of information for its own sake.

### Images versus text

‘A picture is worth a thousand words’ is another cliché, but again captures the spirit of good presentation. Images and graphics can often replace text or they can enhance and support it, perhaps by illustrating the environment in which the study took place, or the lifestyle of the species concerned. For instance, the detailed list of escalating aggressive behaviours in Table 1 of the report in Box 4.2 could be replaced by a series of three to four photographs of the relevant behaviours with simple accompanying labels. A picture or two of fighting crickets might also helpfully embellish some of the introductory, data or summary slides that would comprise a talk, or figure as background embellishments in a poster.

Tables of data published in papers (for example, like that in Table 4.3) can very rarely be included in presentations as they are. Apart from the difficulty of reading the detail once the table becomes larger than a  $3 \times 3$  matrix, the information is not in the best form for the audience to appreciate the interesting differences and similarities in the data. Converting tables to graphs, such as bar charts or means plots, helps to make such differences clear (*see* Figs 4.1, 4.2 and 4.3).

It is rather less obvious that figures from published papers are also not in the best form for a presentation. Unless they make a key point, figures can take up a lot of space. In some cases, therefore, it may be simpler to present summary results in text form. The figures in the example report in Box 4.2, for example, could, if pushed for time or space, be converted from scattergraphs to the summary text format below:

<b>bigger crickets win more often</b>	$r_s = 0.81 \ n = 24^{***}$
<b>and initiate more fights</b>	$r_s = 0.63 \ n = 24^*$
<b>and similarly sized crickets are more likely to escalate</b>	$r_s = -0.71 \ n = 24^{**}$

However, in general, graphs make differences or trends easier to assimilate at a glance and are probably preferable where constraints allow.

### Legibility

The ease with which poster and projected text can be read is affected by many things, some of which are specific to the venue and the local conditions. However, there are a few basic points worth bearing in mind:

**Font sizes and styles.** For slides used in talks, a sizeable font is needed to ensure text can be read comfortably throughout a lecture venue. Our suggestion would be somewhere in the range of 18 pt (for small rooms) to 28 pt (for large theatres), with 24 pt probably being a good all-round working medium. For posters, 14 pt is probably the minimum to go for; where posters are produced in PowerPoint (*see* Box 4.4), and subsequently blown up to A0 size, the font size on screen is likely to be somewhat smaller.

But font *size* is not the only consideration. Many people have various forms of visual impairment or reading difficulties such as dyslexia. The legibility of text to such people can be greatly affected by the *style* of font chosen. In general, sans serif fonts, like *Arial* or *Verdana*, are much clearer to people with visual impairments or dyslexia than serif fonts such as *Times New Roman*, because of the crisper, cleaner lines of the letters, which give them a sharper boundary with the background.

**Contrast and colour combinations.** Contrast between text and background is obviously another issue, as is the degree of comfort in reading it. Black text on a white background has nice high contrast, but it can be rather boring and harsh on the eye in long presentations. Colour combinations such as pale blue or yellow text on a dark blue background, or yellow text on green can be more comfortable, but be careful to avoid combinations that create problems for colour-blind readers (combinations of red and green being an obvious one to avoid – about 8 per cent of males are red/green colour blind). If you want to check the latter, there are various websites and packages that can simulate the appearance of your text to a reader who is colour-blind (e.g. <http://www.vischeck.com/>).

### Getting attention

In both talks and posters, it is necessary to get the attention of people you want to address. You may have a somewhat more captive audience for a talk, but that doesn't stop people's attention wandering, and, if it wanders sufficiently, they may even walk out. There are various things to think about here. One of the first is your title slide or poster header.

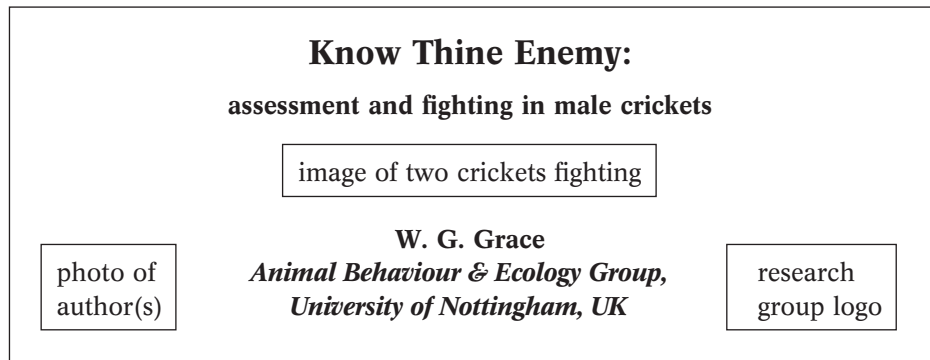
The title of a written paper often tends to be literal and rather pedestrian, though people do liven them up with snappy phrases. Thus, the title of a written paper based on the cricket study in Box 4.2 might be something like.

**The Effect of Body Size on the Escalation of Aggressive  
Encounters between Male Field Crickets *Gryllus bimaculatus***

**W. G. Grace**

**Animal Behaviour & Ecology Research Group, School of Biology,  
University of Nottingham, University Park, Nottingham NG7 2RD, UK**

This is perfectly descriptive of what is to follow, but a lot of text and not very inspiring visually. In the case of a poster, where people are drifting casually past waiting to be hooked by something attractive and interesting-looking, this is unlikely to have much impact. Even in a talk it might suggest a rather dry and cluttered offering is in prospect. Thus it is worth thinking of something catchy as a title and choosing font styles and text layouts that catch the eye. Where fighting in crickets is concerned, something along the lines of the following might do the trick:



Again, some carefully chosen images can be used to brighten it up.

**Gimmicks.** In the scramble for attention, there is often the temptation to descend into gimmickry. The advent of PowerPoint (*see* Box 4.4) has unfortunately brought legion opportunities for this in the slides of spoken presentations, with sound effects, innumerable modes of slide transition, animated cartoons, video windows and many, many more offerings to tempt the over-enthusiastic. While we certainly wouldn't advise eschewing these entirely (some slide transitions, animations and especially the use of video are extremely effective), use them with a great deal of restraint. A gimmicky slide transition with sound effects, for instance, may add a helpful element of surprise the first time, but it will pall to a serious irritation by the third or fourth. In general, the use of gimmicks in slides risks making you look more of a nerd than a slick professional.

To some extent, the case for a bit of carefully judged gimmickry is more arguable in the case of posters, which have to speak out from a background of countless rivals. Most posters are now printed as a single sheet (A0 size) from a file produced by a presentation package such as PowerPoint. However, the advantages of this in terms of ease of mounting and not having to keep track of many separate components can be offset by the greater difficulty of making your poster distinctive. One way in which you can combat this A0 uniformity problem is to have attachments. We have seen convincing recreations of a corner of reed bed (complete with nest and artificial eggs of the study species), clockwork models and miniature sound and video playback systems – all devices that successfully made that poster stand out from the crowd enough to draw us in and begin reading it. But we have also seen efforts (e.g. attached helium balloon banners and clumsy attempts at interactive gadgets) that made the poster look like a tawdry fairground attraction and were decidedly off-putting. Less alternative, and quite useful, extras are contact details, abstracts and preprints of the poster on A4 sheets offered in a manner that can be taken away by interested visitors to the poster – usually from a container attached to the poster.

### Signposting

In a talk, slides are presented in sequence, so you might imagine there is no problem for the audience in following the logical flow of text and images. This is a dangerous assumption. It is easy to be fooled by your own familiarity with the



material into assuming too much on the part of the audience. It is *always* worth having a slide that outlines the structure of your talk and signposts it for people. Thus the famous dictum ‘Stand up, say what you are going to say, say it, then sit down’ is excellent advice. By signposting, of course, we don’t mean listing the conventional sections of a written paper (Introduction, Methods, Results, etc.) Even though your talk may *de facto* follow these sections; you should use the opportunity of a talk structure slide to give a more specific indication of what the talk will be about: for instance, to pursue our cricket example, a structure slide here might go something like:

#### Fighting and assessment in crickets

1. why assess opponents in aggressive contests?
2. aggression in crickets
3. experimental design
4. the evidence for assessment
5. modelling assessment strategies
6. testing the model’s predictions
7. conclusions and the wider view

Again, this can be embellished with a judiciously chosen image in the corner to give it some visual interest. The structure slide can be reprised at various points in the talk to update the audience on where things have got to and where they’re about to go.

A second piece of advice in terms of structuring a spoken presentation is to move on to a new slide or point in an animated list (*see* Box 4.4) only when you are ready to deal with the new material. To judge by the many speakers who move on to the next slide or point before finishing the current one, there is a strong temptation to hurry on. Resist the temptation! The audience’s attention will be taken by the new material, your current point will be lost and confusion is likely to start setting in.

***Signposting in a poster presentation.*** In a poster, text and images tend to be presented in blocks spread over the area of the presentation. Under these circumstances it can sometimes be difficult to determine the intended route through it all. Two obvious ways of overcoming this are to use arrows between successive pieces of text, or to number each piece of text in sequence. Associated images, tables and figures can then be numbered as Plate n, Table n or Figure n accordingly. Numbering blocks of text is probably neater; arrows running through posters can sometimes make them look fussy and cluttered, and you can always use a different coloured font for your numbers to make them stand out.

#### Summaries

In many ways almost the most important part of your presentation is the final impression you leave with your audience/readers. You need to finish with a bit

of a bang in the form of some clear, crisp take-home messages that will stay with those hearing or reading them. Thus, it is vital that you don't just peter out at the end of a rambling discussion and leave the audience/readership to decide for itself what was important about the work. Present a clear *summary* slide or poster panel! The summary should be succinct and memorable, and, in a talk, we strongly advise making it the *last* slide (rather than the common practice of ending with the equivalent of the Acknowledgements section of a written paper). Not only is it then the last slide the audience sees, but it can be displayed for the longest period as it often remains on display during the question session. A summary slide/panel for the study of fighting in crickets might go something like:

### Summary

- Large male *G. bimaculatus* are more likely to win fights
- Fighting is more severe when opponents are closely matched in size
- Both prior assessment and individual confidence influence the decision to fight
- Fighting in the absence of females or reproductive resources may help reduce competition when they are available

### Questions

Questions from your audience/readers are an integral part of both spoken papers and posters and deserve as much consideration as the talk or poster itself. No study is ever the final word, and other people will often have refreshingly new views on how it might be improved or developed. It is often possible to imagine the sort of questions you could be asked, because you will be aware of the study's shortcomings and what remains to be done yourself, and it is well worth thinking through these ahead of your presentation. But by far the best way to prepare for what might come is to do some dummy runs of your talk or poster to groups of people, such as classmates, members of your research group, departmental seminar groups or whatever's handy. This almost always throws up issues you won't have thought of, and is particularly good at highlighting where you're failing to make things clear. When on the spot receiving questions at your presentation proper, listen carefully to what is actually being asked; under stress there can be a tendency to hear the question you are expecting rather than what is really being asked. One way to avoid this trap is to begin a reply by briefly paraphrasing the question back to the asker, not only does this help to establish you are answering the right question, but it also gives a little extra time to marshal your thoughts for the answer.

### 4.5.3 Using Microsoft PowerPoint to prepare slides and posters

Just as with preparing text and organising and analysing data, there are commercial software packages designed for the world of spoken and poster

presentations at meetings. For the same reasons as Excel (flexibility and availability), Microsoft's PowerPoint is the almost universal package of choice here. However, just like Excel, PowerPoint offers a combination of rich opportunities and decided pitfalls. As with Excel, we don't have the space here to give a full-blown tutorial on using PowerPoint, but what follows below and in Box 4.4 is a brief introduction to the package's pluses and minuses.

The most obvious production advantage of PowerPoint is that slides for spoken presentations are produced electronically and not photographically (i.e. as transparencies or diapositives), a process associated with time delays, inflexible results and, often, high cost. Unfortunately, the same advantage doesn't apply to posters. Although PowerPoint allows posters to be created electronically (see Box 4.3), the actual poster still has to be produced through a time-consuming and costly printing process and there is little scope for tinkering without reprinting the entire thing.

One of the most irritating, and potentially disastrous, problems with PowerPoint is that, despite being almost ubiquitous, different versions vary considerably and perform differently on different computer systems, data projectors and printers. As a result, there is no guarantee that what you saw during the creation process will appear during your talk or in your final poster, and, with depressing regularity, your laptop may not work with the particular data projector in the auditorium. Even when all the hardware works together, colours, fonts, layouts, animation effects and audio/visual embellishments are common casualties of different versions of the software. *It is therefore essential to check your presentation thoroughly on the system that will be used at the time well ahead of the giving it for real.* It is also worth checking the compatibility of transfer media ahead of a talk; many times we have seen the miserable consequences of discovering five minutes before being introduced that the host machine won't read the floppy, CD, USB stick or whatever on which the speaker has saved their presentation, or won't connect with the speaker's home network on which it resides. Problems in all respects are particularly likely if you are switching between PC and Macintosh systems, so always check doubly thoroughly here.

With such technical considerations in the background, Box 4.4 introduces some basic facilities offered by PowerPoint for creating slides and posters, but personal experience with using it will lead you to lots of others and your own personal preferences for style and approach.

#### **Box 4.4** A basic guide to PowerPoint

Microsoft's PowerPoint is now almost the default vehicle for preparing and giving spoken and poster presentations. What follows is an extremely basic summary of some of its facilities in this respect. As with most packages, there are usually several ways of achieving the same result (e.g. the keyboard short-cut Control Z undoes an action, as does selecting 'Undo' from the drop-down Edit menu). Most users quickly become familiar with one way of doing things and often remain in blissful ignorance of the others.

## Preparing slides for a talk

### *Slide styles*

When you first boot up PowerPoint the screen offers a default slide format that requests a title and subtitle. You may wish to use this, or you may want a different format. If the latter, click on the 'New' (empty page) icon on the far top left of the screen and a range of alternative slide layouts will appear on the right of the screen (of which the boot-up default layout is the first). Single click on the design you want and it will appear centre screen ready for completion. Where styles have ready-formatted text boxes, such as the first four in the 'Test Layouts' panel, text can be added simply by clicking in the box and typing in. When using the blank open slide in the first 'Content Layouts' format, text must be added by clicking on the 'Text Box' button (the lined panel with the 'A' in the top left corner) on the bottom tool bar, and immediately clicking on the slide to create a panel within which to type text. Of course, the open slide format can be used for many other things besides text, including graphs, photographs, tables, video windows, free-style drawing and so on. Other formats offer different combinations of image, table and text as indicated.

### *Backgrounds*

The form and colour of the background of a slide can be changed by clicking on 'Format' at the top of the screen, followed by 'Background' and selecting the various colour and fill pattern options on offer. Remember to bear in mind the legibility of any text you may want to overlay on the background. In general, the simpler and plainer the background the better. Complicated and fussy images, like photographs of scenery or organisms, usually make very poor, confusing backgrounds.

### *Font size and style*

The font size can be changed by selecting the appropriate number in the 'Font' box in the top tool bar (the white box with a number in it next to the font style [e.g. Arial, Times New Roman] option). Follow the guidelines for font sizes and styles in the main text. The colour of your text can be changed by highlighting it and clicking on the A with a coloured (default black) bar underneath it on the tool bar and choosing a colour from the available options; alternatively click on the symbol and choose a colour prior to typing your text. Be aware of the strengths and pitfalls of different combinations of font and background colours (*see text*).

### *Transitions between slides*

PowerPoint offers many different transition effects when moving from one slide to the next: fade or dissolve in, wipe effects, cover over and many more. To choose one, click on the 'Slide sorter view' button (the one with four small squares in it) at bottom left of the screen. This displays your presentation as rows of slides. Now click on the slide(s) (hold down 'Shift' if you wish to select more than one slide simultaneously) for which you wish the transition effect to apply, then click on 'Transition' in the top tool bar. This produces a

menu of options to the right of the screen, on which you can click to make your selection. PowerPoint will briefly demonstrate the effect as you click it.

#### ***'Build' effects within slides***

As well as customising transitions between slides, you can customise the manner in which text appears in some slide formats. The usual format for this is the numbered or bullet-pointed text slide (the third in the panel of 'Text Layouts' formats above). To choose a style here, click on the 'Normal view' button to the farthest left at the bottom of the screen and move to the slide you wish to customise. Now click on 'Design' in the top tool bar, then on 'Animation schemes' in the options panel that pops up to the right, and finally on your chosen effect. Again PowerPoint will quickly demonstrate the effect.

#### ***Images and movie and sound files***

Images, such as photographs or graphs or tables from statistics packages or scanned from publications, can easily be added to slides in PowerPoint, either from the Windows clipboard, if they've been copied there from another Windows source (just use the 'Paste' option to copy them into the active slide in 'Normal view'), or using the 'Insert' option, if the images are in files saved elsewhere. To use the latter, click on 'Insert' and, if you want, say, a photograph or figure that is saved as a 'jpg' or 'gif' file, click 'Picture' and 'From File' then browse for the file concerned. Double-click on the file name and the image will appear in the slide. You will now need to size and position it as required using the mouse. The 'Picture' option also allows other possibilities, such as inserting 'Clip Art' images and 'Autoshapes' (the standard drawing shapes available from the lower Windows tool bar).

Movie files (such as 'avi' or 'mpeg' files) can be inserted in the same way. From 'Insert' click on 'Movies and Sounds', then 'Movie from File' and browse for the required file. Double-click on it to make it appear in the slide. At this point PowerPoint gives you the option of having the movie play automatically or not until you click on it. Choose which you wish, and size and position the movie window as required. You can now customise the animation in various ways. To make the movie play as soon as the slide appears, right-click on the movie window and choose 'Custom Animation'. In the right-hand column that now appears, click on the downward pointing arrowhead on the blue background immediately next to the title of the movie file and check the 'Start With Previous' option. Click on the arrow again, select 'Timing' and make sure there is a zero second delay (or type in a delay of your choice); then click 'OK'. If you want the movie clip to repeat play until you have finished with the slide, right-click on the movie window again and click on 'Edit Movie Object'. In the resulting dialogue box, check the 'Loop until stopped' option and click 'OK'.

For sound (e.g. 'wav') files, the process is the same as for movies up to the point that a small loudspeaker icon appears on the slide to show that a sound has been inserted. One very useful feature is that by highlighting this icon and using 'Custom Animation', you can make the sound play when you press the Enter key, avoiding the need to position the mouse arrow over the

icon. You can also arrange the slide in an animation sequence triggered by the appearance of a previous item such as a picture of the animal making the sound. An important consideration for both sounds and movies is the size of the file. Even short clips usually exceed 1 MB and take what seems minutes to load. Remember to keep clips as short as possible and check how long they take to load – preferably with the setup that will be used when delivering the talk.

#### *Checking slides and running the presentation*

At any point during the creation of your presentation, you can check what slides will look like when actually running by clicking on the ‘Slide Show’ button, which is the one with the goblet-like icon representing a slide screen on the bottom tool bar. This is also the button to click when you give the presentation itself. Left-clicking the mouse, or pressing the ‘Page Down’ or carriage return keys on the keyboard will move the slide show forwards. Right-clicking the mouse and clicking ‘Previous’, or pressing the ‘Page Up’ key will move the slide show back.

#### **Preparing a poster**

A poster in PowerPoint is usually prepared as a single slide, which is then printed at whatever size is required, normally A0, but A3 and other sizes are occasionally used. Because you are working within a single slide, it is usually necessary to work at something like 150 per cent or 200 per cent normal size, which can be set using the percentage ‘Zoom’ scale in the top tool bar. Adding text is then a matter of clicking on the ‘Text Box’ button on the bottom tool bar and immediately clicking on the slide to create a panel within which to type text (as for open field slides above). Because you are working at magnification, there is a certain amount of judgement to be made about font sizes. First you need to decide how much text there is going to be in total, and how this will be broken up and mixed with figures, tables and images across the poster. Then choose a font size that will accommodate this, but be clearly legible at a comfortable distance once the poster is printed up. You can usually judge this from its appearance in ‘Slide Show’ (*see above*) format. Graphs, tables, photographs, etc. can be inserted as for talk slides via the clipboard or ‘Insert’ options. Text boxes and images can then be moved around with the mouse until the desired layout has been achieved. Don’t forget a good clear banner or box title with all the required information in large, clear font (see suggestions for title slides in the text).


## 4.6 Plagiarism

Before leaving our discussion of reporting work, we have to mention one more issue that has recently been gaining attention in the general education community, and that is *plagiarism*. Plagiarism refers to any attempt to pass off somebody else’s work as your own, and encompasses a wide range of possibilities from

downloading pre-prepared essays from the Internet/World Wide Web, or copying-and-pasting tracts of text or other material from other people's documents into your own without due acknowledgement of the source, to the slightly greyer area of insufficient referencing or use of quotation marks where odd phrases, sentences or ideas have been gleaned from the literature. Science depends on honesty and an ability to attribute ideas and results properly in reporting studies, so quite apart from any dishonest advantage that may be sought in educational assessments, wanton plagiarism undermines the scientific process itself. It is therefore a serious matter. Of course, the possibility of plagiarism has always been with us, but, in recent years, the Internet/World Wide Web has greatly exacerbated the problem by making available a vast amount of material that is downloadable at the push of a button and (at least until very recently) difficult for an external party to trace. Many educational establishments now use sophisticated anti-plagiarism software, which can scan a submitted piece of work and quickly match its contents to other electronic sources, so the arms race against cheats has moved on significantly. However, as with any policing process, it is possible to fall foul of the system inadvertently through carelessness or ignorance. Our advice, therefore, is to take careful note of the reporting conventions in professional journals, books and conferences, and, especially if you're a student, the guidelines given to you by your university, college or school.

## 4.7 Summary

1. Confirmatory analyses are usually presented in summarised form (e.g. summary statistics, scattergrams) as tables or figures or in the text of a report. In all cases, sample sizes (or degrees of freedom), test statistics and probability levels should be quoted. In the case of tables and figures, these can be included within the table or figure itself or within a full, explanatory legend.
2. Results should almost never be presented as raw numerical data because these are difficult for the reader to assimilate. In the exceptional circumstances where the presentation of raw data is helpful, presentation should usually be selective to the points being made and is best incorporated as an Appendix.
3. The axes of figures should be labelled in a way that conveys their meaning clearly and succinctly. Where analyses in different figures are to be compared directly, the axes of the figures should use the same scaling.
4. The legends to tables and figures should provide a complete, self-contained explanation of what they show without the reader's having to search elsewhere for relevant information.
5. Reports of investigations should be structured into clearly defined sections: Introduction, Methods, Results, Discussion, References. Each section has a specific purpose and deals with particular kinds of information. The distinction between them should be strictly maintained. When reports are long




or involved it can be helpful to add an Abstract and/or Conclusions section to highlight the main points and take-home messages.

6. At a time when increasing importance is being attached to the public understanding of science, writing for a general interest readership, as opposed to the more usual professional scientific one, is a useful skill. Success often depends on avoiding technical jargon, or at least explaining it in everyday language, and using familiar analogies or imagery to get attention and facilitate understanding.
7. Information can also be presented in the form of talks or poster papers at meetings. This can be a very important means of disseminating your work and getting yourself known, and, as usual, there are various 'dos' and 'don'ts' that make for good practice in preparing presentations.
8. Be sure to reference sources of information properly in any presentation of your work. Plagiarism (an attempt to pass off somebody else's work as your own) is a serious offence in education and scientific research.

## Reference

Sokal, R. R. and Rohlf, F. J. (1995) *Biometry*, 3rd edition. Freeman, San Francisco.





# Test finder and help guide

## Test finder

- A Looking for a difference or a trend? (if unsure, go to *Help 1*)
- \* Trend – go to B
  - \* Difference – go to C
- B Have you fixed the X-values experimentally (looking to see whether changes in X cause changes in Y) or do you require the equation of the line itself (e.g. for predicting new values)? (if unsure, go to *Help 2*)
- \* No – correlation
  - \* Yes – linear regression
- C Does each data value belong to treatments within one level of grouping, or does it belong to more than one level of grouping? (if unsure, go to *Help 3* and *Help 6*)
- \* One level of grouping (one-way designs) – go to D
  - \* More than one level of grouping (n-way designs) – go to G
- D Are data for each treatment replicated? (if unsure, go to *Help 4*)
- \* No – go to E
  - \* Yes – go to F
- E Data are in the form of counts
- \* No – (not analyzable)
  - \* Yes –  $1 \times n$  chi-squared
- F Data in one treatment are independent of those in other treatments (if unsure, go to *Help 5*)
- \* No
    - > only two treatments – matched-pairs tests, or use one-way repeated-measures ANOVA
    - > two or more treatments – one-way repeated-measures ANOVA
  - \* Yes
    - > only two treatments – *t*-test or Mann–Whitney test or use one-way ANOVA
    - > two or more treatments – one-way ANOVA
- G Are data for treatments within each level of grouping replicated? (if unsure, go to *Help 6*)
- \* No –  $n \times n$  chi-squared
  - \* Yes – *n*-way ANOVA

## Help 1

**Difference or trend?**

- *Difference* predictions are concerned with some kind of difference between two or more groups of data. The groups could be based on any characteristic that can be used to make a clear-cut distinction, e.g. sex, drug treatment, habitat. Thus a difference might be predicted between the growth rates of men and women, or between the development of disease in rats given drug A versus those given drug B versus those given a placebo.
- *Trend* predictions are concerned not with differences between mutually exclusive groupings but with the relationship between two more or less continuously distributed measures, e.g. the relationship between the size of a shark and the size of prey it takes, or the relationship between the amount of rainfall in a growing season and the number of apples produced by an apple tree.

## Help 2

**What sort of trend?**

The basic choice here is between fitting a line (regression) or not (correlation).

Correlation is used whenever we merely want to know whether there is an association between X and Y – there is no real dependent or independent variable, and when plotted as a scattergraph you could equally well plot Y against X as opposed to X against Y.

Regression was developed for situations where the experimenter manipulates the values of the *independent* variable (X) and measures the impact of these manipulations on another *dependent* variable (Y). You are therefore looking for a causal relationship between X and Y – changes in X cause changes in Y. However, because knowing the slope of the relationship between X and Y is useful in many other contexts, the use of regression has expanded to incorporate many cases where the X values are merely measured rather than manipulated. Technically this is wrong, but this usage is so firmly embedded in biological practice that the majority of scientific investigations do it. An example is when you want to use the relationship to predict the value of Y for a particular X value. Comparisons of slopes are also extremely informative, and used extensively in biological research.

## Help 3

**Levels of grouping**

Many difference predictions are concerned with differences at just *one* level of grouping, e.g. differences in faecal egg counts following treatment of mice with one of four different anthelmintic drugs. Here drug treatment is the only level of grouping in which we are interested. However, if we wished, say, to distinguish between the effects of different drugs on male and female mice, we should be dealing with *two* levels of grouping: drug treatment and sex.

## Help 4

**Replication**

Replication simply means that each treatment within a level of grouping has more than one data value in it. The table below shows replicated data as columns of values within each treatment in a one-way (i.e. one level of grouping) design:

	Group		
	Pesticide A	Pesticide B	Control
% mortality of pest	10	30	0
	5	27	1
	3	50	1
	0	6	0
	1	3	2
	20	3	5

## Help 5

**Independence**

Unless specifically allowed for in the analysis, all statistical analyses assume that each data value is independent of all others. Each value in one group is *non-independent* of one from each of the other groups if they have something in common (e.g. it is measured on the same individual, or derives from animals from the same cage – i.e. a given individual is exposed to each treatment in turn, or the same cage provides animals for each treatment). The source of non-independence thus needs to be taken into account in any analysis, and so requires a so-called repeated-measures design.

## Help 6

**Rows and columns**


If data have been collected at two levels of grouping, then each data value can be thought of as belonging to both a row and a column (i.e. to one row/column cell) in a table, where rows refer to one level of grouping (say sex – *see* Help 4) and columns to the other (drug treatment – *see* Help 4). If there are several values per row/column cell, as below for the number of individuals dying during a period of observation:

		Treatment	
		Experimental	Control
Sex	Male	3, 4, 8, 12	23, 24, 12, 32
	Female	1, 0, 2, 9	32, 45, 31, 21

then a two-way analysis of variance is appropriate. If there is just a single *count* in each cell, as in the number of male and female fish responding to an experimental or control odour stimulus:

		Treatment	
		Experimental	Control
Sex	Male	27	91
	Female	12	129

then an  $n \times n$  chi-squared test is appropriate.



## Some self-test questions

(Answers on p. 229)

1. An experimenter recorded the following body lengths of freshwater shrimps (*Gammarus pulex*) in three different lakes.

Body size (mm)		
Lake 1	Lake 2	Lake 3
9.9	10.5	9.6
8.7	12.1	9.0
9.6	11.2	8.7
10.7	9.7	13.2
8.9	8.7	11.9
8.2	11.1	14.0
7.7	10.7	12.9
8.1	11.8	10.8

Faunal diversity in the lakes was known ( $1 < 2 < 3$ ) and the experimenter expected shrimps from more diverse lakes to be smaller because of increased interspecific competition. To test this idea he compared body lengths in each pair of lakes (1 versus 2, 2 versus 3 and 1 versus 3) using Mann–Whitney  $U$ -tests. Is this an appropriate analysis? If not, what would you do instead?

2. How would you decide between correlation and regression analysis when testing a trend prediction?
3. The following is part of the Discussion section of a report into the effects of temperature and weather on the reproductive rate of aphids on bean plants.

While the results show a significant increase in the number of aphids produced as temperature rises, there is a possible confounding effect of the age of the host plant and the rate of flow of nutrients. Indeed, there was a stronger significant positive correlation between nutrient flow rate and the number of aphids produced ( $r_s = 0.84$ ,  $n = 20$ ,  $p < 0.01$ ) than between temperature and production (*see Results*).

Do you have any criticisms of the piece?

4. What does the following tell you about the analysis from which it derives?

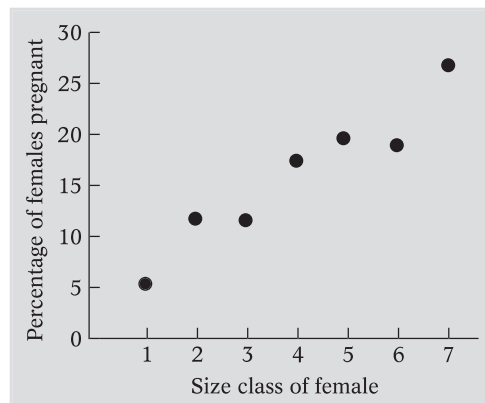
$$H = 14.1, \quad \text{d.f.} = 3, \quad p < 0.01$$

5. An agricultural researcher discovered a significant positive correlation ( $r_s = 0.79$ ,  $n = 112$ ,  $p < 0.01$ ) between daily food intake and the rate of increase in body weight of pigs. What can the researcher conclude from the correlation?
6. A plant physiologist measured the length of the third internode of some experimental plants that had received one of three different hormone treatments. The physiologist calculated the average third internode length for each treatment and for untreated control plants. The data were as follows:

	Treatment			
	Control	Hormone 1	Hormone 2	Hormone 3
Average internode length (mm)	32.3	41.6	38.4	50.2

To see whether there was any significant effect of hormone treatment, the physiologist performed a  $1 \times 4$  chi-squared test with an expected value of 40.6 in each case and three degrees of freedom. Was this an appropriate test?

7. What do you understand by the terms:
- test statistic,
  - ceiling effect,
  - statistical significance?
8. The figure shows a significant positive correlation, obtained in the field, between body size in female thargs and the percentage of females in each size class that were pregnant. From this, the observer concluded that male thargs preferred to mate with larger females. Is such a conclusion justified? Give reasons for your answer.



9. Why are significance tests necessary?
10. The following were the results of an experiment to look at the effect of adding an enzyme to its substrate and measuring the rate at which the substrate was split. In the control Treatment A, no enzyme was added; in Treatment B, 10 mg of enzyme was added; in Treatment C, 10 mg was added but the reaction was cooled; and in Treatment D, 10 mg was added but the reaction was warmed slightly.

Treatment	Treatment	Treatment	Treatment
A	B	C	D
0	20	52	71
1	21	69	92
2	35	100	55
1	15	32	78
0	20		105
0	24		82
			92


What predictions would you make about the outcome of the experiment and how would you analyse the data to test them?

11. A farmer called in an agricultural consultant to help him decide on the best housing conditions (those resulting in the fastest growth) for his pigs. Three types of housing were available (sty + open paddock, crating, and indoor pen). The farmer also kept four different breeds of pig and wanted to know how housing affected the growth rate of each. What analysis might the consultant perform to help the farmer reach a decision?
12. Figures (a) and (b) were used by a commercial forestry company to argue that the effect of felling on the number of bird species living in managed stands (assessed by a single standardised count in each case) was similar in both deciduous and coniferous forest. Would you agree with the company's assessment on this basis?



13. What would a *negative* value of the test statistic  $r$  signify to you?
14. Derive some hypotheses and predictions from the following observational notes:
- Sampled some freshwater invertebrates from three different streams using a hand-net. There were more individuals of each species at some sites than others, both within streams and between them. Also some sites had a more or less even distribution of individuals across all species whereas others had a highly biased distribution with some species dominating the community. Some species occurred in all three streams but they tended to be smaller in some streams than others. A number of predatory dragonfly nymphs were recorded but there was never more than one species in any one sample, even when more than one existed in a stream. Water quality analyses showed that one stream was badly polluted with effluent from a local factory. This stream and one of the others flowed into the third stream, forming a confluence. It was noticed that stones and rocks on the substrate had fewer organisms on or under them in regions of faster flow rate.
15. Is chi-squared used for testing differences or trends?
16. A student in a hall of residence suffered from bed bugs. During the course of a week he was bitten 12 times on his legs, 3 times on his torso, 6 times on his arms and once on his head. Could these data be analysed for site preferences by the bugs? If so, how?
17. An experimenter had counted the number of times kittens showed elements of play behaviour when they were in the presence of their mother or their father and with or without a same-sex sibling. The experimenter had collected ten counts for each condition: (a) mother/sibling, (b) mother/no sibling, (c) father/sibling, (d) father/no sibling, and was trying to decide between a  $2 \times 2$  chi-squared analysis and a  $2 \times 2$  two-way analysis of variance. What would you suggest and why?
18. Why do biologists regard a probability of 5 per cent or less as the criterion for significance? Why not be even stricter and use 1 per cent?
19. A fisheries biologist was interested in the maximum size of prey that was acceptable to adult barracuda. To find out what it was, he introduced six adult barracuda into separate tanks and fed them successively larger species of fish (all known to coexist with barracuda in the wild). He then calculated the mean size of the fish that the barracuda last accepted before refusing a fish as a measure of the maximum size they would take. Is this a sensible procedure?
20. A psychologist argued that since males of a species of monkey had larger brains than the females there was less point in trying to teach females complex problem-solving exercises. Any comments?

21. An ecologist studying populations of voles in different woods suspected from a glance at the data that males from some woods had larger adrenal glands than those from others. Unfortunately, the age of the animals also appeared to differ between the woods. How might the ecologist test for a difference in adrenal glands between woods while controlling for the potential confounding effect of age?
22. What do you understand by an 'order effect'?
23. A parasitologist wished to test for the effect of increasing the amount of food supplement to nestling birds on subsequent parasite burdens as adults. The parasitologist intended to carry out a linear regression analysis with amount of food supplement on the  $x$ -axis and adult parasite burden on the  $y$ -axis. However, when the distribution of parasite burden data was checked for normality, the following was found:




What implications does this have for the regression analysis?

24. What is meant by pseudoreplication?
25. When a botanist compared the frequency of different leaf sizes on a tree to a normal distribution, the significance test comparing the two distributions yielded a test statistic value with an associated probability of 0.0341. Does the botanist's data conform to a normal distribution or not?
26. A behavioural ecologist was interested in the effects of other competing males in the environment on sperm production by focal courting males. Her prediction was that courting males should produce more sperm to transfer to a female if there are more males around to compete for her. She decided



to test her idea by exposing 20 focal males to each of three treatments: (a) no other males present, (b) one other male present, and (c) four other males present. She was careful to randomise the order in which the three treatments were presented to her different subjects and she measured the width of the thorax of each subject male and stimulus female to gauge their body size. The number of sperm produced was estimated as the weight of the spermatophore (package of sperm) transferred to the female. How might she go about analysing the data from the experiment to test her prediction?

27. A psychologist interested in sexual attitudes in men and women from different social environments carried out a questionnaire survey of some 150 subjects. The questionnaire consisted of 40 questions covering a range of aspects of sexual behaviour from behavioural characteristics preferred in potential partners to attitudes towards promiscuity and homosexuality. It also asked for various items of background information, such as financial status, family background, age, religion and ethnic group. How might the psychologist set about preparing this kind of data for analysis, and what analysis might they choose?
28. In a study of the effects of various environmental variables on resistance to an experimental infection in male house mice (*Mus domesticus*), a student had measured four variables with respect to subject males: food availability during development, maternal body weight, local population density and litter size. These were then related to a measure of resistance to an experimental infection with a blood protozoan (clearance rate) by means of four separate regression analyses with clearance rate as the dependent variable in each case. Is this a reasonable approach to the analysis?
29. In a study of the behaviour of foraging bees, individual bees were presented with flowers treated in one of three different ways: (a) not previously visited by any bees, (b) previously visited by the subject forager herself, and (c) previously visited by a different individual forager. Several aspects of each subject's behaviour were measured, including the speed of approach to the flower, the length of time spent hovering in front of the flower before visiting it, whether the flower was touched by the bee, how many times the bee touched the flower with its antennae, the duration of antenna touches, whether the flower was visited at all, the length of any visit, the length of time spent hovering by the flower after a visit, and the speed of departure from the flower. How might these data be analysed to reveal how the bees responded to the different flower treatments?
30. A physiologist was interested in the effect of mobile phones on nerve cell (neurone) function. He therefore measured impulse transmission rates in replicated laboratory preparations of neurones under three different strengths of electromagnetic field, plus a control of zero magnetic field. Each preparation was only used once, and each treatment was replicated 30 times. What kind of analysis is appropriate?




# Appendix I

## Table of confidence limits to the median

Table of non-parametric confidence limits to the median

Sample size ( $n$ )	$r$ (for $p$ approx. 95%)
2	–
3	–
4	–
5	–
6	1
7	1
8	1
9	2
10	2
11	2
12	3
13	3
14	3
15	4
16	4
17	5
18	5
19	5
20	6
21	6
22	6
23	7
24	7
25	8
26	8
27	8
28	9
29	9
30	10

$r$  denotes the number of values in from the extremes of the data set that identifies the 95 per cent confidence limits (*see text*). Modified after Colquhoun (1971) *Lectures on biostatistics*, Clarendon Press, Oxford.



## Appendix II

### How to calculate some simple significance tests

#### Examples of tests for a difference between two or more groups

##### Box A1 The basic calculations for parametric test statistics

If you calculate the test statistic of many parametric statistical tests by hand, the calculations often involve a standard set of core operations on the data. Rather than repeating them in all the relevant boxes, we have placed them here for reference. They all involve calculating a quantity known as the *sum of squares*, which is actually the same operation performed when finding the standard deviation or standard error (see Box 2.1).

1. Follow steps 1–3 of Box 2.1. The resulting quantity is the *sum of squares of  $x$* , conventionally denoted by  $S_{xx}$ .

Sometimes you will have a  $y$ -variable as well as an  $x$ -variable (e.g. in regression, see p. 124). If so:

2. Calculate the sum of all the  $y$ -values in the data set ( $\sum y$ ).
3. Square all the  $y$ -values and sum them, giving ( $\sum y^2$ ).
4. Multiply the  $x$ - and  $y$ -values of each  $x$ - $y$  pair together, and add them together, giving ( $\sum xy$ ).
5. Calculate  $\sum y^2 - [(\sum y)^2/n]$ , giving the *sum of squares of  $y$* ,  $S_{yy}$ .
6. Calculate  $\sum xy - [(\sum x)(\sum y)/n]$  giving the *sum of the cross-products*,  $S_{xy}$ .

##### Box A2a (i) Mean values: how to do a general parametric test for two groups (two-tailed $t$ -test)

1. Frame the prediction. In this case it is the general prediction that the two groups (A and B) will have different mean values ( $A \neq B$ ). Thus the null hypothesis is that the two group means will not differ.
2. Count the number of data values in the first group; this number is referred to as  $n_1$ . If  $n_1 = 1$ , then there is definitely something wrong! (If this number represents a count of the number of items in one of two categories – that form the two

groups – then you should be doing a  $\chi^2$  test. If it is a rank or a constant-interval measurement, then you need to collect some more data!) Count the number of data values in the second group; this number is referred to as  $n_2$  (it should be greater than 1, as before).

3. Calculate the variances of each group separately (see Box 2.1 for how to do this), producing  $s_1^2$  and  $s_2^2$ .
4. Calculate the value of  $P = (n_1 + n_2)/(n_1 n_2)$ .
5. Calculate the value of  $Q = (s_1^2(n_1 - 1) + s_2^2(n_2 - 1))/(n_1 + n_2 - 2)$ .
6. Calculate the value of  $R = \sqrt{PQ}$ . This is the standard error of the difference between the two groups.
7. Calculate the mean values of each group separately, and take the difference  $S = (\mu_1 - \mu_2)$ . Switch the mean values round so as to make the difference positive, since in a general test we are not interested in the direction of the difference, but merely in whether it differs from zero.
8. Calculate the value of the test statistic,  $t = S/R$ .
9. In order to calculate  $t$ , we needed to know the difference in mean values, and the s.e. of this difference, i.e. two prior parameters were required. The degrees of freedom of  $t$  are therefore  $n_1 + n_2 - 2$ .
10. Look up the two-tailed value of  $t$  in Table D of Appendix III for the critical value for your degrees of freedom. If your value is greater than the relevant value in the table, then the difference you found is significant.

### Box A2a

### (ii) Mean values: how to do a specific parametric test for two groups (one-tailed $t$ -test)

1. Frame the prediction, i.e. decide which of the two groups (A and B) is predicted to have the greater mean value. There needs to be some a priori reason (theory, or previous published or gathered data) for this prediction. Suppose that on the basis of your knowledge, you predict that  $A > B$ . The null hypothesis is that the mean values for the two groups do not differ in the predicted direction.
2. Calculate the value of  $t$  as steps 2–9 above, but make sure that the difference in mean values is done in the way that is *predicted* to generate a positive difference. Here you are predicting that  $A > B$ , and hence  $(A - B)$  should generate a positive value of  $t$ . If it does not actually generate a positive difference, then you know automatically that the result is not significant. Note that if the result is an unusually large but negative value of  $t$  (that would have been significant, had you predicted the opposite pattern of mean values), you are not allowed to conclude *anything* other than that your prediction is not supported by the data. This is the cost of a specific (one-tailed) prediction paid in exchange for the benefit of the more powerful test.
3. Look up the critical value of a one-tailed  $t$ -test in Table D of Appendix III using the degrees of freedom you have. If your  $t$ -value is greater than the critical one, then you conclude that the result is significant: the evidence suggests that the group mean predicted to be greater really is so, and you reject the null hypothesis.

**Box A2b****(i) Mean ranks: how to do a general non-parametric test for two groups (Mann–Whitney  $U$ -test)**

1. Count the number of data values in the group with the fewer values (if there is one); this number is referred to as  $n_1$ .
2. Count the number of data values in the other group; this is  $n_2$ .
3. Rank all the values *in both groups combined*. The smallest value takes the lowest rank of 1, the next smallest value a rank of 2 and so on. If two or more values are the same they are called *tied values* and each takes the average of the ranks they would otherwise have occupied. Thus, suppose we have allocated ranks 1, 2 and 3 and then come to three identical data values. If these had all been different they would have become ranks 4, 5 and 6. Because they are tied, however, they each take the same average rank of 5  $((4 + 5 + 6)/3)$ , though – and this is important – the next highest value still becomes rank 7 just as if the three tied values had been ranked separately. If there had been only two tied values they would each have taken the rank 4.5  $((4 + 5)/2)$  and the next rank up would have been 6. We should thus end up with rank values ranging from 1 to  $N$ , where  $N = n_1 + n_2$ .
4. Add up the rank values within each group giving the total  $R_1$  and  $R_2$  respectively.
5. Calculate  $U_1 = n_1 \times n_2 + ((n_1(n_1 + 1))/2) - R_1$ .
6. Calculate  $U_2 = n_1 \times n_2 - U_1$ .  
If  $U_2$  is smaller than  $U_1$  then it is taken as the test statistic  $U$ . If not, then  $U_1$  is taken as  $U$ .
7. We can now check our value of  $U$  against the threshold values in  $U$  tables (a sample is given in Appendix III, Table B). If our value is *less* than the threshold value for a probability of 0.05, we can reject the null hypothesis that there is no difference between the groups. Note that, in this test, we use the two sample sizes,  $n_1$  and  $n_2$ , rather than degrees of freedom to determine our threshold value.

If one of the groups has more than 20 data values in it,  $U$  cannot be checked against the tables directly. Instead, we must use it to calculate another test statistic,  $z$ , and then look this up (some sample threshold values are given in Appendix III, Table C). The calculation is simple:

$$z = \frac{U - (n_1 \times n_2)/2}{[(n_1)(n_2)(n_1 + n_2 + 1)]/12}$$

**Box A2b****(ii) Mean ranks: how to do a specific non-parametric test for two groups**

There is no non-parametric test specially designed to test for a specific difference between two groups. Use the specific non-parametric analysis of variance (*see* Box 3.4a), because this can cope with any number of groups.

## Box A3a

## (i) Mean values: how to do a general parametric one-way analysis of variance (ANOVA)

1. Frame the prediction. In this case, the general prediction being made is to ask whether there are any differences among the mean values of the groups. Therefore the null hypothesis actually tested is that there are no differences among the mean values of the groups.
2. The test considers  $i$  groups of data, each of which contains  $n_i$  data values. The total number of data values in all groups together =  $N = \sum n_i$ .
3. Calculate the mean values of each of the groups,  $\mu_i$ .  
Work out  $T$ , the total sum of squares of all the data (follow Box 3.2, items 1–3).
4. Work out  $S_i$ , the sum of squares for the data of each group separately (again, follow Box 3.3, 1–3).
5. Calculate the error sum of squares,  $SS_{\text{error}} = \sum S_i$ .
6. Calculate the among-groups sum of squares,  $SS_{\text{among}} = T - SS_{\text{error}}$ . This should give the same result as calculating the sum of squares using the mean values  $\mu_i$  rather than the raw data.
7. The d.f.<sub>total</sub>, the total degrees of freedom, is  $N - 1$ .
8. The among-groups degrees of freedom, d.f.<sub>among</sub>, is  $(i - 1)$ .
9. The error degrees of freedom, d.f.<sub>error</sub>, is  $(N - i)$ .
10. The mean square among groups,  $MS_{\text{among}} = SS_{\text{among}}/\text{d.f.}_{\text{among}}$ .
11. The error mean square,  $MS_{\text{error}} = SS_{\text{error}}/\text{d.f.}_{\text{error}}$ .
12. The test statistic,  $F = MS_{\text{among}}/MS_{\text{error}}$ .
13. The degrees of freedom for  $F$  are d.f.<sub>among</sub>, d.f.<sub>error</sub>  
(note that  $F$  has *two* values for degrees of freedom, unlike other test statistics, which have only one).
14. Now look up the critical value of  $F$  in Table G of Appendix III for  $f_1 = \text{d.f.}_{\text{among}}$ , and  $f_2 = \text{d.f.}_{\text{error}}$ .
15. If your value for  $F$  is greater than the critical value, then the result is taken to be significant, and we reject the null hypothesis of equal mean values for all groups. Present the one-way ANOVA laid out in the standard manner (*see* worked example in Appendix II).
16. Note that this is a *general* prediction and, therefore, we can only conclude that there are *some* differences among the means: we are *not* allowed to say which particular pair of mean values are significantly different because we did not make any specific prediction beforehand about the ordering of the mean values, nor did we set up planned specific contrasts between particular sets of groups. Many researchers would like to be able to make such *post hoc* tests, however, and a plethora of different methods are available and are often used. However, they should only be used as a rough guide to what the differences might be, preferably to set up a priori hypotheses for a new data set to test. Many people use the least significant difference, or LSD test. All this does is to take the  $MS_{\text{error}}$  from the ANOVA (which is an estimate of the variation within groups) and construct a 95 per cent confidence limit with it using:

$$\text{LSD} = t_{N-i} \times \sqrt{[(2/n) \times MS_{\text{error}}]}$$

where  $t_{N-i}$  is the 5 per cent threshold value of  $t$  for  $N - i$  degrees of freedom. This looks a simple enough procedure, but in fact the LSD test is only accurate for a priori contrasts. It should, therefore, only be used as a rough guide in *post hoc* testing.

### Box A3a (ii) Mean values: how to do a specific parametric one-way ANOVA

Unfortunately there is effectively no parametric equivalent of the test for a particular rank order of mean ranks, as in Box A3b(ii). Although one does exist, called isotonic regression (*see* Gaines and Rice, 1990), it is very obscure and hardly ever used; it also involves computer-intensive randomisations of the data, and is well beyond the scope of this book. More common (but still unusual) is the use of a priori contrasts. If done in advance of obtaining the data, these allow  $(i - 1)$  contrasts to be made. We explain this technique here.

1. Each contrast consists of one subset of the  $i$  groups (A) contrasted against another subset (B) (for example, a control (A) versus all treated groups (B)): each therefore involves effectively creating *two* groups out of the data, and testing them using a  $t$ -test. Such contrasts can themselves be either general ( $A \neq B$ ) or specific (e.g.  $A > B$ ).
2. Frame the predictions. Here you do this by formulating the contrasts that you want to make *in advance of collecting the data*. You are allowed  $(i - 1)$  contrasts. Express each contrast as an inequality, and make the left-hand side greater or equal to zero. For example, if you have three groups (A, B and C), and want to test whether group A is different from the other two groups (i.e. from the average of  $B + C$ ), then this general hypothesis is that:

$$A \neq (B + C)/2$$

and hence

$$2A - B - C \neq 0, \text{ i.e. } +2(A) - 1(B) - 1(C) \neq 0$$

This gives us a set of coefficients to apply to the mean values of each group; in this case they are +2, -1 and -1 for the mean values of groups A, B and C, respectively.

Groups that you want to leave out of the contrast have a coefficient of zero. Thus another example might be that you are predicting that the control group A has a greater mean value than group C; this specific hypothesis is that:

$$A > C$$

and hence (making the left-hand side greater than zero)

$$+1(A) + 0(B) - 1(C) > 0$$

The coefficients are labelled as  $\lambda_i$  for the  $i$  groups.

3. You can check that you have a valid set of coefficients, because if so, then they sum to zero, i.e.  $\sum \lambda_i = 0$ .

In addition, each contrast must be independent of all the others. This is the main reason why a maximum of  $(i - 1)$  contrasts are allowed, since it is not possible for more than that to be independent. Statistically, if two contrasts are

independent they are called orthogonal. If you have equal sample sizes per group (as every well-designed experiment should have!), you can easily check whether any two contrasts ( $a$  and  $b$ ) are independent since the following must be true:

$$\sum \lambda_{ia} \times \lambda_{ib} = 0$$

In the case of unequal sample sizes of  $n_i$  per group, the equivalent formula is  $\sum n_i \times \lambda_{ia} \times \lambda_{ib}$ , and there are additional complications later on in the calculation. Strictly it is not essential for every contrast to be independent, but if they are not (as in this example), you need to adjust the threshold probability of significance: instead of 0.05, it becomes  $1 - (0.95)^{1/r}$ , where  $r$  = the number of non-independent contrasts you make.

4. Do a general parametric one-way ANOVA first (*see* Box 3.5a(i)), since technically you are breaking the among-group differences down into specific independent contrasts (*decomposing the sums of squares* in statistical jargon). You will need the error mean square,  $MS_{\text{error}}$ , to test each contrast.
5. For each contrast, obtain the coefficients ( $\lambda_i$ ) and using the mean values ( $m_i$ ) for each of the original groups, do the following calculation, which results in a sum of squares for the contrast. This contrast has one degree of freedom, and hence is also a mean square (MS):

$$L = \sum \lambda_i \times m_i$$

If you are testing a *specific* hypothesis about the contrast, check at this point that  $L$  is positive. If it is, this means that the data follow the predicted pattern (analogous to the positive value of a one-tailed  $t$ -test). If it is not positive, then you know already that your hypothesis will not be supported by the data, and will not be significant.

$$MS_L = L^2 n / \sum \lambda_i^2$$

where  $n$  is the sample size of each group.

If you have unequal sample sizes for your groups, follow the method of Sokal and Rohlf (1995, pp. 528–529), but you should probably use a computer package to do this calculation for you!

6. Now form the variance ratio
 
$$F = MS_L / MS_{\text{error}}$$
 where  $MS_L$  is the mean square for the contrast, and  $MS_{\text{error}}$  is the error mean square from the one-way ANOVA. The test statistic is then
 
$$t = \sqrt{F}$$
7. The degrees of freedom of the  $t$ -test for the contrast is  $a(n - 1)$ , where  $a$  is the number of groups involved in the contrast (excluding the groups left out, those with coefficients of zero).
  - (a) If your hypothesis is *general*, then look up the critical value of a two-tailed  $t$ -test in Table D of Appendix III.
  - (b) If your hypothesis is *specific*, then look up the critical value of a one-tailed  $t$ -test in Table D of Appendix III.
8. Repeat steps 2–7 for each contrast.
9. Present the results of the one-way ANOVA, as well as the contrasts, laid out in the standard manner.



### Box A3b (i) Mean ranks: how to do a general non-parametric one-way ANOVA

1. Formulate the prediction. In this case, the general prediction being made is to ask whether there are any differences at all among the mean ranks of the groups. Therefore the null hypothesis actually tested is that there are no differences among the group mean ranks.
2. The test considers  $i$  groups of data, each of which contains  $n_i$  data values. The total number of data values in all the groups together =  $N = \sum n_i$ .
3. Rank all the values across all the groups combined (as in the  $U$ -test), giving low rank scores to low values. Once again, tied values are given the average of the ranks they would have been ascribed had they been slightly different. Where there are lots of tied values relative to the sample size, you may need to apply a tie-correction factor, but in this case it is better to get the calculation done by a computer.
4. Sum the ranks in each group, giving  $R_i$  in each case.
5. The test statistic is  $H$ , where

$$H = \frac{12}{N(N+1)} \times (\sum R_i^2/n_i) - 3(N+1)$$

6. The degrees of freedom are  $(i - 1)$ .
7. Look up the significance of the calculated  $H$  value as if it were a  $\chi^2$ , in Table A of Appendix III (although  $H$  is not actually a  $\chi^2$ , its value is distributed in the same way, so it is as if we were using  $\chi^2$ ). There is no standard layout for a general non-parametric ANOVA; just quote the test statistic, its degrees of freedom and the probability (*see* worked example in Appendix II).
8. Note that this is a *general* prediction and, therefore, if significant we can only conclude that there are *some* differences among the mean ranks: we are *not* allowed to say which particular pair of mean values are significantly different because we did not make any specific prediction beforehand about the ordering of the mean values, nor did we set up planned specific contrasts between particular sets of groups. Many researchers would like to be able to make such *post hoc* tests, however, and some methods are available (but not often used). They should only be used as a rough guide to what the differences might be, preferably to set up a priori hypotheses for a new data set to test. Sokal and Rohlf (1995, p. 431) and Day and Quinn (1989) have some recommendations.

### Box A3b (ii) Mean ranks: how to do a specific non-parametric one-way ANOVA

There are two ways of making specific predictions about the mean ranks of your groups. The one we favour here uses all the groups in a single a priori prediction of their rank order. The alternative is to use a priori contrasts (*see* point 9, below).

1. Formulate the specific prediction by specifying a particular rank order of the mean ranks of the groups, based on some a priori knowledge (theory, or previous published or gathered data), in advance of obtaining the data. The null hypothesis is that the rank order does not follow the prediction.

2. The test considers  $i$  groups of data, each of which contains  $n_i$  data values. There are  $N$  data values in total in all the groups ( $= \sum n_i$ ).
3. Rank all the values across all the groups combined (as in the  $U$ -test), giving low rank scores to low values. Once again, tied values are given the average of the ranks they would have been ascribed had they been slightly different. Where there are lots of tied values relative to the sample size, you may need to apply a tie-correction factor, but in this case it is better to get the calculation done by a computer.
4. Sum the ranks in each group, giving  $R_i$  in each case.
5. Assign the predicted rank order to the groups, from the lowest (rank = 1) to the highest (rank =  $i$ ). This rank order then provides the  $\lambda_i$  coefficient values. Using these  $\lambda_i$  values, calculate:
 

the observed  $L = \sum \lambda_i R_i$ ,

the expected  $E = (N + 1)(\sum n_i \lambda_i)/2$ ,

the variance  $V = (N + 1)(N \sum n_i \lambda_i^2 - (\sum n_i \lambda_i)^2)/12$ .
6. Calculate the test statistic,  $z$ , as:  $z = (L - E)/\sqrt{V}$ . (Note that  $z$  is a standardised statistic which does not have any degrees of freedom.)
7. Look up the value of  $z$  in Table C of Appendix III, where you will see that the critical value for this specific (one-tailed) test is 1.64. If your value is greater than this, then the result is significant. If it is significant, we then reject the null hypothesis: there is evidence that the mean ranks fall into the predicted rank order.
8. What do you do if the result is *not* significant? If the mean ranks in fact fall in the opposite direction to your prediction, the value of  $z$  will be negative and quite possibly greater in absolute magnitude than 1.64. You *cannot* conclude *anything* about this, since your predicted rank order was not supported. You benefited from a gain in power over a general test, but the cost was that you could not conclude anything from a failure to reject the null hypothesis. You certainly cannot go and test an alternative rank order: this would now be *post hoc* since you have seen the actual pattern of the mean ranks. What you can do, following on from a non-significant specific test, is to ask the question: my predicted rank order was not supported, but is there evidence of *any* differences among groups in the data? In other words, you can go ahead and do a general test for any differences.
9. An alternative method is to use a priori contrasts, similar to the parametric case of Box A3a(ii). If done in advance of obtaining the data, these allow  $(i - 1)$  independent contrasts to be made, each one consisting of a subset of the groups contrasted against another subset. Follow the method of Box A3a(ii), points 1–2, to obtain the coefficients for the contrast you want to make, and then create two new groups by adding together the data for all the original groups that have the same sign (+ or –) coefficient. These two artificial groups are then tested using either a Mann–Whitney  $U$ -test (general prediction only), or preferably a non-parametric one-way ANOVA (general or specific prediction).

Note that the coefficients for these specific contrasts are *not* the same thing as the coefficients that specify the rank order for the test outlined above, points 1–8. You cannot use positive, negative and zero coefficients in the rank order test, imagining that you are doing a specific contrast.

10. There is no standard layout for a specific non-parametric one-way ANOVA; just quote the test statistic for each prediction, its degrees of freedom (if appropriate) and the probability. Note that a  $z$ -test does not have degrees of freedom, whereas specific contrasts have one degree of freedom (using either Mann–Whitney  $U$ -tests or non-parametric one-way ANOVA).

### Box A4a (i) Mean values: how to do a general parametric two-way ANOVA

This is a little more complicated than a parametric one-way ANOVA, but in principle it is the same. Using the freshwater/marine fish example:

- The two-way design is cast as  $ij$  cells (here 4) formed from  $i$  columns (here  $i = 2$ ) and  $j$  rows (here  $j = 2$ ) in a table. Each cell has a number of replicate measurements,  $n$  per cell.  
Formulate the predictions. In this case for example, (a) *rows*: marine fish differ in growth rate from freshwater fish; (b) *columns*: male fish differ in growth rate from female fish; (c) *interaction*: water type and sex of fish interact to determine growth rate. The appropriate null hypotheses are that there are no differences among the rows, or among the columns, and that there are no interactions of any kind.
- Work out the mean values of all the replicates in each cell ( $B_{ij}$ ), and the overall means of each column ( $C_i$ ) and row ( $R_j$ ) in the two-way design. The grand mean of all the data is  $M$ .
- Calculate  $SS_{\text{rows}}$ , the sum of squares of the row mean values (not the raw data), following the method given in Box 3.1. Do the same for the column mean values to find  $SS_{\text{cols}}$ , and the cell means to find  $SS_{\text{grps}}$ .
- Calculate the interaction sum of squares,  $S_{\text{int}} = SS_{\text{grps}} - SS_{\text{rows}} - SS_{\text{cols}}$ .
- Calculate the sum of squares of each cell separately, and add them up to give  $SS_{\text{error}}$ .
- The degrees of freedom for columns  $d.f._{\text{cols}} = (i - 1)$ , for rows  $d.f._{\text{rows}} = (j - 1)$ , for the interaction  $d.f._{\text{int}} = (i - 1)(j - 1)$ . The error degrees of freedom  $d.f._{\text{error}} = ij(n - 1)$ .
- Calculate the mean squares:  $MS_{\text{rows}} = SS_{\text{rows}}/d.f._{\text{rows}}$ ;  $MS_{\text{cols}} = SS_{\text{cols}}/d.f._{\text{cols}}$ ;  $MS_{\text{int}} = SS_{\text{int}}/d.f._{\text{int}}$ ;  $MS_{\text{error}} = SS_{\text{error}}/d.f._{\text{error}}$ .
- Calculate the test statistic for each component:  

$$F_{\text{rows}} = MS_{\text{rows}}/MS_{\text{error}}$$
 with degrees of freedom of  $d.f._{\text{rows}}$ ,  $d.f._{\text{error}}$   

$$F_{\text{cols}} = MS_{\text{cols}}/MS_{\text{error}}$$
 with degrees of freedom of  $d.f._{\text{cols}}$ ,  $d.f._{\text{error}}$   

$$F_{\text{int}} = MS_{\text{int}}/MS_{\text{error}}$$
 with degrees of freedom of  $d.f._{\text{int}}$ ,  $d.f._{\text{error}}$
- Look up these values of  $F$  with the appropriate degrees of freedom in Table G of Appendix III to see whether they exceed the relevant critical values. Present the two-way ANOVA laid out in the standard manner.

10. Note that if the interaction is significant, then the row and the column effects don't really mean much, because the difference among the row groups will then depend on which column group it is, and vice versa.
11. You should be aware that there are two kinds of grouping factors that create the 'ways' of a one- or two-way ANOVA. The distinction doesn't make a difference in a one-way parametric analysis, but it does in a two-way. First, there are *fixed* factors, where the groups are fixed by the experiment, and are not intended to represent a few of a great range of possibilities; the hypothesis being tested is about whether real differences exist among the actual groups in the experiment. Many experimentally created groupings are of this type, such as a treated and a control group (which are not representative of a range of different possible groups). A *random* factor, on the other hand, has groups that are merely a sample of all possible groups, and are intended to represent this range of possibilities; here one is not really interested in whether there are differences among these particular groups, but rather in how much of the variation is contained among as opposed to within the groupings. Rearing sets of animals in different cages would be a good example: cage differences are not really treatments, but merely random variation that is not interesting but nevertheless must be allowed for. Sokal and Rohlf (1995, Section 8.4) and Underwood (1997) have good discussions of this distinction.

**Box A4a****(ii) Mean values: how to do a specific parametric two-way ANOVA**

As in the one-way case (Box 3.5a(ii)), we cannot specify a particular rank order that we expect, but we can make a priori contrasts. If these contrasts are specified in advance of obtaining the data, they allow  $(i - 1)$  contrasts to be made among column groups,  $(j - 1)$  contrasts among row groups, and  $(i - 1)(j - 1)$  contrasts involving the interaction of both row and column groups. As before, each contrast consists of a subset of the groups contrasted against another subset.

1. Formulate the predictions that you want to make in the form of contrasts, before collecting the data. You do this by expressing the prediction as an inequality, just as in Box 3.9b(ii). From rearranging the inequality, you obtain the relevant set of coefficients,  $\lambda_i$ , for each contrast. Unlike Box 3.9b(ii), however, you can either cast each contrast as a general or a specific contrast (using  $\neq$  or  $<$ , respectively, in the inequality, just as in Box 3.5a(ii), point 2).
2. Each contrast must be valid and independent of all the others (using the checks exactly as detailed in Box 3.5b(ii), point 3).
3. Do a general parametric two-way ANOVA first (*see* Box 3.9a(i)), since you are technically breaking the among-group differences down into specific independent contrasts ('decomposing the sums of squares' in statistical jargon). You will need the  $MS_{\text{error}}$  term from this analysis.
4. For each contrast, obtain the coefficients ( $\lambda_i$ ) and the mean values ( $m_i$ ) for each group, and follow exactly the method detailed for the one-way case in Box 3.5a(ii), points 5 and 6, to obtain the test statistic,  $t$ , and its significance for your contrast.
5. Repeat step 4 for each contrast.
6. Present the results of the two-way ANOVA, as well as the contrasts, laid out in the standard manner.

## Box A4b

## (i) Mean ranks: how to do a general non-parametric two-way ANOVA

1. Rank the data values in all cells combined and add up the ranks in each cell to give a rank total  $R$  for that cell.
2. We can now use these rank totals to test our *general* predictions:
  - (a) *Marine fish differ from freshwater fish.*  
Sum the rank totals for each *column* (see table on p. 119) separately giving an  $R_i$  value for marine and an  $R_j$  value for freshwater environments. Now calculate  $H$  as in the one-way analysis of variance (Box 3.4b), using the column  $R_i$  values and their appropriate  $n_i$  values (here there are eight values making up each  $R_i$ , hence  $n_i = 8$ ), and check the resulting value in an appropriate table (e.g. Appendix III, Table A) as if it were  $\chi^2$  for  $(i - 1)$  degrees of freedom.
  - (b) *Male fish differ from female fish.*  
Sum the rank totals for each *row* separately giving an  $R_i$  value for male and an  $R_j$  value for female fish and calculate  $H$  again as above. Again,  $n_i = 8$  in our example.
  - (c) *There is an interactive effect of water type and sex on the growth rate of fish.*  
This is an open-ended general prediction which combines both (a) and (b) above. It is asking whether there is any interaction between levels of grouping in determining growth rate. The calculation of  $H$  is exactly as above but the  $\sum R_i^2/n_i$  term includes the rank totals for *all* the cells in the table instead of just the columns or the rows to give  $H_{\text{tot}}$ . Here  $n_i$  is again the number of values making up each  $R_i$  ( $= 4$  in our example).  $H$  for the interaction,  $H_{\text{int}}$ , can then be calculated as:  $H_{\text{int}} = H_{\text{tot}} - H_{\text{water}} - H_{\text{sex}}$ , where  $H_{\text{water}}$  and  $H_{\text{sex}}$  refer to  $H$  values from (a) and (b) above. The degrees of freedom for  $H_{\text{int}}$  are  $\text{d.f.}_{\text{tot}} - \text{d.f.}_{\text{water}} - \text{d.f.}_{\text{sex}}$ , in this case, therefore,  $3 - 1 - 1 = 1$  ( $\text{d.f.}_{\text{tot}}$  is the total  $i$  (number of group means) minus 1).

## Box A4b

## (ii) Mean ranks: how to do a specific non-parametric two-way ANOVA

Testing our *specific* predictions is a little more complicated but still relatively straightforward. We give three illustrations below:

- (a) A prediction about the columns, e.g. *marine fish grow faster than freshwater fish* (or vice versa). In terms of the means in the table above this predicts that  $(B + D) > (A + C)$ . Another possibility is the opposite prediction, that freshwater fish grow faster than marine, i.e.  $(A + C) > (B + D)$ .
- (b) A prediction about the rows, e.g. *male fish grow faster than female fish* (or vice versa). This predicts  $(A + B) > (C + D)$  (or, conversely, for the prediction that females grow faster than males,  $(C + D) > (A + B)$ ).
- (c) A prediction about interaction, e.g. *the effect of water type on growth rate will be greater in male fish than in female fish*. This predicts that  $(A - B) > (C - D)$ . The converse (the effect will be greater in females) would, of course, predict  $(C - D) > (A - B)$ . This class of prediction is thus concerned with the *interaction* between water type and sex.

These are the predictions we can make about the relative sizes of the means; how do we arrive at the coefficients for testing? The procedure is as follows:

1. The first step is to rearrange the various predicted inequalities so that all the means are on the left, thus:

(a) The prediction about the effect of water type becomes  $-A + B - C + D > 0$  (for growth in marine  $>$  growth in fresh water) or  $+A - B + C - D > 0$  (for growth in fresh water  $>$  growth in marine).

We then substitute 1 with the appropriate sign for each letter so that we arrive at  $-1, +1, -1, +1$  or  $+1, -1, +1, -1$ , respectively.

(b) In the same way, the prediction about the effect of sex becomes  $+A + B - C - D > 0$  (for males  $>$  females) or  $-A - B + C + D > 0$  (for females  $>$  males) and the coefficients thus  $+1, +1, -1, -1$  or  $-1, -1, +1, +1$ .

(c) The interaction predictions must also be framed in this way. Thus the prediction  $(A - B) > (C - D)$  becomes  $+A - B - C + D > 0$  and the coefficients  $+1, -1, -1, +1$ . The prediction  $(C - D) > (A - B)$  becomes  $-A + B + C - D$  and the coefficients therefore  $-1, +1, +1, -1$ .

2. Then, testing these predictions

(a) *Marine fish grow faster than freshwater fish.*

Remember this means we are testing the prediction  $-A + B - C + D > 0$ . The coefficients  $\lambda_i$  thus become  $-1, +1, -1, +1$  so that the rank totals for each cell are weighted as follows:

$$\begin{aligned} L &= \sum \lambda_i R_i \\ &= (-1)(R_{\text{freshwater/male}}) + (+1)(R_{\text{marine/male}}) \\ &\quad + (-1)(R_{\text{freshwater/female}}) + (+1)(R_{\text{marine/female}}) \end{aligned}$$

$E$  and  $V$  can then be calculated as:

$$\begin{aligned} E &= (N + 1)(\sum n_i \lambda_i) / 2 \\ V &= (N + 1)[N \sum n_i \lambda_i^2 - (\sum n_i \lambda_i)^2] / 12 \end{aligned}$$

The test statistic  $z$  can then be calculated as:

$$z = (L - E) / \sqrt{V}$$

and checked against a table of  $z$ -values (see Appendix III, Table C).

(b) *Male fish grow faster than female fish.*

Now we are testing the prediction  $+A + B - C - D > 0$ , so  $\lambda_i$  becomes  $+1, +1, -1, -1$ . The calculation of  $L$ ,  $E$  and  $V$  and then the test statistic  $z$  can proceed as above, but with the new  $\lambda_i$  weightings.

(c) *The effect of water type is greater in males than in females.*

This tests the interaction prediction  $+A - B - C + D > 0$  using  $\lambda_i$  of  $+1, -1, -1, +1$ . Once again, follow the calculations above for  $L$ ,  $E$ ,  $V$  and  $z$ .

There is no standard layout for a non-parametric two-way analysis of variance; just quote the test statistic, its degrees of freedom and the probability.

## Examples of tests for a trend

### Box A5 How to calculate a correlation coefficient

Set out the two sets of data values to be correlated in pairs (remember, for each value in set 1 there must be a corresponding value in set 2). Thus, if we were looking for a correlation between height and weight in people, the data would be set out as below:

Person	Weight (kg)	Height (m)
1	63	1.8 (pair 1)
2	74	1.7 (pair 2)
3	60	1.9 (pair 3)
4	71	1.8 (pair 4)
.	.	.
.	.	.
.	.	.
		etc.

It may be that several values of one measure are paired with the same value of the other, for example when measuring some behaviour in several individuals from the same social group and using these values in a correlation of time spent doing the behaviour and group size. In this case, the data might be as follows:

Observation	Time spent in behaviour (s)	Group size
1	15.3	3 (pair 1)
2	17.1	3 (pair 2)
3	18.0	5 (pair 3)
4	6.0	5 (pair 4)
5	31.1	5 (pair 5)
.	.	.
.	.	.
.	.	.
		etc.

We can now calculate either a parametric (Pearson) or non-parametric (Spearman rank) correlation:

### Box A5 (i) Pearson correlation coefficient (parametric)

This tests for a *linear* association between two (bivariate-) *normally* distributed variables.

1. Formulate the prediction, either as a general (any) or a specific (positive or negative) association. You can test for either a general (two-tailed) or specific (one-tailed) correlation coefficient by using different threshold values of the test statistic.
2. Calculate  $S_{xx}$ ,  $S_{yy}$  and  $S_{xy}$  as shown in Box A1.
3. Calculate the test statistic,  $r = S_{xy} / \sqrt{(S_{xx} \times S_{yy})}$ .
4. Look up the threshold value for  $r$  using  $(n - 2)$  degrees of freedom in Table E of Appendix III, using either the one-tailed or two-tailed levels of significance, as appropriate to your hypothesis.

**Box A5****(ii) Spearman rank correlation coefficient (non-parametric)**

This tests for a *monotonic* (i.e. continuously increasing or decreasing) association between two variables, and works with ranking or constant-interval measurements, making no assumptions about the normality of the data. As with the Pearson coefficient, we can test for either a general (two-tailed) or specific (one-tailed) correlation by using different threshold values of the test statistic.

1. Rank the values for the first measure only, then rank the values for the second measure only.
2. Subtract second-measure ranks from first-measure ranks (giving  $d_i$ ) then square the resulting differences ( $d_i^2$ ) and calculate the Spearman coefficient as:

$$r_s = 1 - [(6\sum_i d_i^2)/(n^3 - n)]$$

where  $n$  is the number of pairs of data values.

3. If  $n$  is between 4 and 20 (4 is a minimum requirement), consult Appendix III, Table F for the appropriate sample size to see whether the calculated  $r_s$  value is significant. Remember that general (two-tailed) or specific (one-tailed) tests have different threshold values for  $r_s$ . Thus you must know whether you are looking for any association at all (general), or just a positive, or just a negative one (specific).
4. If  $n$  is greater than 20, a different test statistic,  $t$ , is calculated from  $r_s$  as:

$$t = r_s \sqrt{[(n - 2)(1 - r_s^2)]}$$

$t$  can then be checked against its own threshold values (Appendix III, Table D) for  $n - 2$  degrees of freedom.

A significant calculated value for  $r_s$  or, with large samples,  $t$ , allows us to reject the null hypothesis of no trend in the relationship between our two measures.

**Box A6****(i) How to do a parametric linear regression**

1. Calculate  $S_{xx}$ ,  $S_{yy}$  and  $S_{xy}$  as in Box A1.
2. Calculate the slope of the line as:  $b = S_{xy}/S_{xx}$ .
3. Calculate the intercept of the line on the  $y$ -axis as:  $a = \bar{y} - b\bar{x}$ , where  $\bar{y}$  is the mean of the  $y$ -values and  $\bar{x}$  is the mean of the  $x$ -values.

The line can now be fitted by calculating  $y = a + bx$  for some sample  $x$ -values and drawing it on the scattergram.

To calculate the standard error of the slope:

4. Calculate the *variance of  $y$  for any given value of  $x$*  as:

$$s_{y/x}^2 = [1/(n - 2)][S_{yy} - S_{xy}^2/S_{xx}].$$

5. The *standard error of the slope* is then:  $\sqrt{(s_{y/x}^2/S_{xx})}$ .



6. To find the test statistic  $F$ , calculate the following:

$$\text{Regression sum of squares (RSS)} = (S_{xy})^2/S_{xx}$$

$$\text{Deviation sum of squares (DSS)} = S_{yy} - (S_{xy})^2/S_{xx}$$

$$\text{Regression mean square (RMS)} = \text{RSS}/1$$

(1 is the value always taken by the *regression degrees of freedom*.)

$$\text{Deviation mean square (DMS)} = \text{DSS}/(n - 2).$$

( $n - 2$  is the value taken by the *deviation degrees of freedom*.)

$F$  is now calculated simply as  $F = \text{RMS}/\text{DMS}$  and its value can be checked against critical values in  $F$ -tables (Appendix III, Table G) for 1 ( $f_1$ ) and  $n - 2$  ( $f_2$ ) degrees of freedom.

7. To find  $y$  for new values of  $x$ :

Having established our regression equation, we might well want to predict  $y$  for other values of  $x$  that lie within the range we actually used in the analysis. Once we had then gone away and *measured*  $y$  for our new  $x$ -value we should want to see whether it departed significantly from its predicted value. Three steps are needed:

- (a) calculate the predicted  $y$ -value using the equation  $y = a + bx$  as when fitting the regression line, but this time use the new  $x$ -value ( $x'$ ) in which you are interested;
- (b) calculate the standard error (s.e.) of the predicted  $y$ -value as follows:

$$\text{s.e.} = \sqrt{[(s_{y/x}^2)(1 + 1/n + d^2/S_{xx})]}$$

where  $d = x' - \bar{x}$ ;


- (c) calculate the test statistic  $t$  as:

$$t = \frac{\text{observed } y - \text{predicted } y}{\text{s.e.}}$$

and look up the calculated value of  $t$  in  $t$ -tables (Appendix III, Table D) for  $n - 2$  degrees of freedom (where  $n$  is the number of pairs of data values in the regression). If  $t$  is significant it means the measured value of  $y$  departs significantly from the value predicted by the regression equation and might lead to interesting questions as to why.

### Box A6 (ii) How to do a non-parametric regression

A non-parametric approach to regression is little used (*see* Sokal and Rohlf, 1995, p. 539). There seem to be few advantages over using a simple rank correlation in cases where you do not want to predict a value of  $y$  from  $x$ , or to discover the actual equation of their relationship. However, some use is made of a non-parametric procedure (spline regression) fitted to a set of points to produce a smoothed description of highly non-linear irregular relationships (*see* Schluter, 1988 and Pentecost, 1999).



## Appendix III Significance tables

**Table A** Critical values of chi-squared at different levels of  $p$ . To be significant, calculated values must be *greater* than those in the table for the chosen level (0.05, 0.01, 0.001) of  $p$  and the appropriate number of degrees of freedom.

Degrees of freedom	Probability, $p$		
	0.05	0.01	0.001
1	3.841	6.635	10.83
2	5.991	9.210	13.82
3	7.815	11.34	16.27
4	9.488	13.28	18.47
5	11.07	15.09	20.51
6	12.59	16.81	22.46
7	14.07	18.48	24.32
8	15.51	20.09	26.13
9	16.92	21.67	27.88
10	18.31	23.21	29.59

**Table B** Critical values of Mann–Whitney  $U$  at  $p = 0.05$ . To be significant, values must be *smaller* than those in the table for appropriate sizes of  $n_1$  and  $n_2$ .

$n_1 \backslash n_2$	3	4	5	6	7	8	9	10	15	20
2	–	–	0	0	0	0	0	0	1	2
3	0	0	1	2	2	2	2	3	5	8
4		1	2	3	4	4	4	5	10	13
5			4	5	6	7	7	8	14	20
6				7	8	10	10	11	19	27
7					11	12	12	14	24	34
8						15	15	17	29	41
9							17	20	34	48
10							20	23	39	55
15								34	39	90
20								48	55	127



**Table D** Critical values of  $t$  at different levels of  $p$ . To be significant at the appropriate level of probability, values must be *greater* than those in the table for the appropriate degrees of freedom.

Degrees of freedom	Probability, $p$						(one-tailed) (two-tailed)
	0.05 0.10	0.025 0.05	0.01 0.02	0.005 0.01	0.001 0.002	0.0005 0.001	
1	6.314	12.71	31.82	63.66	318.3	636.6	
2	2.920	4.303	6.965	9.925	22.33	31.60	
3	2.353	3.182	4.541	5.841	10.21	12.92	
4	2.132	2.776	3.747	4.604	7.173	8.610	
5	2.015	2.571	3.365	4.032	5.893	6.869	
6	1.942	2.447	3.143	3.707	5.208	5.959	
7	1.895	2.365	2.998	3.499	4.785	5.408	
8	1.860	2.306	2.896	3.355	4.501	5.041	
9	1.833	2.262	2.821	3.250	4.297	4.781	
10	1.812	2.228	2.764	3.169	4.144	4.587	
11	1.796	2.201	2.718	3.106	4.025	4.437	
12	1.782	2.179	2.681	3.055	3.930	4.318	
13	1.771	2.160	2.650	3.012	3.852	4.221	
14	1.761	2.145	2.624	2.977	3.787	4.140	
15	1.753	2.131	2.602	2.947	3.733	4.073	
16	1.746	2.120	2.583	2.921	3.686	4.015	
17	1.740	2.110	2.567	2.898	3.646	3.965	
18	1.734	2.101	2.552	2.878	3.610	3.922	
19	1.729	2.093	2.539	2.861	3.579	3.883	
20	1.725	2.086	2.528	2.845	3.552	3.850	

**Table E** Critical values for the Pearson product-moment correlation coefficient  $r$ . Values must be *greater* than those in the table to be significant at the indicated level of probability.

Degrees of freedom	Probability, $p$					(one-tailed) (two-tailed)
	0.05 0.1	0.025 0.05	0.01 0.02	0.005 0.01	0.0005 0.001	
1	0.988	0.997	1.000	1.000	1.000	
2	0.900	0.950	0.980	0.990	0.999	
3	0.805	0.878	0.934	0.959	0.991	
4	0.729	0.811	0.882	0.917	0.974	
5	0.669	0.755	0.833	0.875	0.951	
6	0.622	0.707	0.789	0.834	0.925	
7	0.582	0.666	0.750	0.798	0.898	
8	0.549	0.632	0.716	0.765	0.872	
9	0.521	0.602	0.685	0.735	0.847	
10	0.497	0.576	0.658	0.708	0.823	
11	0.476	0.553	0.634	0.684	0.801	
12	0.458	0.532	0.612	0.661	0.780	
13	0.441	0.514	0.592	0.641	0.760	
14	0.426	0.497	0.574	0.623	0.742	
15	0.412	0.482	0.558	0.606	0.725	
16	0.400	0.468	0.543	0.590	0.708	
17	0.389	0.456	0.529	0.575	0.693	
18	0.378	0.444	0.516	0.561	0.679	
19	0.369	0.433	0.503	0.549	0.665	
20	0.360	0.423	0.492	0.537	0.652	
25	0.323	0.381	0.445	0.487	0.597	

**Table F** Critical values for the Spearman rank correlation coefficient  $r_s$ . Values must be *greater* than those in the table to be significant at the indicated level of probability.


$n$	Probability, $p$				
	0.05 0.10	0.025 0.05	0.01 0.02	0.005 0.01	(one-tailed) (two-tailed)
4	1.000				
5	.900	1.000	1.000		
6	.829	.886	.943	1.000	
7	.714	.786	.893	.929	
8	.643	.738	.833	.881	
9	.600	.700	.783	.833	
10	.564	.648	.745	.794	
11	.536	.618	.709	.755	
12	.503	.587	.671	.726	
13	.484	.560	.648	.703	
14	.464	.538	.622	.675	
15	.443	.521	.604	.654	
16	.429	.503	.582	.635	
17	.414	.485	.566	.615	
18	.401	.472	.550	.600	
19	.391	.460	.535	.584	
20	.380	.447	.520	.570	
21	.370	.435	.508	.556	
22	.361	.425	.496	.544	
23	.353	.415	.486	.532	
24	.344	.406	.476	.521	
25	.337	.398	.466	.511	

**Table G** Critical values of  $F$  at  $p = 0.05$  (upper) and  $p = 0.01$  (P146). To be significant, values must be *greater* than those in the tables for the appropriate degrees of freedom ( $f_1$  and  $f_2$ ).

		$p = 0.05$																		
$f_2 \backslash f_1$		1	2	3	4	5	6	7	8	9	10	12	15	20	30	$\infty$				
1	161.4	199.5	215.7	224.6	230.2	234.0	236.8	238.9	240.5	241.9	243.9	245.9	248.0	250.1	254.3					
2	18.51	19.00	19.16	19.25	19.30	19.33	19.35	19.37	19.38	19.40	19.41	19.43	19.45	19.46	19.50					
3	10.13	9.55	9.28	9.12	9.01	8.94	8.89	8.85	8.81	8.79	8.74	8.70	8.66	8.62	8.53					
4	7.71	6.94	6.59	6.39	6.26	6.16	6.09	6.04	6.00	5.96	5.91	5.86	5.80	5.75	5.63					
5	6.61	5.79	5.41	5.19	5.05	4.95	4.88	4.82	4.77	4.74	4.68	4.62	4.56	4.50	4.36					
6	5.99	5.14	4.76	4.53	4.39	4.28	4.21	4.15	4.10	4.06	4.00	3.94	3.87	3.81	3.67					
7	5.99	4.74	4.35	4.12	3.97	3.87	3.79	3.73	3.68	3.64	3.57	3.51	3.44	3.38	3.23					
8	5.32	4.46	4.07	3.84	3.69	3.58	3.50	3.44	3.39	3.35	3.28	3.22	3.15	3.08	2.93					
9	5.12	4.26	3.86	3.63	3.48	3.37	3.29	3.23	3.18	3.14	3.07	3.01	2.94	2.86	2.71					
10	4.96	4.10	3.71	3.48	3.33	3.22	3.14	3.07	3.02	2.98	2.91	2.85	2.77	2.70	2.54					
11	4.84	3.98	3.59	3.36	3.20	3.09	3.01	2.95	2.90	2.85	2.79	2.72	2.65	2.57	2.40					
12	4.75	3.89	3.49	3.26	3.11	3.00	2.91	2.85	2.80	2.75	2.69	2.62	2.54	2.47	2.30					
13	4.67	3.81	3.41	3.18	3.03	2.92	2.83	2.77	2.71	2.67	2.60	2.53	2.46	2.38	2.21					
14	4.60	3.74	3.34	3.11	2.96	2.85	2.76	2.70	2.65	2.60	2.53	2.46	2.39	2.31	2.13					
15	4.54	3.68	3.29	3.06	2.90	2.79	2.71	2.64	2.59	2.54	2.48	2.40	2.33	2.25	2.07					
16	4.49	3.63	3.24	3.01	2.85	2.74	2.66	2.59	2.54	2.49	2.42	2.35	2.28	2.19	2.01					
17	4.45	3.59	3.20	2.96	2.81	2.70	2.61	2.55	2.49	2.45	2.38	2.31	2.23	2.15	1.96					
18	4.41	3.55	3.16	2.93	2.77	2.66	2.58	2.51	2.46	2.41	2.34	2.27	2.19	2.11	1.92					
19	4.38	3.52	3.13	2.90	2.74	2.63	2.54	2.48	2.42	2.38	2.31	2.23	2.16	2.07	1.88					
20	4.35	3.49	3.10	2.87	2.71	2.60	2.51	2.45	2.39	2.35	2.28	2.20	2.12	2.04	1.84					

Table G (Cont.)

		$p = 0.01$																		
$f_2 \backslash f_1$		1	2	3	4	5	6	7	8	9	10	12	15	20	30	$\infty$				
1	4052	4999	5403	5625	5764	5859	5928	5982	6022	6056	6106	6157	6209	6261	6366	99.50				
2	98.50	99.00	99.17	99.25	99.30	99.33	99.36	99.37	99.39	99.40	99.42	99.43	99.45	99.47	99.50	26.13				
3	34.12	30.82	29.46	28.71	28.24	27.91	27.67	27.49	27.35	27.23	27.05	26.87	26.69	26.50	26.13	13.46				
4	21.20	18.00	16.69	15.98	15.52	15.21	14.98	14.80	14.66	14.55	14.37	14.20	14.02	13.84	13.46	9.02				
5	16.26	13.27	12.06	11.39	10.97	10.67	10.46	10.29	10.16	10.05	9.89	9.72	9.55	9.38	9.02	6.88				
6	13.75	10.92	9.78	9.15	8.75	8.47	8.26	8.10	7.98	7.87	7.72	7.56	7.40	7.23	6.88	5.65				
7	12.25	9.55	8.45	7.85	7.46	7.19	6.99	6.84	6.72	6.62	6.47	6.31	6.16	5.99	5.65	4.86				
8	11.26	8.65	7.59	7.01	6.63	6.37	6.18	6.03	5.91	5.81	5.67	5.52	5.36	5.20	4.86	4.31				
9	10.56	8.02	6.99	6.42	6.06	5.80	5.61	5.47	5.35	5.26	5.11	4.96	4.81	4.65	4.31	3.91				
10	10.04	7.56	6.55	5.99	5.64	5.39	5.20	5.06	4.94	4.85	4.71	4.56	4.41	4.25	3.91	3.60				
11	9.65	7.21	6.22	5.67	5.32	5.07	4.89	4.74	4.63	4.54	4.40	4.25	4.10	3.94	3.60	3.36				
12	9.33	6.93	5.95	5.41	5.06	4.82	4.64	4.50	4.39	4.30	4.16	4.01	3.86	3.70	3.36	3.17				
13	9.07	6.70	5.74	5.21	4.86	4.62	4.44	4.30	4.19	4.10	3.96	3.82	3.66	3.51	3.17	3.00				
14	8.86	6.51	5.56	5.04	4.69	4.46	4.28	4.14	4.03	3.94	3.80	3.66	3.51	3.35	3.00	2.87				
15	8.68	6.36	5.42	4.89	4.56	4.32	4.14	4.00	3.89	3.80	3.67	3.52	3.37	3.21	2.87	2.75				
16	8.53	6.23	5.29	4.77	4.44	4.20	4.03	3.89	3.78	3.69	3.55	3.41	3.26	3.10	2.75	2.65				
17	8.40	6.11	5.18	4.67	4.34	4.10	3.93	3.79	3.68	3.59	3.46	3.31	3.16	3.00	2.65	2.57				
18	8.29	6.01	5.09	4.58	4.25	4.01	3.84	3.71	3.60	3.51	3.37	3.23	3.08	2.92	2.57	2.49				
19	8.18	5.93	5.01	4.50	4.17	3.94	3.77	3.63	3.52	3.43	3.30	3.15	3.00	2.84	2.49	2.42				
20	8.10	5.85	4.94	4.43	4.10	3.87	3.70	3.56	3.46	3.37	3.23	3.09	2.94	2.78	2.42					



## Answers to self-test questions

1. No, this is not an appropriate analysis because it involves the multiple use of a two-group test (Mann–Whitney  $U$ -test). A significant result becomes more likely by chance the greater the number of two-group comparisons that are made. An appropriate analysis would be a non-parametric one-way analysis of variance, which allows a test of the specific prediction that body size will be greatest in Lake 1, intermediate in Lake 2 and least in Lake 3.
2. Regression analysis is appropriate when testing for cause and effect in the relationship between  $x$ - and  $y$ -values, and the  $x$ -values are established in the experiment. The data are measured on some kind of constant interval scale that allows a precise, quantitative relationship to be calculated. Because it depends on establishing a quantitative relationship, predicting new values of one variable from new values of the other also demands regression analysis. In other cases, correlation analysis is necessary. Non-parametric correlation analysis is appropriate for both these kinds of data, and others where  $x$ -values are merely measured, but yields only the sign and magnitude of the relationship and makes no assumptions about the cause-and-effect relationship between  $x$ - and  $y$ -values.
3. The investigator has introduced a new analysis into the Discussion (the correlation between nutrient flow rate and aphid production). The analysis should, of course, be in the Results section.
4. The information tells you that the investigator carried out a non-parametric one-way analysis of variance and that they tested a general prediction, thus using the test statistic  $H$  instead of  $z$ . The three degrees of freedom tells you that the analysis compared four groups and the  $p$ -value that  $H = 14.1$  at d.f. = 3 is significant at the 1 per cent level.
5. The researcher can conclude that there is a significant positive association between daily food intake and growth rate, but can't necessarily infer that increased growth rate is caused by greater food intake; it could be that faster growing pigs simply eat more food as a result.
6. No, a chi-squared test is not appropriate here because the data are constant interval measurements and not counts.



7. (a) A test statistic is calculated by a significance test and its value has a known probability of occurring by chance for any given sample size or number of degrees of freedom.
- (b) A ceiling effect occurs where observational or experimental procedures are too undemanding to allow a prediction to be tested; all samples approach the maximum value.
- (c) Statistical significance refers to cases where the probability that a difference or trend as extreme as the one observed could have occurred by chance, if the null hypothesis of no difference or trend is true, is equal to or less than an accepted threshold probability (usually 5 per cent, but sometimes 1 or 10 per cent).
8. No, the observer cannot conclude that male thargs prefer larger females just from this. It could be, for instance, that larger females are simply more mature and thus more likely to conceive. Alternatively, depending on how and when size was measured, pregnant females may be larger precisely because they are pregnant!
9. Significance tests provide a generally accepted, arbitrary yardstick for deciding whether a difference or trend is interesting. The yardstick is the probability that the observed difference or trend could have occurred by chance when there wasn't really such a difference or trend in the population. Random variation in sampling will mean that differences or trends will crop up from time to time just by chance.
10. A reasonable prediction would be that the rate of reaction would be lowest in Treatment A because no enzyme had been added and highest in the warmed enzyme/substrate mixture of Treatment D. By the same rationale, Treatment C should have a lower rate than Treatment B because it is cooler. The predicted order is thus  $A < C < B < D$ . A suitable significance test would be a specific form of a non-parametric one-way analysis of variance using  $z$  as the test statistic.
11. The consultant could try a two-way analysis of variance. The two levels of grouping would be 'housing condition' and 'breed' with three groups at the first level and four at the second. A dozen or so samples for each combination of housing and breed would be useful, though the same number of samples should be used in each case since general rather than specific predictions are being tested. One thing the consultant should be careful to do is distribute pigs from different families arbitrarily across housing conditions so that any effect of housing on growth rate is not confounded with family-specific growth rates (related pigs might grow at a similar rate because they inherited similar growth characteristics). The analysis would indicate any independent effect of housing and breed and any interaction between the two in influencing growth rate.

12. Although the figures look very similar, they are deceptive because their y-axes are scaled differently. The drop in numbers in felled deciduous forest represents 49 per cent of the number in unfelled forest. In coniferous forest it represents only 38 per cent. In proportional terms, therefore, the impact of felling seems to be greater in deciduous forest. However, the fact that the analysis is based on only single counts means it should act only as an exploratory analysis leading to a properly replicated confirmatory analysis.

13. A negative value of  $r$  indicates a negative correlation, i.e. the value of  $y$  decreases as that of  $x$  increases. The sign of the coefficient is ignored when checking against threshold values of significance.

14. Some predictions can be derived as follows:

(a) *Difference predictions*

**Observation** The distribution of individuals across species varied between different sites.

**Hypothesis** *Differences in the degree of dominance of species within a community vary with the ability of species to compete with others for limited resources.*

**Prediction** *Dominant species will be those whose individuals win in contests with individuals of other species over the resource they occupy.*

**Observation** Individuals of some species are smaller in some streams than in others and some streams are more polluted.

**Hypothesis** *Pollution results in reduced body size among some freshwater species.*

**Prediction** *The body size of any given species will be smallest in the most polluted stream and largest in the least polluted stream.*

(b) *A trend prediction*

**Observation** Fewer organisms were seen attached to the substrate in fast-flowing parts of the streams.

**Hypothesis** *Flow rate influences the ability of organisms to settle on the substrate.*


**Prediction** *If clean substrate is provided in areas of different flow rate, then fewer of the organisms drifting by will settle the faster the flow.*

15. Differences.

16. Yes, a  $1 \times 4$  chi-squared analysis could be carried out, but care would have to be taken in calculating expected values because of the different surface areas of the body sites and their possibly different degrees of vulnerability.

17. Since the experimenter had collected ten counts for each combination of 'parent' and 'sibling' treatment, a two-way analysis of variance would provide most information. It would allow either general or specific predictions about the effects of 'parent' and 'sibling' treatments and the interaction between them to be tested. The ten values in each case could be totalled and used in a  $2 \times 2$  chi-squared analysis but this would only test for an overall combined effect of the treatments; much useful information would thus be lost in comparison with the analysis of variance.
18. The threshold probability of 0.05 is an arbitrarily agreed compromise between the risk of accepting a null hypothesis in error (as would happen by setting the threshold  $p$ -value too high) and the risk of rejecting it in error (by setting the threshold too low). For many situations in biology, a threshold of  $p = 0.01$  would result in an unnecessary risk of accepting the null hypothesis when it was not true.
19. Clearly this is not a sensible procedure because it confounds the size of prey with the amount each barracuda has already eaten. Barracuda may give up at a certain size of fish simply because they are satiated, not because the fish is too big. Also size is confounded with species, some of which may be distasteful or be unpalatable in other ways. Again, therefore, barracuda may reject a fish for reasons other than size.
20. There are at least three logical flaws in this line of reasoning. First, if males are generally larger than females then their brains will be proportionally larger too; any comparison of brain sizes should thus be on a relative scale. Second, should the brains of males turn out to be relatively larger, there is no a priori reason to suppose this will affect learning or any other ability. Third, even if differences in brain size do produce differences in learning, training may overcome any such differences.
21. The analysis requires a test for a difference. There is one level of grouping (different woods), so a one-way analysis of variance of the effect of wood on adrenal gland size would be appropriate. The potentially confounding factor of age may be a problem, but as long as the distributions of adrenal gland size and age are normal, the ecologist could do a parametric analysis of variance and include age as a covariate. The result would then test for an effect of wood controlling for any confounding effect of age.
22. An order effect may occur when there is a systematic confounding of experimental treatment and time. Thus, testing for a difference in the performance of chicks in three different learning environments would suffer from an order effect if birds always encountered environment 1 first, then environment 2 and then environment 3. Any difference between environments may simply be due to being earlier or later in the sequence experienced by the birds (or earlier or later in the day/week/season if individual birds experienced only one of the environments).

23. On the face of it, things don't look good for the parasitologist's regression analysis. The parasite-burden data clearly aren't normal, so violate an important assumption about the  $y$ -axis variable in a linear regression. However, all is not lost. It may be possible to normalise the data by performing an appropriate transformation. Indeed, a  $\log_{10}$  transformation normalises the data quite nicely, as the figure below shows (a Kolmogorov-Smirnov one-sample test yielded the outcome  $z = 1.113$ ,  $n = 120$ ,  $p = 0.157$ , showing the parasite data did not depart significantly from a normal distribution). The parasitologist can therefore carry on with the planned analysis using the transformed data.




24. Pseudoreplication occurs where there is non-independence between measures purporting to be independent replicates. An extreme case would be taking repeated measurements of a particular character from the same individual, but more subtle pseudoreplication can arise, for example, when animals from the same cage or litter are treated as independent samples. Animals sharing a cage can have a profound influence on each other behaviourally, physiologically and even morphologically, and those from the same litter obviously share a genetic and experiential background. The number of cages or litters, rather than the number of individuals, thus determines the sample size in any statistical analysis.
25. A bit sneaky this one. Unlike most significance tests of difference, you are looking for a *non-significant* outcome when you compare the distribution of data with a normal distribution. If the comparison is not significant, it means your data do not depart from normality more than you would expect by chance. In the example the probability associated with the test outcome

was 0.0341, which is less than 0.05. The data thus differ significantly from normal, so the botanist will have to use transformed data if they want to do a parametric significance test, or use a non-parametric test instead.

26. The behavioural ecologist is testing for a difference between three experimental treatments (since the treatments involve the number of males, it could also be argued to be testing for a trend, but it is not clear a priori that any response by focal males would relate linearly to the number of competitors, so the analysis is probably best treated as one of difference). So an analysis of variance (ANOVA) of some kind would be the obvious choice. A crucial feature of the experimental design, however, is that each focal male was tested in each of the treatments, so we have a *repeated measures* design. Assuming data for spermatophore weight satisfies the test for normality, we can use a parametric repeated-measures ANOVA. Since males may be tempted to invest more when confronted with larger females (because they are likely to have more eggs to fertilise), female thorax widths could be included as a covariate in the analysis to control for this. Since the design is a repeated-measures one, we do not have to control for male size – the treatment effects are being measured within each male. If spermatophore weight is not normally distributed, then a non-parametric repeated-measures Friedman ANOVA could be used, but it isn't possible to include covariates in this case.
27. An obvious problem facing the psychologist is that the questionnaire data relating to attitudes will contain a lot of responses reflecting different aspects of sexual behaviour, many of which may be intercorrelated. So carrying out an analysis of each response individually would run into serious problems of non-independence and inflated estimates of significance. One approach to overcome this would be to subject the 40 responses to a principal components analysis to distil the data down to a set of mutually uncorrelated composite variables that each reflect different aspects of attitudes towards sex. The 40 separate responses may well boil down to two or three composite variables that can then be used as dependent variables in an analysis of variance by, for example, sex and ethnic group, with perhaps financial status (e.g. salary) and age as covariates.
28. As long as clearance rate is normally distributed, the idea of using regression analysis is reasonable, but what is problematic is the use of four separate analyses. The independent variables are very likely to be intercorrelated (e.g. maternal weight may well depend on local food availability), so analysing them as if they are all independent of each other risks a Type 1 error (see Box 3.14). A better approach would be to conduct a *multiple* regression where the independent variables are all included in a single analysis and their mutual intercorrelations taken into account in assigning significance.
29. Since many behavioural variables were recorded that are (probably) intercorrelated, the best approach would be to perform a *principal components*

*analysis* first to reduce the many variables to their main features. It is likely that this will produce two or three components (each a composite of the original variables) that can be interpreted in terms of combinations of the bees' responses. Each of these components can be used as a dependent variable in a one-way analysis of variance of flower treatment (with three levels: unmanipulated control, previously visited by the subject bee, previously visited by a different bee).

30. The treatment has four levels (three strengths of electromagnetic field and a control), and apart from the control, no information is available about the relative magnitudes of the experimental fields. Each response measure is independent of the others because each neurone preparation was used only once; it is thus clearly not a repeated-measures design. A simple one-way analysis of variance would thus be appropriate. The only issue is whether a general or a specific prediction should be tested. Unless there is a good a priori reason for expecting a particular directional effect of field strength, the only possible specific prediction is the contrast of control versus electromagnetic levels (i.e. coefficients of +3 [control] and -1 for each treatment): this is a sensible choice of specific contrast.



# Index

- Abbreviations 173
  - Abstract (section of report) 172
  - Academic journals 3, 166, 172, 196
    - citations 3, 170–2
  - Acclimation 6
  - Acknowledging sources 195
  - Acorns 1
  - Additive effects 118
  - Aggression 11, 12, 18–19, 20, 29, 42–3, 45, 66, 147–8, 162–3, 169, 174–9, 180, 185, 188–9, 190, 191
    - escalating 147, 175–9
  - Alcohol consumption 93–4
  - Alcohol dehydrogenase 106–10
  - Alternative explanation 2
  - Analogy 181–2
  - Analysis 12, 19–39, 49–151
    - confirmatory 38, 44, 49–151, 152
    - corroborative 166
    - exploratory 12, 19–39, 47, 152, 231
    - multivariate 59
  - Analysis of covariance (ANCOVA)
    - (*see also* SPSS) 99–101
  - Analysis of variance (ANOVA) (*see also* AQB and SPSS) 68, 84–94, 103–20, 143, 145, 146, 150, 153–6, 160, 173, 198, 200, 211–19, 229, 230, 232, 234, 235
    - fixed factors in 108, 110
    - interactions 110, 114, 117, 120, 143, 160–1
    - multifactor 110, 121, 133
    - multivariate 133
    - one-way 84–94, 101, 211–16, 231
    - post-hoc* multiple comparisons within 87
    - random factors in 108
    - two-way 103–20, 216–19, 231, 232
  - Animal Behavior Society (ABS) 8
  - Anogenital sniffing 31, 32
  - Anthelmintic 45, 199
  - Antibiotic 102
  - Aphid 201
  - Appendix (section of report) 173
  - AQB (statistical package) 67, 69, 72
    - correlation analysis 121
    - general parametric one-way analysis of variance 85
    - general non-parametric two-way analysis of variance 117
    - general parametric two-way analysis of variance 106–8
  - Kruskal-Wallis test 91
  - linear regression 126–7
  - Mann-Whitney *U*-test 76–7
  - 1 × 2 chi-squared test 70
  - 1 × n chi-squared test 102
  - one-tailed *t*-test 75
  - one-tailed paired samples *t*-test 81
  - repeated-measures non-parametric analysis of variance 97–8
  - repeated-measures parametric analysis of variance 95–6
  - specific non-parametric one-way analysis of variance 93
  - specific parametric one-way analysis of variance 88–9
  - specific parametric two-way analysis of variance 112–14
  - 2 × 2 chi-squared test 103
  - two-tailed *t*-test 73
  - two-tailed paired samples *t*-test 80
  - Wilcoxon matched-pairs signed ranks test 82
- Asking questions 3, 7
    - sources 10–12

- Aspergillus* 126–9  
 Association for the Study of Animal Behaviour (ASAB) 8  
 Assumptions 6  
 Audience 184–5  
 Avian malaria 10
- Babesia microti* 16, 21, 22, 40, 47, 51  
 Background 5, 166  
     essential 5  
 Bacterial culture 65  
 Bacteriophage 155  
 Bank vole (*Clethrionomys glareolus*) 16, 19, 21, 46, 51, 144, 180–2  
 Bar chart (*see also* Presenting information) 19–22, 25, 27, 28, 48, 153, 155, 157  
 Basal metabolic rate (BMR) 99  
 Binomial test 79  
 Biosis 3, 5  
 Blackthorn (*Prunus spinosa*) 15–16  
 Blocking 79, 81, 82, 95, 107  
 Blood group 5  
 Blood sample 144  
 Blood smear 16  
 Body size 99–101, 137, 140, 147, 166  
 ‘Box and whisker’ plot 26, 27  
 Bumblebee (*Bombus* sp.) 85–7
- Carcinogen 6  
 Casual reflection 1  
 Catastrophe theory 141  
 Caterpillar 15–16  
 Cattle 15–16  
 Cause and effect 46, 120–1, 124–5, 126, 151  
 Cell 6, 15  
     biology 15  
     blood 16, 19, 52  
     physiology 6  
 Chaffinch (*Fringilla coelebs*) 9, 10  
 Charles Darwin 9  
 Chemical ablation 146  
 Chicken (*Gallus gallus domesticus*) 69, 102, 160–1  
 Chi-squared ( $\chi^2$ ) test (*see also* AQB and SPSS) 57, 60, 68–72, 91, 101–5, 150, 155, 160–2, 198, 200, 223, 231
- Chromosomal abnormality 44  
 Clover (*Trifolium repens*) 15–16  
 Comparative anatomy 15  
 Competition 40, 129  
     interspecific 40  
 Conclusions 170  
     section of report 172–3  
 Conference 12, 182, 196  
 Confidence limit 22, 25–6, 87, 163  
     to median 207  
     to regression 164  
     *r*-value 26  
 Confounding effects 141, 143, 151, 201  
 Constraint 5  
 Contrasts 67, 93, 95, 111–15  
     *a priori* 87  
     coefficients 88, 119  
     Helmert 89, 114  
     hierarchical set 88  
     independent (orthogonal) 67, 88, 110  
     planned 87  
     single degree of freedom 118  
     specific 235  
 Control 79–80, 95, 98, 154, 155, 156, 169, 200  
 Coping 6  
 Copulation 3  
 Correlation (*see also* AQB and SPSS)  
     56, 80, 121–4, 135–6, 137, 140, 164, 166, 220, 229  
     Pearson product-moment correlation (*r*) 121–2, 151, 220–1, 225  
     Spearman rank correlation (*r<sub>s</sub>*) 121, 147, 148, 151, 221, 226  
 Corticosterone 31, 37, 40, 46, 145, 146, 180  
 Cortisol 4  
 Courtship 3  
 Covariate 99–101, 121, 145, 232, 234  
 Cricket 6, 20, 22, 66  
     field (*Gryllus bimaculatus*) 18–19, 29, 42–3, 45, 47, 147–8, 174–9, 185, 188–9, 190, 191, 194  
 Curiosity 11  
*Cynips divisa* 129–33  
 Cytotoxic drugs 158
- Dandelion (*Taraxacum officinale*) 15–16



- Data (*see also* Measurement) 9
  - collection 9, 15–19, 169
  - counts 68–72, 101, 155
  - entry 26, 30–4, 60, 103, 107, 109, 113, 117, 126, 128
  - ‘eyeballing’ 19, 27, 38, 51, 138
  - independence 20, 72, 78, 164
  - multivariate 133
  - rank (ordinal) 26
  - raw 23, 152, 169, 196
  - reduction 125
  - sample 34
- Deer 15–16
- Degrees of freedom 68, 77, 88, 89, 165, 230
- Development 6, 11
- Developmental instability 180
- Difference (*see* Prediction)
- Discussion (section of report) 170, 229
- Display 9, 10
  - anti-predator 10
  - food defence 10
- Distastefulness 142
- Diversity 201
- Dogrose (*Rosa canina*) 15–16
- Dominance 45, 69
  - Mendelian 69
  - social 45, 162
- Dragonfly 6
- Duodenal ulcer 5
  
- Ectoparasite 16, 145
- Effect size 53, 138–9
- Endangered species 8
- Enzyme activity 106–10
- Epiphyte 159
- Escherichia coli* 155
- Ethics 7–8
  - codes of practice 8
  - dilemmas 8
  - guidelines 8
- Excel (Microsoft) 29, 67, 85
  - calculating new variables in 32–3
  - data entry using 30–4
  - exploratory analyses in 34–7
  - pivot tables 35–7
  - reading into SPSS 33
  - using formulae in 29, 32, 34–7, 38
- Exoskeleton 148
  
- Experience 3
  - personal 3
- Experiment 1, 6, 44, 49, 141, 142, 147, 180
- Experimental design 106
- Experimental manipulation 11, 15, 120
- Explanatory legend 153
  
- Fecundity 146
- Field course 15, 182
- Fish 119
- Fitting a model 60
- Flea 16, 145, 181
- Flight speed 121
- Floor and ceiling effects 140–1, 144, 151, 230
- Fluctuating asymmetry 180, 181
- F*-ratio 73, 74, 85–90, 103–21, 125–33, 160–1, 165, 227–8
- Free radicals 95
- Frequency distribution 27–9, 35, 37, 38, 55, 57
  - bimodal 57
  - binomial distribution 59
  - kurtosis 35, 57
    - leptokurtic 57
    - platykurtic 57
  - normal (Gaussian) distribution 38, 55–8, 153, 234
    - bivariate 121
    - multivariate test 96
    - testing for (*see also* Residuals) 58
  - Poisson distribution 59
  - skewness 35, 57
  - unimodal 28
- Friedman test (*see also* SPSS) 234
- Fungicide 49–52
  
- Gall 129–33
- Gamete 69
- Gammarus pulex* 201
- General test (*see also* Two-tailed test) 65–7, 73, 150
- GenStat (statistical package) 29, 59
- Germination 161
- Getting attention 189
- Gibberellin 153–6, 165
- GLIM (statistical package) 29, 59
- Goldenrod (*Solidago canadensis*) 15–16

- Google 3  
 Growth rate 119, 126, 153–6, 165  
 Guppy (*Poecilia reticulata*) 3, 5  
  
 Heavy metal 17–18, 42, 146  
 Herbicide 156  
 Herbivore 15–16, 40, 46, 141  
 Heterozygosity 69  
 Histogram 28, 29, 38, 60  
 Histology 7, 11, 141  
 Hormone 4  
   effect on resistance 146  
   slow-release implant 145  
   steroid 4  
   ‘stress’ 145  
 Host-parasite relationships 45, 46–7, 144–5  
 House sparrow (*Passer domesticus*) 121–3, 162–3  
 Hunger 144  
 Hypothesis 2, 9, 12, 48, 149, 150, 152  
   alternative 12  
   development of 9  
   framing 12, 14–48  
   mutually exclusive 10  
   null 43–4, 52, 66, 70, 72, 73, 75, 79, 85, 95, 97, 115, 117, 127, 138, 150  
   ‘straw man’ 10  
   testable 15  
   testing 2, 9–10, 44, 49–151  
  
 Imagery 181–2  
 Immune depression 4  
 Immune responsiveness 47  
 Information 3  
   acquiring 3  
   unregulated 3  
 Insemination 6  
 ‘Intelligent layman’ 180  
 Internet 3, 4–5, 30, 67, 195  
 Interquartile range 26, 38  
 Intracellular probe 6  
 Introduction (section of report) 168  
 Invertebrates 7, 40  
 Investigation 12  
   independent 12  
 Investigative skills 2  
  
 Jargon 180–1, 197  
 Jonckheere-Terpstra test 93  
   standardized 94  
  
 Kolmogorov-Smirnov test 57, 60, 62  
   one-sample 57, 60, 233  
 Kurtosis (*see* Frequency distribution)  
  
 Laboratory 5, 143  
 Lay summary 172  
 Least significant difference (LSD) 87  
 Levene’s test 73  
 Life history 6  
 Light absorbance 88  
 Literature database (*see also* Web of Science, Biosis, PubMed) 3, 4–5, 172  
 Locust 23, 24  
  
 Mann-Whitney *U*-test 68, 72, 76–8, 84, 150, 153, 155, 198, 210, 224, 229  
 Mathematical model 12, 173  
 Mathematical performance 110  
 Mating 3, 6, 43, 185  
   preference 5, 43  
 Mate-guarding 6  
 Mauchly’s test of sphericity 97  
 Mean 22–5, 27, 28, 35, 38, 59, 73–6, 85, 108, 114, 120, 153, 163, 164  
   arithmetic 56  
   confidence in 24  
   predicted order of 66, 67, 75, 88, 110  
   sample 23  
   true 23  
 Measurement 14–19, 49, 58, 137, 148, 151  
   constant interval 55, 58, 72, 99, 110, 121, 124, 147  
   independence of 58, 121  
   methods of 169  
   nominal/classificatory 54  
   rank (ordinal) 26, 54, 58, 72, 124, 147–8  
   types of 54–5  
   units of 166–7  
 Median 22, 25–6, 27, 35, 38, 76–7, 163, 164  
   predicted rank order 93–4, 97, 117–18

- Mendelian inheritance 69–72
- Methods (section of report) 168–9
- Migration 11
- Minitab (statistical package) 29
- Minnow (*Phoxinus phoxinus*) 163
- Mode 28
- Molecular genetics 15
- Monotonic relationship 123
- Moorhen (*Gallinula chloropus*) 165
- Mortality 45
  - larval 45
- Motivation 9, 110–16
  - sexual 9
- Mouse (*Mus musculus*) 4, 6, 29, 32, 181
  - laboratory 4, 6, 82, 158
- Multiple comparisons (*see* Analysis of variance)
- Mutagen 44
- National Centre for the Three Rs (UK) 8
- Natural selection 181
- Navigation 11
- Nematode 16, 45, 47, 146
  - eggs 41, 45, 46, 147
  - gut 16, 141, 145
  - larvae 147
  - soil-dwelling 17–18, 41–2, 45
- Niche 41
- Nitrogen 79
- Non-independence (*see also* Data and Measurement) 78, 104, 141–2
- Non-parametric tests (*see* Significance tests)
- Normal distribution (*see* Frequency distribution)
- Nutrient 5, 142
  - flow rate 142
  - transfer 5
- Oak (*Quercus rober*) 129–33
- Observation 1, 14, 18–19, 49, 141, 146
  - casual 11
  - exploratory 2, 11
  - open-ended 47
- Observational notes 14–19, 47
- Observer 14
- Objectivity 2
- Oestrous cycle 6
- One-tailed test (*see also* Specific test) 57, 65–7, 126–7, 139–40, 150
- Online search 3, 4–5, 172
- Open-endedness 2
- Order effect 140, 232
- Organophosphate 17–18, 41–2, 45, 146
- Outlier values 26, 56
- Ovipositor 18
- Ox-eye daisy (*Leucanthemum vulgare*) 15–16
- Palatability 144
- Parameter 51
- Parametric tests (*see* Significance tests)
- Parasite 3, 106–10, 144, 163, 181
  - passaging 6
  - worms 3, 6, 16, 140–1
    - eggs 3, 6, 16, 17–18
- Percentages 57, 68
- Periwinkle (*Littorina* sp.) 76–8
- Personal communication 172
- Personal observation 172
- Phloem 142
- Plagiarism 195–6
  - anti-plagiarism software 196
- Plantain (*Plantago lanceolata*) 15–16
- Poppy (*Papaver rhoeas*) 15–16
- Population 42, 56
  - size 42
  - true 56
- Poster presentation 182–95
- Post hoc* tests (*see* Analysis of variance)
- Potato 49–52
- PowerPoint (Microsoft) 30, 192–5
  - general advice for using 192–5
  - preparing a poster in 195
- Power test 8, 137–40
  - a priori* 139
  - G\*Power (package) 139
  - post hoc* 140
- Practical 2, 15
  - open-ended 2
  - self-driven 2
  - set-piece 2
- Practicality 7
- Prediction 9, 10, 12, 48, 149, 150
  - background 3
  - framing 14–48
  - general 65–7, 76–7, 84, 117, 120, 155

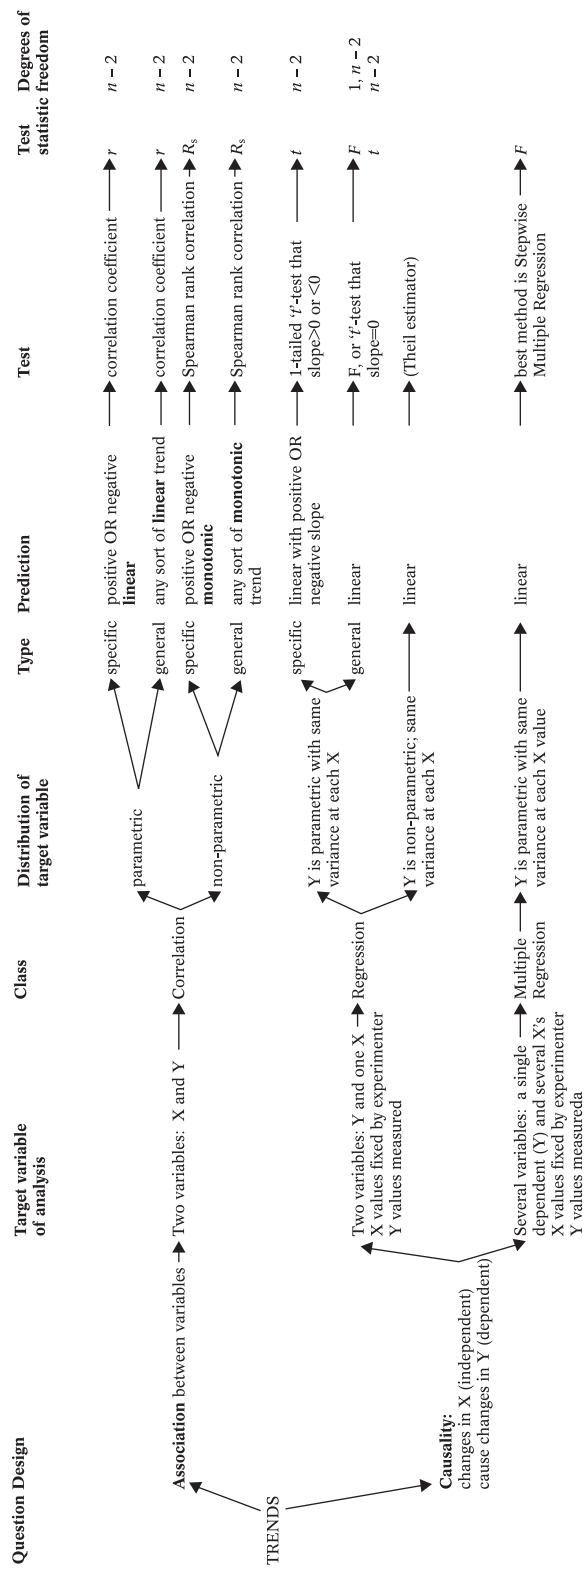
- Prediction (*continued*)  
 of a difference 44–8, 49–151, 199, 231  
 of a trend 44–8, 49–151, 199, 231  
 preparation 3, 6  
 refinement 149  
 specific 65–7, 76–7, 84, 113, 120, 122, 155, 229
- Presenting written information 2, 152–96  
 as figures and tables 152–65  
 analyses of difference 153–62  
 analyses of trends 162–5  
 labelling 157  
 scaling 157  
 in the text of a report 165–6  
 standard format for written report 166, 167–73, 180  
 example report 174–9  
 writing for a general readership 180–2  
 example report 183
- Principal components analysis 133–7, 234
- Probability 51–3, 153  
 exact 52, 83, 91  
 one-tailed 76, 94, 122  
 threshold 52, 66, 68, 138, 230  
 presented as asterisks 155  
 two-tailed 75
- Proportions 57
- Pseudoreplication 141, 233
- Psychology 15
- Public Understanding of Science 180, 196
- Published literature 12
- PubMed 3
- Questions (in spoken presentations) 194
- R (statistical package) 29, 59, 88
- R.A. Fisher 5
- Ragwort (*Senecio jacobaea*) 15–16
- Rainforest 159
- Randomness 59
- Rat (*Rattus norvegicus*) 45, 106–10
- Rattlesnake (*Crotalus crotalus*) 168, 170
- Received wisdom 5
- Recruitment 40
- References (section of a report) 170–2
- Regression (*see also* AQB and SPSS) 55, 122, 143, 146, 147, 148, 151, 229, 233, 234  
 equation 126, 130  
 line 62, 162, 164  
 of best fit 125, 162, 164  
 slope 125, 126–7  
 linear 63, 124–33, 221–2  
 multiple 125, 129–33  
 non-parametric (spline) 222  
 stepwise 130–3
- Repeated measures analysis (*see also* AQB and SPSS) 81, 117, 142–3, 199
- Replication 79, 137, 151, 162, 163, 199
- Report 2  
 written 2
- Research talk 12
- Research project 15
- Residuals 60  
 saving 64, 130  
 standardised 63  
 testing for normality 60–5, 73, 78, 88, 93, 106, 110, 117, 126, 130  
 unstandardised 63
- Resistance 145, 155
- Results (section of report) 165, 229
- Richard Dawkins 181, 184
- Round robin testing 84
- Sample 8, 51  
 chance bias in 51, 230  
 independent 91  
 related 81, 84  
 size (*see also* Power test) 8, 53, 57, 68, 77, 117, 137–40, 153, 158, 159, 165, 230
- SAS (statistical package) 59
- Scale 55, 88  
 logarithmic 88  
 non-arbitrary 55
- Scattergram 19–22, 38, 48, 49, 53, 124, 126
- Science 1  
 doing 1–13
- Scientific training 2
- ‘Selfish gene’ metaphor 181

- Sex ratio 11
- Shapiro-Wilk test 60, 62, 63, 73, 75, 76, 79, 82, 85, 88, 93, 99, 106, 110, 117, 122, 126
- Siamese fighting fish (*Betta splendens*) 29, 78
- Significance tests 51–151, 152, 230  
 choice of 170  
 non-parametric 58–65  
 parametric 55–8
- Skewness (*see* Frequency distribution)
- Slug 15–16, 39, 142–3
- Snail 15–16
- Social experience 31
- Song 173
- Specific test (*see also* One-tailed test, AQB and SPSS) 65–7, 75–6, 150
- Speculation 1, 169  
 formalized 1
- Sperm 6
- Spermatophore 6, 234
- Spreadsheet 29, 60
- S-Plus (statistical package) 29, 59
- Spoken presentation 182–95
- SPSS (Statistical Packages for the Social Sciences) 29, 30, 59, 60, 67, 69, 72  
 analysis of covariance 99–101  
 calculating new variables in 34  
 calculating summary statistics 37–8  
 correlation analysis 122–3  
 data view 33  
 exploratory analyses using 37–8  
 Friedman test 98–9, 117  
 general parametric one-way analysis of variance 85–7  
 general non-parametric two-way analysis of variance 117–18  
 general parametric two-way analysis of variance 106–10  
 linear regression 127–9  
 $1 \times 2$  chi-squared test 70–1  
 $1 \times n$  chi-squared test 102  
 one-tailed exact test 78  
 one-tailed *t*-test 75  
 one-tailed paired samples *t*-test 81  
*post-hoc* multiple comparisons 87  
 principal components analysis 133–6  
 reading Excel files into 33–4  
 repeated-measures parametric analysis of variance 96–7  
 specific non-parametric one-way analysis of variance 93–4  
 specific parametric one-way analysis of variance 89–90  
 specific parametric two-way analysis of variance 114–16  
 stepwise multiple regression 130–3  
 testing for normality in 60–5  
 $2 \times 2$  chi-squared test 104–5  
 two-tailed *t*-test 73–4  
 two-tailed paired samples *t*-test 80  
 variable view 33
- Squirrel 1
- Staining 7
- Standard deviation 24, 56  
 squared 57
- Standard error 22–5, 27, 38, 85, 97, 108, 114, 153  
 common 107  
 of regression line 126, 127
- Statgraphics (statistical package) 29
- Statistic 51
- Statistical power (*see also* Power test) 58, 67, 108, 137–40
- Statistical significance 49–151  
 by chance 84
- Statistics (*see also* AQB and SPSS) 13, 19  
 confirmatory 38  
 summary (descriptive) 19–39, 48, 80, 99, 153, 162, 164  
 tests of significance 26
- Steven Jay Gould 181
- Stickleback (*Gasterosteus aculeatus*) 169
- Stomach cancer 103
- Stork 123, 124
- Stress 182
- Suffering 7, 8
- Sums of squares 208
- Survival 42
- ‘Take-home’ message 173, 184, 196
- Tapeworm 181
- ‘Telegraphese’ 186
- Territory 9, 11, 45, 99

- Testosterone 4, 46, 144, 145, 146, 180  
Test statistic 68, 153, 159, 165, 230  
The 'Three Rs' (Replacement, Reduction, Refinement) 8  
T-helper cells 158  
Tick 16, 21, 47, 51, 145  
Time-lapse photography 144  
Transformation 57, 63, 145, 233  
    arcsine square root 58, 76  
    logarithmic 58, 63, 126, 233  
    square root 58, 63  
Trend (*see* Prediction)  
Trophic group 45  
*t*-test 67, 72–6, 84, 125–33, 139, 150, 198, 225  
    one-tailed 209–10  
    paired samples 79–81  
    two-tailed 208–9  
Turkey 54  
Two-tailed test (*see also* General test, AQB and SPSS) 57, 65–7, 126–7, 150  
Type I and Type II error 138–40, 234  
Undirected fact gathering 9  
Vaginal smear 6  
Variability 22  
Variance 59  
    equality (homogeneity) of variances 60, 73, 86–7, 99, 107, 115  
    proportion explained ( $r^2$ ) 126  
    adjusted  $r^2$  130  
Victor Hugo 186  
Vital dye 5  
  
Web of Science 3, 4–5  
Welch test 87  
Wilcoxon matched-pairs signed ranks test 79, 82–3  
Within-subjects effects 97  
Wood mouse (*Apodemus sylvaticus*) 99  
Worldwide Web (*see* Internet)  
  
Yahoo 3  
  
 $z$  57, 76–7, 78, 83, 93, 118, 157, 224, 229, 233  
Zero 55  
    arbitrary 55  
    true 55



# Quick Test Finder: Tests for Trends



**Question Design**

**Target variable of analysis**

**Class**

**Distribution of target variable**

**Type**

**Prediction**

**Test**

**Test statistic**

**Degrees of freedom**



