

palgrave
macmillan

Serious Stats

A Guide to
Advanced
Statistics
for the
Behavioral
Sciences

Thorn
Baguley



SERIOUS STATS

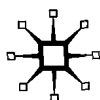
A GUIDE TO ADVANCED STATISTICS FOR THE
BEHAVIORAL SCIENCES

THOM BAGULEY

Professor of Experimental Psychology, Nottingham Trent University, UK



palgrave
macmillan



© Thomas Simon Baguley 2012

All rights reserved. No reproduction, copy or transmission of this publication may be made without written permission.

No portion of this publication may be reproduced, copied or transmitted save with written permission or in accordance with the provisions of the Copyright, Designs and Patents Act 1988, or under the terms of any licence permitting limited copying issued by the Copyright Licensing Agency, Saffron House, 6–10 Kirby Street, London EC1N 8TS.

Any person who does any unauthorized act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

The author has asserted his right to be identified as the author of this work in accordance with the Copyright, Designs and Patents Act 1988.

First published 2012 by
PALGRAVE MACMILLAN

Palgrave Macmillan in the UK is an imprint of Macmillan Publishers Limited, registered in England, company number 785998, of Houndmills, Basingstoke, Hampshire RG21 6XS.

Palgrave Macmillan in the US is a division of St Martin's Press LLC, 175 Fifth Avenue, New York, NY 10010.

Palgrave Macmillan is the global academic imprint of the above companies and has companies and representatives throughout the world.

Palgrave® and Macmillan® are registered trademarks in the United States, the United Kingdom, Europe and other countries.

ISBN 978–0–230–57717–6 hardback
ISBN 978–0–230–57718–3 paperback

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources. Logging, pulping and manufacturing processes are expected to conform to the environmental regulations of the country of origin.

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

A catalog record for this book is available from the Library of Congress.

10 9 8 7 6 5 4 3 2 1
21 20 19 18 17 16 15 14 13 12

Printed in China

SERIOUS STATS

Preface

About this book

This book is a bridging text for students and researchers in the human and behavioral sciences. The ideal reader will have some familiarity with inferential statistics – perhaps as part of an undergraduate degree in a discipline such as psychology, cognitive science or ergonomics – and be interested in deepening their understanding or learning new material. This book aims to bridge the gap between a reader's existing understanding of statistics and that required to apply and interpret more advanced statistical procedures.

I have also tried to make the book a helpful resource for experienced researchers who wish to refresh their statistical knowledge or who have good understanding of a 'narrow' but fairly advanced topic such as analysis of variance. I hope it will allow these readers to expand from islands of existing expertise to new territory.

The book starts with a review of basic inferential statistics, beginning with descriptive statistics, probability distributions and statistical inference (in the form of confidence intervals and significance tests). If you are already familiar with these topics I would encourage you to look through these chapters to refresh your understanding. In addition, this material may be presented in a slightly different way (e.g., from a different perspective or in greater depth).

Later chapters introduce core topics such as, regression, correlation and covariance, effect size, and statistical power. Unless you have advanced training in statistics it is likely that you will benefit from looking closely at this material – it is fundamental to an appreciation of later content. Two further chapters consider the messiness inherent in working with real data (particularly data from human participants). The approach I adopt is to give a taster of some methods for exploring and dealing with messy data, rather than provide a comprehensive recipe for checking and solving every possible problem. This is both for practical reasons (as each of these chapters could be a book in its own right) and because the best approach in any particular situation depends on what you are trying to do and the context from which the data are drawn.

Later chapters cover what I consider to be advanced material: multiple regression, analysis of variance, analysis of covariance, and the general linear model. Before covering these topics I review alternatives to classical, frequentist inference (and significance tests in particular). In order to get the most out of the more advanced material in the book, you will need to understand the problems inherent in relying (solely) on a p value from a significance test for inference. I also think it important to go beyond criticism of the p value approach and present viable alternatives. Three are presented here: Bayesian, likelihood, and information theoretic approaches to inference. There are important connections (and distinctions) between these three approaches. In this chapter, I sacrifice depth for breadth (though there is sufficient material to run a range of analyses using each approach).

The final chapters explore the most challenging topics. Also included are chapters on interaction effects and contrasts. These topics are extremely important for researchers in the human and behavioral sciences, but are often covered only briefly (if at all) in introductory classes. My goal here is to remedy this deficit. The final two chapters introduce generalized linear models

(for discrete outcomes) and multilevel models (with emphasis on repeated measures models). I have tried to emphasize the links between these advanced topics and the general linear model and to demonstrate what they offer over and above simpler models.

If you want to learn and understand what is covered, it is essential that you have a go at applying it. Each chapter contains worked examples. Many of these use real data sets. These are necessarily a bit messy and don't always lead to clear answers (and several data sets are chosen because they have interesting or unusual quirks). My aim is to illustrate some of the challenges of working with real data sets (and the importance of data exploration and model checking). In other cases I have resorted to creating artificial data sets to illustrate a particular point, or to make it easier to conduct calculations by hand. These data sets are carefully constructed to meet the requirements of the example – though you will sometimes encounter real data sets with similar properties. In general, the early examples use hand calculation while later examples require you to use a computer.

Hand calculation can sometimes help you to understand how an equation works or demystify a (supposedly) complex technique. This will depend on your confidence and ability with basic mathematical operations. This doesn't work with every procedure, and in many examples I explain how to use a computer package to provide intermediate values that, when put together in the right way, illustrate what is going on. From time to time the mathematics is sufficiently challenging that I merely describe the gist of what is happening (and rely on the computer to provide a complete solution). Where necessary, I refer interested readers to a more detailed mathematical account of what is taking place.

The contents of the book differ from the coverage of a typical introductory or intermediate statistics course in the behavioral or human sciences. One difference is the breadth of coverage, which runs from descriptive statistics to generalized linear and multilevel models. Another is the reduced emphasis on null hypothesis significance tests and increased emphasis on confidence intervals or other inferential tools. Several topics have more prominence than you might expect: graphical methods; effect size; contrasts and interactions. Other topics have less emphasis or are presented differently: psychometrics, multivariate analysis of variance, non-parametric statistics; and pairwise comparisons. I have chosen to focus on univariate methods – methods for the analysis of a single outcome measure (though there may be many predictor variables). Covering multivariate statistics (in the sense of modeling multiple outcomes) and psychometrics would probably have doubled the page count. Nonparametric statistics are covered, but in an atypical way. Several methods, often considered to be 'nonparametric' (e.g., bootstrapping, kernel density estimation, the rank transformation and robust regression), are integrated into the text at appropriate points. I have, in particular, avoided describing a large number of rank transformation tests in detail. My preference is to emphasize the link between parametric and so-called non-parametric approaches and to encourage consideration of robust methods as alternatives to the usual (e.g., least squares) models.

If there is a single message to take away from this book, it is that statistical modeling is not a set of recipes or instructions. It is the search for a model or set of models that capture the regularities and uncertainties in data, and help us to understand what is going on.

Software

Many of the statistical tools described in the book require specialist software to run them. Nearly all the examples were implemented in the free, open source statistical programming

environment R (R Core Development Team, 2011). This will run on PC, Mac and Linux operating systems. Installation on a Mac or PC is generally very easy. For details on downloading and installing R see: <http://cran.r-project.org/>

R works slightly differently on PC, Mac and Linux machines and for this reason I have, for the most part, avoided referring to platform-specific features of R (e.g., resizing windows, printing, opening or saving files). There are many online guides and books that run through the basics of installing and running R (in addition to the information available when you download R for the first time).

At the end of each chapter, I provide a detailed description of the R code used to reproduce the examples in that chapter. Data sets and R scripts for each chapter are available with the online resources for this book. I assume no previous knowledge of R (and only limited familiarity with statistical computing). Although I am not trying to teach R *per se*, I have tried to include enough explanation of how R works for readers, if they wish, to learn it as they go along. To this end, the complexity of the R code being used increases gradually from Chapter 1 through to Chapter 18. In some cases I have glossed over fine grain technical details about how R works (e.g., the difference between ‘modes’ of vector) and used generic terms such as ‘string’ alongside R-specific terms such as ‘data frame’.

As well as R code, I provide very brief notes on relevant SPSS syntax at the end of each chapter (where relevant). SPSS is the most widely used statistics package in the human and behavioral sciences (in the UK at least). These notes are included for two reasons: (1) to reveal some of the hidden features and capabilities of SPSS, and (2) to highlight the advantages of using R alongside or in place of SPSS.

Mathematics

To get the most from this book you will need to have basic mathematical competence. I assume readers will have mastered basic arithmetic (e.g., addition, subtraction, division and multiplication) and be familiar with concepts such as, fractions, decimals, rounding, negative numbers, squares and square roots, and cubes and cubed roots. Knowing the order in which to apply arithmetic operations (e.g., PEMDAS or BODMAS) is also necessary. You should also have some understanding of percentages, probabilities and ratios (and perhaps reciprocals, exponents, logarithms and factorials) and simple algebra. If you are rusty on any of these topics don’t worry too much – as there will be some ‘refresher’ material in each chapter.

I would expect readers to be able to answer the following arithmetic problems without much difficulty:

$$\begin{array}{lll} 4 + 3 \times 5 = ? & 6(3 - 1) = ? & \sqrt{9} = ? \\ 10^2 = ? & \frac{9}{2} = ? & 0.5 \times 0.5 = ? \end{array}$$

I would also expect readers to understand what the following equations mean, and (perhaps with a little help) be able to solve them:

$$\begin{array}{lll} 3x + 1 = 10 \Rightarrow x = ? & 4! = ? & \sqrt{2} \times \sqrt{2} = ? \\ P(A) + P(\sim A) = ? & \sqrt[3]{27} = ? & -2 \times -3 = ? \end{array}$$

In addition to mathematical competence, I anticipate some familiarity with data collection, exploration and analysis. This may include experimental design, descriptive statistics and simple graphical methods such as line graphs or scatter plots (and technical terms such as *x*-axis and *y*-axis). Again, the text contains refresher material on most of these topics.

Many of the examples refer to 'hand calculation'. This is a fuzzy concept, but I take it to mean reproducing the calculations step-by-step as if doing them by hand. It therefore includes using paper and pencil, mental arithmetic, pocket calculators or spreadsheets (if used in the correct way). In fact, a spreadsheet is one of the best ways to organize hand calculations to understand what they are doing (provided you know how to use one and are careful to set out all the intermediate steps).

Boxed sections and online supplements

As well as learning features (such as examples, sections on R code, and SPSS notes), there are two types of boxed section used throughout the book. One type covers key concepts or important ideas that are referred to in several different chapters. These are referred to as 'key concepts' and numbered by chapter and serial position within chapter (e.g., Key Concept 2.1 is the first key concept box in Chapter 2). The other is a more traditional boxed section that is used to improve the flow of the text and contains material that is relatively self-contained (and that is generally referred to again only within that chapter). These are referred to as 'boxes' and numbered by chapter and serial position within chapter (e.g., Box 1.2 is the second boxed section in Chapter 1).

In addition to the boxed sections there are five online supplements. Supplements 1, 2 and 5 cover advanced topics that are not central to the text, but will be very useful for some readers (meta-analysis, dealing with missing data and loglinear models). Supplements 3 and 4 provide more detail on peripheral topics that are mentioned in the main text (replication probabilities and pseudo- R^2 measures).

Acknowledgments

I would like to thank the following people who have contributed suggestions for the book, read chapters or provided encouragement: Jaime Marshall; Andrew Dunn; Danny Kaye; Mark Andrews; Stanislav Kolenikov; Dave Atkins; Tim Wells; Tom Dunn; Mark Torrance; James Houston and Zoltan Dienes.

THOM BAGULEY

Contents overview

<i>List of tables</i>	xii
<i>List of figures</i>	xiv
<i>List of boxes</i>	xviii
<i>List of key concepts</i>	xix
<i>Preface</i>	xx
1 Data, samples and statistics	1
2 Probability distributions	37
3 Confidence intervals	74
4 Significance tests	118
5 Regression	158
6 Correlation and covariance	205
7 Effect size	234
8 Statistical power	277
9 Exploring messy data	304
10 Dealing with messy data	339
11 Alternatives to classical statistical inference	363
12 Multiple regression and the general linear model	423
13 ANOVA and ANCOVA with independent measures	472
14 Interactions	527
15 Contrasts	590
16 Repeated measures ANOVA	622
17 Modeling discrete outcomes	667
18 Multilevel models	724
<i>Notes</i>	785
<i>References</i>	798
<i>Author index</i>	817
<i>Subject index</i>	821

Contents

<i>List of tables</i>	xii
<i>List of figures</i>	xiv
<i>List of boxes</i>	xviii
<i>List of key concepts</i>	xix
<i>Preface</i>	xx
1 Data, samples and statistics	1
1.1 Chapter overview	2
1.2 What are data?	2
1.3 Samples and populations	3
1.4 Central tendency	6
1.5 Dispersion within a sample	17
1.6 Description, inference and bias	25
1.7 R code for Chapter 1	27
1.8 Notes on SPSS syntax for Chapter 1	34
1.9 Bibliography and further reading	36
2 Probability distributions	37
2.1 Chapter overview	38
2.2 Why are probability distributions important in statistics?	38
2.3 Discrete distributions	42
2.4 Continuous distributions	48
2.5 R code for Chapter 2	65
2.6 Notes on SPSS syntax for Chapter 2	72
2.7 Bibliography and further reading	73
3 Confidence intervals	74
3.1 Chapter overview	75
3.2 From point estimates to interval estimates	75
3.3 Confidence intervals	76
3.4 Confidence intervals for a difference	86
3.5 Using Monte Carlo methods to estimate confidence intervals	93
3.6 Graphing confidence intervals	100
3.7 R code for Chapter 3	103
3.8 Notes on SPSS syntax for Chapter 3	115
3.9 Bibliography and further reading	117
4 Significance tests	118
4.1 Chapter overview	119
4.2 From confidence intervals to significance tests	119
4.3 Null hypothesis significance tests	120
4.4 <i>t</i> tests	125
4.5 Tests for discrete data	130

13.6	ANOVA, ANCOVA and multiple regression	511
13.7	R code for Chapter 13	511
13.8	Notes on SPSS syntax for Chapter 13	522
13.9	Bibliography and further reading	526
14	Interactions	527
14.1	Chapter overview	528
14.2	Modeling interaction effects	528
14.3	Interactions in regression: moderated multiple regression	529
14.4	Polynomial regression	540
14.5	Factorial ANOVA	542
14.6	ANCOVA and homogeneity of covariance	562
14.7	Effect size in factorial ANOVA, ANCOVA and multiple regression	564
14.8	Statistical power to detect interactions	568
14.9	Problems with interactions in ANOVA and regression	569
14.10	R code for Chapter 14	571
14.11	Notes on SPSS syntax for Chapter 14	586
14.12	Bibliography and further reading	589
15	Contrasts	590
15.1	Chapter overview	591
15.2	Contrasts and the design matrix	591
15.3	Interaction contrasts	605
15.4	Post hoc contrasts and correction for multiple testing	609
15.5	Contrasts of adjusted means in ANCOVA	611
15.6	The role of contrasts in other statistical models	612
15.7	R code for Chapter 15	613
15.8	Notes on SPSS syntax for Chapter 15	620
15.9	Bibliography and further reading	621
16	Repeated measures ANOVA	622
16.1	Chapter overview	623
16.2	Modeling correlated or repeated measures	623
16.3	ANOVA with repeated measures	623
16.4	Combining independent and repeated measures: mixed ANOVA designs	638
16.5	Comparisons, contrasts and simple effects with repeated measures	642
16.6	MANOVA	647
16.7	ANCOVA with repeated measures	650
16.8	R code for Chapter 16	656
16.9	Notes on SPSS syntax for Chapter 16	664
16.10	Bibliography and further reading	666
17	Modeling discrete outcomes	667
17.1	Chapter overview	668
17.2	Modeling discrete outcomes in the general linear model	668
17.3	Generalized linear models	669
17.4	Logistic regression	672
17.5	Modeling count data	694
17.6	Modeling discrete outcomes with correlated measures	706

17.7	R code for Chapter 17	708
17.8	Notes on SPSS syntax for Chapter 17	720
17.9	Bibliography and further reading	722
17.10	Online supplement 4: Pseudo- R^2 and related measures	723
17.11	Online supplement 5: Loglinear models	723
18	Multilevel models	724
18.1	Chapter overview	725
18.2	From repeated measures ANOVA to multilevel models	725
18.3	Multilevel regression models	731
18.4	Building up a multilevel model	741
18.5	Crossed versus nested random factors	762
18.6	Multilevel generalized linear models	766
18.7	R code for Chapter 18	769
18.8	Notes on SPSS syntax for Chapter 18	782
18.9	Bibliography and further reading	784
	<i>Notes</i>	785
	<i>References</i>	798
	<i>Author index</i>	817
	<i>Subject index</i>	821

List of tables

3.1	Approximate and exact 99% confidence intervals for binomial proportions	85
4.1	Correct responses by cue and the difference between cues for the TOT data	130
5.1	Hypothetical data from a simple location memory experiment	166
5.2	Intermediate quantities for calculating the slope and intercept for a simple linear regression on the location memory data	167
5.3	Predicted, residuals, squared predicted, residual and observed Y for the location memory data	172
6.1	Reproduction of location memory data (from Table 5.1) in order of increasing presentation time	223
7.1	Buchner and Mayr (2004) negative priming data	245
7.2	Effect of priming a city name on judged closeness of city	248
7.3	Sample size per group and summary statistics for three hypothetical studies with two independent group designs	251
7.4	Mortality before and after surgical checklist intervention is introduced	261
7.5	The largest possible effect for the observed marginal totals of the surgical checklist data	263
8.1	Approximate correction factors for σ as a function of n and required degree of certainty	298
10.1	Comparison of 95% CIs of a difference in means, trimmed mean or medians for the comprehension data using independent t and a range of robust methods	346
12.1	A small data set with two orthogonal predictors X_1 and X_2 and a continuous outcome Y	429
12.2a	Analysis of variance (ANOVA) table format for a multiple linear regression model	438
12.2b	Formulas for computing a multiple regression ANOVA table	439
12.3a	An incomplete ANOVA table for the multiple linear regression of data from Table 12.1	440
12.3b	A completed ANOVA table for the multiple regression of data in Table 12.1	440
12.4	Comparison of full and reduced models for intention to vote data	455
13.1	Dummy codes for a categorical variable with three categories	475
13.2	Effect codes for a categorical variable with three categories	475
13.3	Dummy and effect indicator codes and mean description quality scores by group for the diagram data	477
13.4a	ANOVA table for a one-way ANOVA with independent measures	480
13.4b	Formulas for a one-way ANOVA with independent measures and J levels	480
13.5	Raw description quality score by group for the diagram data	484
13.6	One-way ANOVA table for the diagram data	485
13.7	Performance of procedures for controlling familywise error for five <i>a priori</i> tests on the diagram data	494
13.8	Maximum number of true null hypotheses if $i - 1$ have been rejected	498
13.9	Adjusted p values for five <i>a priori</i> tests	499
13.10	Dayton's PCIC procedure using AIC_C applied to the diagram ANOVA data	502

13.11	The structure of the ANOVA table for a one-way ANCOVA	505
13.12	The ANOVA table for a one-way ANCOVA of the diagram data	509
14.1	Raw data and means for a 2×2 factorial ANOVA design	544
14.2a	Dummy coding for 2×3 factorial ANOVA	551
14.2b	Effect coding for 2×3 factorial ANOVA	551
14.3	The ANOVA table for a two-way independent measures design	552
14.4	The ANOVA table for a three-way independent measures design	553
14.5	Summary data for the prospective memory study	554
14.6	ANOVA table for the proactive memory data	555
15.1	Sets of linear, quadratic and (where possible) cubic integer contrast weights	602
16.1	One-way ANOVA with repeated measures	629
16.2	Two-way ANOVA with repeated measures	629
16.3	Table for one-way ANOVA on the pride data	634
16.4	Comparison of Loftus-Masson and Cousineau-Morey 95% CIs	637
16.5	Two-way ANOVA with mixed measures	639
16.6	Cell means by condition and emotion for the pride data	640
16.7	Two-way ANOVA with mixed measures for the pride data	641
17.1	Regression coefficients, standard errors, Wald statistics and odds ratios for ordered logistic regression of the traffic data	693

List of figures

1.1	Frequency of eye colour for a sample of ten people	9
1.2	The main features of a box plot	20
2.1	Probability mass function for the number of heads observed from 10 tosses of a fair coin	40
2.2	Cumulative distribution function for the number of heads observed from 10 tosses of a fair coin	40
2.3	Probability mass function for a Poisson distribution with $\lambda = 3.5$	46
2.4	Probability mass functions for Poisson distribution with different rates	47
2.5	Histograms for 100,000 simulated samples from binomial distributions	50
2.6	Probability density function for a normal distribution	52
2.7	Normal distributions with differing parameters	53
2.8	Probability density and cumulative distribution functions for the standard normal (z) distribution	54
2.9	Probability density function for the standard lognormal	56
2.10	An illustration of the asymmetry of the binomial distribution	58
2.11	Probability density function for a chi-square distribution with 1 df	59
2.12	Probability density functions for chi-square distributions	60
2.13	Probability density of t_1 and t_{29} relative to z	61
2.14	Probability density of F	63
3.1	The 5% most extreme values (tail probabilities) of the standard normal (z) distribution	79
3.2	The effective width of a 95% CI for a difference	88
3.3	A taxonomy of sampling plans for resampling statistics	94
3.4	Frequency of observed percentage accuracy scores for the comprehension data	96
3.5	Comparison of observed and a single bootstrap sample by condition for the comprehension data	97
3.6	Distribution of mean between-group differences for 9,999 bootstrap samples from the comprehension data	99
3.7	Mean percentage accuracy and 95% confidence intervals by group for the comprehension data	101
3.8	Mean percentage accuracy for the comprehension data	103
4.1	Areas representing (a) two-sided and (b) one-sided p values for a z test	124
4.2	Degree class awarded by gender in 2009 to students in England and Wales	138
5.1	The straight line described by the equation $Y = 2 + 0.5X$	161
5.2	A graphical method to calculate the slope of a straight line	161
5.3	Lines of best fit for simple linear regression	163
5.4	Scatter plots of the location memory data	167
5.5	Residuals depicted as the vertical distance between observed and predicted Y for a simple linear regression of the location memory data	169
5.6	How the regression line for the location memory data changes after adding a new observation (the black triangle) at the mean of X	175

5.7	How the regression line for the location memory data changes after adding a new observation (the black triangle) with high leverage and a large residual	175
5.8	Regression line with 95% confidence bands for the location memory data	182
5.9	The impact of forcing the intercept through the origin on the estimate of the regression slope	183
5.10	Regression line for a simple linear regression of percentage accuracy by group (dummy coded) for the comprehension data	189
5.11	Curvilinear fits to the location memory data	191
6.1	Samples of $N = 7000$ from a bivariate normal distribution	212
6.2	The impact of selecting the extremes of X (extreme groups analysis) on the unstandardized slope (upper panels) and the standardized slope (lower panels) of a regression line	216
8.1	The relationship between statistical power and one-sided α	282
8.2	The relationship between statistical power and α when the separation of H_0 and H_1 is small	283
8.3	The influence of (a) larger and (b) smaller standard errors of the mean on statistical power ($1 - \beta$)	284
8.4	Comparing central and noncentral t distributions with 6 and 29 degrees of freedom	286
8.5	Statistical power as a function of n per group for effect sizes of $\delta = 1$ and $\delta = .667$	292
9.1	The influence of sample size ($n = 10$ or $n = 100$) on the sample standard deviation ($\hat{\sigma}$)	313
9.2	The influence of bin size on a histogram	315
9.3	True histograms for nine random samples ($n = 30$) from the same normal population ($\mu = 20$ and $\sigma = 4$)	316
9.4	True histograms with lines plotting kernel density estimates or the population distribution for either a normal distribution (upper panels) or χ^2 distribution (lower panels)	317
9.5	Box plots of samples from a (positively skewed) lognormal distribution	319
9.6	Box plots of random samples from a (symmetrical) normal distribution	319
9.7	A basic box plot of the Hayden (2005) data	320
9.8	Histogram and kernel density estimate for the Hayden (2005) data	321
9.9	The Hayden (2005) data displayed (a) as a bean plot, and (b) as a density estimate with a rug	322
9.10	Normal probability plots of samples from normal, leptokurtotic or positively skewed distributions	326
9.11	Normal probability plot of standardized residuals from a t test of the comprehension data versus the expected quantiles from a normal distribution	327
9.12	Two data sets with identical simple linear regression fits	329
9.13	Detecting a violation of the homogeneity of variance assumption in simple linear regression	329
9.14	Standardized residuals versus X for two data sets	330
9.15	Scatter plots of samples from bivariate normal distributions	331
10.1	Comparing least squares and robust regression fits for the data in Figure 9.13	348
10.2	Large ($n = 10,000$) samples from a Poisson distribution	353
11.1	The variability of 95% CIs, p and p_{rep} for 20 simulated samples of $n = 25$ from a normal population	375

11.2	Functions for (a) the likelihood and (b) the probability of observing five heads from ten tosses of a fair coin	378
11.3	Two ways of scaling a likelihood function for 19 successes from a binomial distribution with 25 trials;	379
11.4	The 1/8 and 1/32 likelihood intervals for 19 successes from a binomial distribution with 25 trials	382
11.5	Poisson likelihood intervals for (a) $\hat{\lambda} = 3$, and (b) $\hat{\lambda} = 0$	383
11.6	The 1/8 and 1/32 likelihood intervals for the sleep loss data	387
11.7	How different priors influence the posterior distribution for the same sample data	391
11.8	Probability density for an inverse χ^2 distribution with 1 <i>df</i>	398
11.9	The JZS standard Cauchy prior and the unit-information prior for effect size	399
12.1	The best-fitting regression plane for a data set with two predictors	428
12.2	Venn diagrams depicting shared and non-shared contribution R^2 in multiple regression	443
12.3	Scatter plot matrix of outcome and predictor variables for the intention to vote data	453
13.1	Deviations of level means from grand mean for a one-way ANOVA of the diagram data	479
13.2	Normal probability plot of standardized residuals for one-way ANOVA of the diagram data	486
13.3	Tukey's HSD tests with robust standard errors	496
13.4	ANCOVA with two groups and one covariate	504
14.1	Density plots of predictors and outcome for lottery data analysis	532
14.2	Plots of the bilinear interaction of perceived personal finances and future income on money required to part with a lottery ticket	538
14.3	Simple slope of personal finances (with 90% confidence bands) as a function of future income	539
14.4	Examples of linear, quadratic and cubic functions of X	541
14.5	Linear and quadratic fits to the curvilinear Anscombe data	542
14.6	Interaction plot of cell means for the data in Table 14.1	545
14.7	Alternative interaction plot of cell means for the data in Table 14.1	546
14.8	Examples of potential patterns of cell means in 2×2 ANOVA	549
14.9	Alternative plots for a 5×3 ANOVA	556
14.10	Interaction plot of the cell means by condition and age for the prospective memory data	559
14.11	Three-way interaction plots for ANOVA with and without interaction	561
14.12	How a ceiling effect can influence detection of an interaction	570
16.1	Deviations from the grand mean for a one-way repeated measures ANOVA	624
16.2	Two-tiered error bar plot for the one-way ANOVA of the pride data	638
16.3	Cell means by condition and emotion for the pride data	641
17.1	A sigmoidal curve produced by the inverse of the logistic function	674
17.2	Probability distribution functions comparing the logistic distribution with a scaled normal distribution	675
17.3	Comparing linear and logistic regression fits for a simulated data set with a single continuous predictor	680
17.4	Predicted probability of expenses problem by parliamentary majority	682