Analysis of Variance and Covariance: How to Choose and Construct Models for the Life Sciences

Analysis of variance (ANOVA) is a core technique for analysing data in the life sciences. This reference book bridges the gap between statistical theory and practical data analysis by presenting a comprehensive set of tables for all standard models of analysis of variance and covariance with up to three treatment factors. The book will serve as a tool to help post-graduates and professionals define their hypotheses, design appropriate experiments, translate them into a statistical model, validate the output from statistics packages and verify results. The systematic layout makes it easy for readers to identify which types of model best fit the themes they are investigating, and to evaluate the strengths and weaknesses of alternative experimental designs. In addition, a concise introduction to the principles of analysis of variance and covariance is provided, alongside worked examples illustrating issues and decisions faced by analysts.

PATRICK DONCASTER is a Reader in Ecology in the School of Biological Sciences at the University of Southampton.

ANDREW DAVEY is an Environmental Scientist at the UK Water Research centre (WRc).

Analysis of Variance and Covariance

How to Choose and Construct Models for the Life Sciences

C. PATRICK DONCASTER and ANDREW J. H. DAVEY

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521684477

First published 2007

Printed in the United Kingdom at the University Press, Cambridge

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

ISBN 978-0-521-86562-3 hardback
ISBN 978-0-521-68447-7 paperback

Contents

Preface

Hypothesis testing in the life sciences often involves comparing samples of observations, and analysis of variance is a core technique for analysing such information. Parametric analysis of variance, abbreviated as 'ANOVA', encompasses a generic methodology for identifying sources of variation in continuous data, from the simplest test of trend in a single sample, or difference between two samples, to complex tests of multiple interacting effects. Whilst simple one-factor models may suffice for closely controlled experiments, the inherent complexities of the natural world mean that rigorous tests of causality often require more sophisticated multi-factor models. In many cases, the same hypothesis can be tested using several different experimental designs, and alternative designs must be evaluated to select a robust and efficient model. Textbooks on statistics are available to explain the principles of ANOVA and statistics packages will compute the analyses. The purpose of this book is to bridge between the texts and the packages by presenting a comprehensive selection of ANOVA models, emphasising the strengths and weaknesses of each and allowing readers to compare between alternatives.

Our motivation for writing the book comes from a desire for a more systematic comparison than is available in textbooks, and a more considered framework for constructing tests than is possible with generic software. The obvious utility of computer packages for automating otherwise cumbersome analyses has a downside in their uncritical production of results. Packages adopt default options until instructed otherwise, which will not suit all types of data. Numerous problems can arise from incautious use of any statistics package, be it of the simplest or

the most sophisticated type. In this book we will anticipate all of the following common issues:

- Wrong model or insufficient terms requested for the desired hypothesis (page 2 onwards);
- Wrong error terms calculated by default or wrongly requested (page 2 onwards);
- Data unsuitable for analysis of variance (page 14);
- Unwise pooling of error terms by default or design (page 38);
- Default analysis of effects that have no logical test (e.g., several designs in Chapter 7);
- In unbalanced designs, inappropriate default adjustment to variance estimates (page 237);
- In mixed models, undesired default use of an unrestricted model (page 242);
- Inappropriate application of analysis of variance (page 250).

Armed with precise knowledge of the structure of a desired analysis, the user can evaluate outputs from a statistics package and correct inconsistencies or finish the analysis by hand. The main chapters of this book are designed to provide the relevant information in a clearly accessible format. They are preceded by an introduction to analysis of variance that provides the context of experimental design, and followed by further topics that treat issues arising out of design choices.

Scope and approach

Whilst there is no computational limit to the complexity of ANOVA models, in practice, designs with more than three treatment factors are complicated to analyse and difficult to interpret. We therefore describe all common models with up to three treatment factors for seven principal classes of ANOVA design:

1 One-factor – replicate measures at each level of a single explanatory factor;
2 Nested – one factor nested in one or more other factors;
3 Factorial – fully replicated measures on two or more crossed factors;
4 Randomised blocks – repeated measures on spatial or temporal groups of sampling units;
5 Split plot – treatments applied at multiple spatial or temporal scales;
6 Repeated measures – subjects repeatedly measured or tested over time;
7 Unreplicated factorial – a single measure per combination of two or more factors.

For each class of ANOVA, we describe its particular applications, highlight its strengths and weaknesses, and draw comparisons with other classes. We then present a series of models covering all reasonable combinations of fixed and random factors. For each model we provide the following information:

- The model equation and its associated test hypothesis;
- A table illustrating the allocation of factor levels to sampling units;
- Illustrative examples of applications;
- Any special assumptions and guidance on analysis and interpretation;
- Full analysis of variance tables.

A systematic approach, with consistent layout and notation, makes it easy for readers to evaluate alternative models and to identify which type of model best fits the themes they are investigating.

Examples bring statistics to life as they show how particular models can be applied to answer real-life questions. Throughout the book we develop a series of examples to illustrate the similarities and differences between different ANOVA models. More detailed worked examples are also given to illustrate how the choice of model follows logically from the design of the experiment and determines the inferences that can be drawn from the results.

A multitude of statistics packages are available on the market and it is beyond the scope of this book to describe the analysis of ANOVA models in each. Rather, we encourage readers to become familiar with the approach taken by their favourite package, and to interpret its outputs with the help of the tables in the book and the sample datasets on our website.

How to use this book

The book is a reference tool to help experimental and field biologists define their hypotheses, design an appropriate experiment or mensurative study, translate it into a statistical model, analyse the data and validate the resulting output. As such, it is intended to be a companion throughout the scientific process. At the planning stage, the documented tables allow users to make informed choices about the design of experiments or fieldwork, with particular regard to the need for replication and the different scales of replication across space or over time. Different designs are directly comparable, facilitating the task of balancing costs of replication against benefits of predictive power and generality. At the analysis stage, the book shows how to construct ANOVA models with

the correct *F*-ratios for testing the hypotheses, gives options for *post hoc* pooling of error terms, and highlights the assumptions underlying the predictions. Finally, by appreciating the methods used by computer packages to perform ANOVAs, users can check that their input model is appropriately structured and correctly formatted for the desired hypothesis, can verify that the output has tested the intended hypothesis with the correct error degrees of freedom, and can draw appropriate conclusions from the results.

Who should use this book?

The book is aimed at researchers of post-graduate level and above who are planning experiments or fieldwork in the life sciences and preparing to ask questions of their data. We assume that readers are familiar with the basic concepts of statistics covered by introductory textbooks (e.g., Dytham 2003; Ruxton and Colegrave 2003; McKillup 2006, amongst many). Numerous very readable texts already exist to explain the theory and mechanics underpinning analysis of variance (e.g., Kirk 1994; Underwood 1997; Crawley 2002; Grafen and Hails 2002; Quinn and Keough 2002), and we recommend that readers consult such texts in addition to this book. We expect the users of this book to analyse their data with a statistics package suitable for analysis of variance, and we assume that they will employ its tutorial and help routines sufficiently to understand its input commands and output tables.

Companion website

The book is supported by a website at www.soton.ac.uk/~cpd/anovas, which provides additional tools to help readers analyse and interpret the ANOVA models presented here. The website includes:

- *Analyses of example datasets*. The analyses illustrate how the raw data translate into tested hypotheses for each of the ANOVA models in this book. Datasets can be freely downloaded to verify the output from the reader's own statistics package.
- *Model selection and comparison tools*. A dichotomous key to the main classes of ANOVA model is provided to help readers select the right kind of ANOVA design for their needs, and a hyperlinked summary of all the ANOVA models in the book is presented to facilitate the comparison of alternative models.

Acknowledgements

We owe a great debt of thanks to the numerous people who gave us their time to discuss statistical concepts, make suggestions for improving structure and content, and correct our errors. Although we have not succeeded in following all of their recommendations, those that we have implemented have improved the book beyond recognition. We are especially grateful to Henrique Cabral, Stefano Cannicci, Tasman Crowe, Alan Grafen, Jeremy Greenwood, Laura Guichón, Stephen Hawkins, Jan Geert Hiddink, Mark Huxham, Fernando Milesi, Jonathan Newman, Susanne Plesner Jensen, Will Resetarits, Graeme Ruxton, Martin Skov, Richard Thompson, the undergraduates and postgraduates in biology and environmental sciences at Southampton University. Finally, AD would like to thank Kirsty Davey for her unfailing support during the writing of this book.

While researching this book, we have been inspired and guided by a number of standard-setting accounts of statistical methods, notably in books by Robert Sokal and F. James Rohlf, Tony Underwood, Gerry Quinn and Michael Keough, Alan Grafen and Rosie Hails, and Michael Crawley. We thank these authors, along with all those referenced in the book, for the tuition that we have gained from their works.