


OXFORD

experimental design for the life sciences

third edition



Graeme D. Ruxton
Nick Colegrave

Experimental design for the life sciences


THIRD EDITION

Graeme D. Ruxton

University of Glasgow

Nick Colegrave

University of Edinburgh



OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan South Korea Poland Portugal
Singapore Switzerland Thailand Turkey Ukraine Vietnam

Published in the United States
by Oxford University Press Inc., New York

© Graeme D. Ruxton and Nicholas Colegrave, 2011

The moral rights of the authors have been asserted
Database right Oxford University Press (maker)

First edition 2003

Second edition 2006

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,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose the same condition on any acquirer

British Library Cataloguing in Publication Data

Data available


Library of Congress Cataloging in Publication Data

Data available

Typeset by MPS Limited, a Macmillan Company
Printed in Great Britain on acid-free paper by
CPI Antony Rowe, Chippenham, Wiltshire

ISBN 978-0-19-956912-0

1 3 5 7 9 10 8 6 4 2



OXFORD
UNIVERSITY PRESS

Experimental design for the life sciences

To Hazel and Becky

Preface

How to read this book

This book is an introduction to experimental design. We mean it to be a good starter if you have never thought about experimental design before, and a good tune-up if you feel the need to take design more seriously. It does not come close to being the last word on experimental design. We cover few areas of design exhaustively, and some areas not at all. We use the Bibliography to recommend some good books that would facilitate a deeper consideration of design issues. That said, it is also important to realise that the basic ideas of design covered in this book are enough to carry out a wide range of scientific investigations. Many scientists forge a very successful career using experiments that never take them outside the confines of the material covered in this book. This book will also help you tackle more advanced texts, but if you absorb all that we discuss here, then you may find that you know all the experimental design that you feel you need to know.

This book is about how to design experiments so as to collect good quality data. Of course, those data will almost certainly need statistical analysis in order to answer the research questions that you are interested in. In order to keep the size of this book manageable, we do not enter into the details of statistical analysis. Fortunately, there are a huge number of books on statistical analysis available: we even like some of them! We recommend some of our favourites in the Bibliography. We also provide some pointers in the text to the types of statistical tests that different designs are likely to lead to.

We often illustrate our points through use of examples. Indeed, in some cases we found that the only way to discuss some issues was through examples. In other cases, our point can be stated in general terms. Here we still use examples to amplify and illustrate. Although we think that you should aim to read the book right through like a novel at least once, we have tried to organize the text to make dipping into the book easy too.

Our first job will be to remove any doubt from your mind that experimental design is important, and so we tackle this in Chapter 1. Chapter 2 discusses how good designs flow naturally from clearly stated scientific questions. Almost all experiments involve studying a sample and extrapolating conclusions about that sample more widely; how to select a good sample is the key theme of Chapter 3. The nitty-gritty of some simple designs forms the basis of Chapter 4. There is no point in formally designing an elegant experiment if you then make a poor job of actually collecting data from your sample individuals, so Chapter 5 is about some tips on taking effective measurements. Chapter 6 is a compendium of slightly more specialized points, which do not fit naturally into the other chapters but at least some of which ought be useful to your branch of the life sciences.

This is *not* a maths book. Nowhere in this book will you find the equation for the normal distribution, or any other equation for that matter. Experimental design is not a subsection of maths, you don't need to be a mathematician to understand simple but effective designs, and you certainly don't need to know maths to understand this book.

Good experimental design is vital to good science. It is generally nothing like as difficult as some would have you believe: you can go a long way with just the few simple guidelines covered in this book. Perhaps most amazingly of all, it is possible to derive enjoyment from thinking about experimental design: why else would we have wanted to write this book!

On the second edition

The coverage has not changed dramatically from the first edition, although all sections have been rewritten—sometimes extensively so—for greater clarity. We have also attended to filling in a few gaps, most notably in substantially increasing our consideration of the special challenges associated with using human subjects. However, the big change in this edition is the way in which the material is presented. As well as extending the use of some of the features of the previous edition, we also include a number of new features that we hope will substantially increase the ease of use of this book. We introduce these in the following pages.

On the third edition

We have very much retained the format of the second edition and cover much of the same material. However, we have again taken advantage of feedback to rewrite sections for improved clarity. This should be especially apparent in Chapter 3, where we have extensively revised and expanded our coverage of the key subjects of between-individual variation, replication, and statistical power, using an entirely new example for illustration throughout the chapter. We have also tried to rectify some omissions from the previous editions. In particular, we have always discussed in Chapter 2 how to go from a biological question to hypotheses, and then the design of an experiment. However, in previous editions, we had ducked the issue of where questions come from in the first place. Since this is an issue that early-career scientists in particular have challenged us to address, we now offer some general guidance on this issue. Another notable inclusion in this edition is greater consideration of human biomedical studies. In particular, we have added sections on how correlative studies can be particularly important in studying human epidemiology and on the specialist terminology used in the biomedical literature to describe experimental designs. We have also increased our discussion of the advantages of Latin square designs and increased our coverage of how to make measurements (for example, introducing the concepts of sensitivity and specificity of a test). Finally, we have increased our use of figures throughout, and have greatly enhanced the captions of each figure to allow figures to be understood without consultation of the surrounding text.

Learning features

Key definitions

There is a lot of jargon associated with experimental design and statistical analysis. We have not tried to avoid this. Indeed, we have deliberately tried to introduce you to as much of the jargon as we can. By being exposed to this jargon, we hope that it will become second nature to you to use—and understand—it. This should also make reading more advanced texts and the primary literature less daunting. However, to make negotiating the minefield of jargon more straightforward, we have increased the number of definitions of key terms provided and increased the level of detail in all definitions. Each key word or phrase is emboldened at the point where we first use it, and is given a clear definition in a box nearby.

Reverse causation

The second problem of correlational studies is **reverse causation**. We see a relationship between factors A and B. It might be that factor A influences factor B, when in fact it is change in B that influences A. For example, imagine a survey shows that those who use recreational drugs also consider themselves to have a financial problem. It is tempting to conclude that a drug habit is likely to cause financial problems. A more plausible explanation is that people who have financial problems are more likely to turn to drugs (perhaps as a way to temporarily escape their problems).

Reverse causation is mistakenly concluding that variable A influences variable B when actually it is B that influences A.

Statistics boxes

To emphasize the important link between good experimental design and statistics we now include a number of statistical boxes. These boxes should help you to see how thinking about design helps you think about statistics and vice versa. We have added the boxes at points where we think that keeping the statistics in mind is particularly useful and include pointers in the main text to each box.

STATISTICS BOX 4.1 Types of measurement

There are different types of data. It is common to split data into the following types:

- **Nominal scales:** a collection of categories into which experimental units can be placed. Categories should be mutually exclusive but there can be more than one category. Examples include species or sex.

Self-test questions

The more you think about experimental design, the easier designing robust experiments becomes. To get you thinking while you are reading this book, we now include a number of self-test questions in every chapter. Often there will not be a clear right or wrong answer to a question, but suggested answers to all questions can be found at the back of the book.

Q 3.1 Are we safe in restricting our sample to married couples?

We found the same pattern in 98 out of 100 couples we would expect to find. What we have done is **replicate** our observation. If difference between replicate measurements. The more times we make the more likely it is that we are observing a real pattern. All statistical tests are really just a way of formalizing the idea that the phenomenon the less likely it is to be occurring simply by chance.

Take-home messages

To help to consolidate your thinking, we end most sections with a take-home message. By including these throughout the text we hope to give you an immediate opportunity to evaluate your understanding of a section before moving on.

To get more power from your subjects, increasing the number of subjects will generally increase the number of measurements (trees in this case).

Bias your sampling by getting more measurements at the level where you expect the greatest effect.

Using unbalanced groups for ethical reasons

BOX 6.1 Some design considerations for human biomedical studies

Imagine researchers are studying whether the incidence of schizophrenia differs between ethnic groups. One method might be to select 20 towns and for each use census data to estimate the incidence of schizophrenia. Say it is found that those towns with a higher proportion of white ethnicity also had a higher per-capita rate to estimate the incidence of schizophrenia. This is an inappropriate conclusion to draw because the incidence of schizophrenia is known to be higher in towns with a higher proportion of white ethnicity.

Boxes

We have aimed to produce a book that can be read from cover to cover, and so have tried to keep unnecessary details out of the main text. However, in areas where we feel that more details or more examples will lead to a fuller understanding of a concept we have included supplementary boxes.


Chapter outlines

Every chapter begins with an outline of the main points covered by that chapter. These outlines should help you to prepare for what is ahead. The outlines will also provide an easy way for you to dip in and out of the book during subsequent readings.

Chapter summaries

Every chapter finishes with a summary of the most important points covered in that chapter. By reading through this list you should be able to reassure yourself that you have got everything out of the chapter that you can, or go back and read again sections that are not yet clear.

manipulate clutch size, by removing or adding eggs to zebra finches, to monitor the survival of the mother after the chicks have fledged. We could potentially carry out such an experiment in the field. How should we choose? Let's begin with the lab. Maybe the question to ask is whether your study organism will be comfortable in the lab. We are unlikely to have a problem as zebra finches will breed readily in the lab. We have been looking at the same question in a bird that is more challenging to keep in captivity (like an albatross), we might have had problems. The suitability of the lab is hugely variable, and must be given careful thought. If an animal is not well disposed to captivity, then there are studies is



Ethical issues

We cannot overemphasize the importance of ethical issues when designing studies involving living organisms (including humans) or material derived from living organisms. In this new edition we highlight sections of particular relevance to ethics by placing a symbol in the margin at the appropriate point.

Flow chart

The exact process of designing an experiment will vary considerably between studies. Nevertheless, there are key stages in the design process that will apply to most, if not all, studies. The flow chart at the end of this book is intended to summarize and guide you through the main stages of designing an experiment. We have indicated at each point in the chart the sections of the book that are most relevant.

Acknowledgements from the first edition

Graeme has been moaning at the representatives of scientific publishers for years that there was not a book like this available. However, it was Sarah Shannon of Blackwell's that suggested that he stop moaning and write one himself. It was also Sarah who suggested that Nick would be an ideal co-author. We may have ended up at another publisher, but we are both very grateful to Sarah for her support and advice at the inception of this book.

At OUP, Jonathan Crowe has been everything we could have hoped for in an editor. He commented on two drafts, and found just the right combination of criticism and enthusiasm to spur us on. He also found excellent reviewers (Roland Hall, Keith McGuiness, and two others), who, in their own diverse ways, have also really helped to improve the finished work. Also at OUP, John Grandidge and Emily Cooke were models of friendly efficiency. Many students from the 2001–2002 BSc Zoology and Aquatic Bioscience courses at Glasgow, and the Zoo4 QZ course at Edinburgh provided useful feedback. In particular, Sandie King provided particularly full and useful comments. Fiona McPhie read the whole thing and gave honest feedback on some bad jokes that we've thankfully removed.

The text was carefully proof-read by Hazel Ruxton and the late David Ruxton, who managed to remove embarrassing numbers of errors.

Many people have influenced our thoughts on experimental design over the years, and as a result some of the examples will have been borrowed in whole or part. While it is impossible to trace and thank the originators of all these ideas there are some people we should certainly mention. Graeme learnt a huge amount about designing experiments from Will Cresswell and Rik Smith. Kate Lessells and Andrew Read have both, over the years, forced Nick to think far harder about experiments than he ever really intended; they will both find large elements of themselves throughout the book. Anyone else that recognizes, and wishes to claim as their own, examples contained in this book is welcome to contact us. If the publishers ever let us write a second edition, we will make sure you get the proper credit.

The figures were drawn by the imaginative and efficient Liz Denton. It's been an absolute pleasure working with her. Stuart Humphries kindly provided the drawing of a bird for Figure 2.1. Thanks also to Ruedi Nager, Stuart Humphries, and James Brown of the Scottish Agricultural Science Agency for cover illustrations.

Lastly we'd like to thank each other. Co-writing a book is much easier than working alone. We both found that the other person can really pick you up whenever inspiration disappears and writers' block is looming. But best of all, no matter what mistakes or shortcomings you find in this book, we'll both be able to blame it on the other guy!

Acknowledgements from the third edition

Thanks again to Liz Denton for the figures. At OUP, Dewi Jackson has helped to guide us through the process of putting the new edition together. Thanks to all the students and colleagues who have offered feedback on previous editions. Thanks to Hazel and Becky for still being with us through all the editions. Thanks also to three new additions since we first wrote the book—Katherine, Amelie, and Isla—for enriching our lives beyond measure, and providing happy distraction from these revisions.

Contents

1 Why you should care about design	1
1.1 Why experiments need to be designed	1
1.2 The costs of poor design	2
1.2.1 Time and money	2
1.2.2 Ethical issues	3
1.3 The relationship between experimental design and statistics	4
1.4 Why good experimental design is particularly important to life scientists	5
1.4.1 Random variation	5
1.4.2 Confounding factors	6
Summary	7

2 Starting with a well-defined hypothesis	8
2.1 Why your experiment should be focused: questions, hypotheses, and predictions	8
2.1.1 An example of moving from a question to hypotheses, and then to an experimental design	10
2.1.2 An example of multiple hypotheses	11
2.1.3 Where do (good) ideas come from in the first place?	13
2.2 Producing the strongest evidence with which to challenge a hypothesis	14
2.2.1 Indirect measures	15
2.2.2 Considering all possible outcomes of an experiment	15
2.3 Satisfying sceptics	17
2.4 The importance of a pilot study and preliminary data	17
2.4.1 Making sure that you are asking a sensible question	18
2.4.2 Making sure that your techniques work	19

2.5	Experimental manipulation versus natural variation	21
2.5.1	An example hypothesis that could be tackled by either manipulation or correlation	21
2.5.2	Arguments for doing a correlational study	23
2.5.3	Arguments for doing a manipulative study	24
2.5.4	Situations where manipulation is impossible	26
2.6	Deciding whether to work in the field or the laboratory	29
2.7	<i>In vivo</i> versus <i>in vitro</i> studies	31
2.8	There is no perfect study	31
	Summary	32
3	Between-individual variation, replication, and sampling	34
3.1	Between-individual variation	34
3.2	Replication	35
3.3	Pseudoreplication	45
3.3.1	Explaining what pseudoreplication is	45
3.3.2	Common sources of pseudoreplication	48
3.3.3	Dealing with pseudoreplication	51
3.3.4	Accepting that sometimes pseudoreplication is unavoidable	53
3.3.5	Pseudoreplication, third variables, and confounding variables	54
3.3.6	Cohort effects, confounding variables, and cross-sectional studies	54
3.4	Randomization	55
3.4.1	Why you often want a random sample	56
3.4.2	Haphazard sampling	56
3.4.3	Self-selection	57
3.4.4	Some pitfalls associated with randomization procedures	58
3.4.5	Randomizing the order in which you treat replicates	59
3.4.6	Random samples and representative samples	59
3.5	Selecting the appropriate number of replicates	60
3.5.1	Educated guesswork	61
3.5.2	Formal power analysis	61

3.5.3	Factors affecting the power of an experiment	61
3.5.4	Relation between power and type I and type II errors	63
<i>Summary</i>		74

4 Different experimental designs **75**

4.1	Controls	75
4.1.1	Different types of control	76
4.1.2	Blind procedures	77
4.1.3	Making sure the control is as effective as possible	79
4.1.4	The ethics of controlling	80
4.1.5	Situations where a control is not required	80
4.2	Completely randomized and factorial experiments	81
4.2.1	Experiments with several factors	82
4.2.2	Confusing levels and factors	85
4.2.3	Advantages and disadvantages of complete randomization	85
4.3	Blocking	88
4.3.1	Blocking on individual characters, space, and time	89
4.3.2	The advantages and disadvantages of blocking	90
4.3.3	Paired designs	91
4.3.4	How to select blocks	91
4.3.5	Covariates	92
4.4	Within-subject designs	93
4.4.1	The advantages of a within-subject design	93
4.4.2	The disadvantages of a within-subject design	94
4.4.3	Isn't repeatedly measuring the same individual pseudoreplication?	95
4.4.4	With multiple treatments, within-subject experiments can take a long time	96
4.4.5	Which sequences should you use?	97
4.5	Split-plot designs (sometimes called split-unit designs)	98
4.6	Thinking about the statistics	100
<i>Summary</i>		102

5 Taking measurements	104
5.1 Calibration	105
5.2 Inaccuracy and imprecision	106
5.3 Intra-observer variability	108
5.3.1 Describing the problem	108
5.3.2 Tackling the problem	109
5.3.3 Repeatability	109
5.3.4 Remember, you can be consistent but still consistently wrong	110
5.4 Inter-observer variability	112
5.4.1 Describing the problem	112
5.4.2 Tackling the problem	112
5.5 Deciding how to measure	113
5.5.1 Defining categories	113
5.5.2 How fine-scale to make interval/ratio measurements	113
5.6 Observer effects	114
5.7 Recording data	115
5.7.1 Don't try to record too much information at once	115
5.7.2 Beware of shorthand codes	116
5.7.3 Keep more than one copy of your data	116
5.7.4 Write out your experimental protocol formally and in detail, and keep a detailed field journal or lab book	116
5.7.5 Don't over-work	117
5.8 Computers and automated data collection	117
5.9 Floor and ceiling effects	118
5.10 Observer bias	119
5.11 Taking measurements of humans and animals in the laboratory	119
5.12 Sensitivity and specificity	120
Summary	123

6 Final thoughts	124
6.1 How to select the levels for a treatment	125
6.2 Subsampling: more woods or more trees?	127
6.3 Using unbalanced groups for ethical reasons	129
6.4 Other sampling schemes	132
6.4.1 Sequential sampling	132
6.4.2 Stratified sampling	132
6.4.3 Systematic sampling	134
6.5 Latin square designs	134
6.6 More on interactions	136
6.6.1 Covariates can interact too	136
6.6.2 The importance of interactions (and the interaction fallacy)	138
6.6.3 Simpson's paradox	141
6.7 Dealing with human subjects	142
6.7.1 Deception	144
6.7.2 Collecting data without permission	144
6.7.3 Confidentiality	145
6.7.4 Discretion	145
6.7.5 Randomized response	146
6.7.6 Ethical guidelines	147
6.7.7 Volunteers	147
6.7.8 Honesty of subjects	147
6.7.9 There is no perfect study: a reprise	148
6.8 Special terminology of experimental design used in the biomedical literature	148
<i>Summary</i>	152
<i>Sample answers to self-test questions</i>	155
<i>Flow chart on experimental design</i>	167
<i>Bibliography</i>	172
<i>Index</i>	177



Why you should care about design

1

1.1 Why experiments need to be designed

When life sciences students see the phrase 'experimental design', it can either ease them gently into a deep sleep or cause them to run away screaming. For many, experimental design conjures up unhappy memories of mathematics or statistics lessons, and is generally thought of as something difficult that should be left to statisticians. Wrong on both counts! Designing simple but good experiments doesn't require difficult maths. Instead, experimental design is more about common sense, biological insight, and careful planning. Having said that, it does require a certain type of common sense, and there are some basic rules. In this book, we hope to steer you towards thinking more effectively about designing experiments.



Designing effective experiments needs thinking about biology more than it does mathematical calculations.

So why are many life scientists so averse to thinking about design? Part of the reason is probably that it is easy to think that time spent designing experiments would be better spent actually doing experiments. After all, the argument goes, we are biologists so let's concentrate on the biology and leave the statisticians to worry about the design and analysis. This attitude has given rise to a number of myths that you can hear from the greenest student or the dustiest professor.

Myth 1: It does not matter how you collect your data, there will always be a statistical 'fix' that will allow you to analyse it.

It would be wonderful if this was true, but it is not. There are a large number of statistical tests out there, and this can lead to the false impression that there must be one for every situation. However, all statistical tests make assumptions about your data that must be met before the test can be meaningfully applied. Some of these assumptions are very specific to the particular test. If you cannot meet these, there may be a substitute test that assumes different characteristics of your data. But you may find that this alternative test only allows you to use your data to answer a different scientific question