

DISCOVERING STATISTICS USING R

NDY FIELD JEREMY MILES ZOE FIELD

GAGE

Los Angeles | London | New Delhi
Singapore | Washington DC

CONTENTS

•Kggwr
BSSm

Preface	xxi
How to use this book	xxv
Acknowledgements	xxix
Dedication	xxx
Symbols used in this book	xxxii
Some maths revision	xxxiv

1 Why is my evil lecturer forcing me to learn statistics? 1

1.1. What will this chapter tell me? ©	1
1.2. What the hell am I doing here? I don't belong here ©	2
1.3. Initial observation: finding something that needs explaining ©	4
1.4. Generating theories and testing them ©	4
1.5. Data collection 1: what to measure ©	7
1.5.1. Variables©	7
1.5.2. Measurement error ©	11
1.5.3. Validity and reliability ©	12
1.6. Data collection 2: how to measure ©	13
1.6.1. Correlational research methods ©	13
1.6.2. Experimental research methods ©	13
1.6.3. Randomization©	17
1.7. Analysing data ©	19
1.7.1. Frequency distributions ©	19
1.7.2. The centre of a distribution ©	21
1.7.3. The dispersion in a distribution ©	24
1.7.4. Using a frequency distribution to go beyond the data ©	25
1.7.5. Fitting statistical models to the data ©	28
What have I discovered about statistics? ©	29
Key terms that I've discovered	29
Smart Alex's tasks	30
Further reading	31
Interesting real research	31

2 Everything you ever wanted to know about statistics (well, sort of) 32

2.1. What will this chapter tell me? ©	32
2.2. Building statistical models ©	33

2.3.	Populations and samples ©	36
2.4.	Simple statistical models ©	36
2.4.1.	The mean: a very simple statistical model ©	36
2.4.2.	Assessing the fit of the mean: sums of squares, variance and standard deviations ©	37
2.4.3.	Expressing the mean as a model ©	40
2.5.	Going beyond the data ©	41
2.5.1.	The standard error ©	42
2.5.2.	Confidence intervals ©	43
2.6.	Using statistical models to test research questions ©	49
2.6.1.	Test statistics ©	53
2.6.2.	One- and two-tailed tests ©	55
2.6.3.	Type I and Type II errors ©	56
2.6.4.	Effect sizes ©	57
2.6.5.	Statistical power ©	58
	What have I discovered about statistics? ©	59
	Key terms that I've discovered	60
	Smart Alex's tasks	60
	Further reading	60
	Interesting real research	61
3	The R environment	62
3.1.	What will this chapter tell me? ©	62
3.2.	Before you start ©	63
3.2.1.	The R-chitecture ©	63
3.2.2.	Pros and cons of R ©	64
3.2.3.	Downloading and installing R ©	65
3.2.4.	Versions of R ©	66
3.3.	Getting started ©	66
3.3.1.	The main windows in R ©	67
3.3.2.	Menus in R ©	67
3.4.	Using R ©	71
3.4.1.	Commands, objects and functions ©	71
3.4.2.	Using scripts ©	75
3.4.3.	The R workspace ©	76
3.4.4.	Setting a working directory ©	77
3.4.5.	Installing packages ©	78
3.4.6.	Getting help ©	80
3.5.	Getting data into R ©	81
3.5.1.	Creating variables ©	81
3.5.2.	Creating dataframes ©	81
3.5.3.	Calculating new variables from existing ones ©	83
3.5.4.	Organizing your data ©	85
3.5.5.	Missing values ©	92
3.6.	Entering data with R Commander ©	92
3.6.1.	Creating variables and entering data with R Commander ©	94
3.6.2.	Creating coding variables with R Commander ©	95
3.7.	Using other software to enter and edit data ©	95
3.7.1.	Importing data ©	97
3.7.2.	Importing SPSS data files directly ©	99

3.7.3.	Importing data with R Commander©	101
3.7.4.	Things that can go wrong ©	102
3.8.	Saving data©	103
3.9.	Manipulating data ©	103
3.9.1.	Selecting parts of a dataframe ©	103
3.9.2.	Selecting data with the subset() function ©	105
3.9.3.	Dataframes and matrices ©	106
3.9.4.	Reshaping data®	107
	What have I discovered about statistics? ©	113
	R packages used in this chapter	113
	R functions used in this chapter	113
	Key terms that I've discovered	114
	Smart Alex's tasks	114
	Further reading	115
4	Exploring data with graphs	116
4.1.	What will this chapter tell me? ©	116
4.2.	The art of presenting data ©	117
4.2.1.	Why do we need graphs ©	117
4.2.2.	What makes a good graph? ©	117
4.2.3.	Lies, damned lies, and ... erm ... graphs©	120
4.3.	Packages used in this chapter©	121
4.4.	Introducing ggplot2 ©	121
4.4.1.	The anatomy of a plot ©	121
4.4.2.	Geometric objects (geoms) ©	123
4.4.3.	Aesthetics©	125
4.4.4.	The anatomy of the ggplotQ function ©	127
4.4.5.	Stats and geoms ©	128
4.4.6.	Avoiding overplotting ©	130
4.4.7.	Saving graphs ©	131
4.4.8.	Putting it all together: a quick tutorial ©	132
4.5.	Graphing relationships: the scatterplot©	136
4.5.1.	Simple scatterplot ©	136
4.5.2.	Adding a funky line©	138
4.5.3.	Grouped scatterplot ©	140
4.6.	Histograms: a good way to spot obvious problems ©	142
4.7.	Boxplots (box-whisker diagrams) ©	144
4.8.	Density plots ©	148
4.9.	Graphing means®	149
4.9.1.	Bar charts and error bars ©	149
4.9.2.	Line graphs©	155
4.10.	Themes and options ©	161
	What have I discovered about statistics? ©	163
	R packages used in this chapter	163
	R functions used in this chapter	164
	Key terms that I've discovered	164
	Smart Alex's tasks	164
	Further reading	164
	Interesting real research	165

5	Exploring assumptions	166
5.1.	What will this chapter tell me? ©	166
5.2.	What are assumptions? ©	167
5.3.	Assumptions of parametric data ©	167
5.4.	Packages used in this chapter ©	169
5.5.	The assumption of normality ©	169
5.5.1.	Oh no, it's that pesky frequency distribution again: checking normality visually ©	169
5.5.2.	Quantifying normality with numbers ©	173
5.5.3.	Exploring groups of data ©	177
5.6.	Testing whether a distribution is normal ©	182
5.6.1.	Doing the Shapiro-Wilk test in R ©	182
5.6.2.	Reporting the Shapiro-Wilk test ©	185
5.7.	Testing for homogeneity of variance ©	185
5.7.1.	Levene's test ©	186
5.7.2.	Reporting Levene's test ©	188
5.7.3.	Hartley's F_{max} the variance ratio ©	189
5.8.	Correcting problems in the data ©	190
5.8.1.	Dealing with outliers ©	190
5.8.2.	Dealing with non-normality and unequal variances ©	191
5.8.3.	Transforming the data using R ©	194
5.8.4.	When it all goes horribly wrong ©	201
	What have I discovered about statistics? ©	203
	R packages used in this chapter	204
	R functions used in this chapter	204
	Key terms that I've discovered	204
	Smart Alex's tasks	204
	Further reading	204
6	Correlation	205
6.1.	What will this chapter tell me? ©	205
6.2.	Looking at relationships ©	206
6.3.	How do we measure relationships? ©	206
6.3.1.	A detour into the murky world of covariance ©	206
6.3.2.	Standardization and the correlation coefficient ©	208
6.3.3.	The significance of the correlation coefficient ©	210
6.3.4.	Confidence intervals for r ©	211
6.3.5.	A word of warning about interpretation: causality ©	212
6.4.	Data entry for correlation analysis ©	213
6.5.	Bivariate correlation ©	213
6.5.1.	Packages for correlation analysis in R ©	214
6.5.2.	General procedure for correlations using R Commander ©	214
6.5.3.	General procedure for correlations using R ©	216
6.5.4.	Pearson's correlation coefficient ©	219
6.5.5.	Spearman's correlation coefficient ©	223
6.5.6.	Kendall's tau (non-parametric) ©	225
6.5.7.	Bootstrapping correlations ©	226
6.5.8.	Biserial and point-biserial correlations ©	229

6.6.	Partial correlation ©	234
6.6.1.	The theory behind part and partial correlation ©	234
6.6.2.	Partial correlation using R ©	235
6.6.3.	Semi-partial (or part) correlations ©	237
6.7.	Comparing correlations ©	238
6.7.1.	Comparing independent r s ©	238
6.7.2.	Comparing dependent r s ©	239
6.8.	Calculating the effect size ©	240
6.9.	How to report correlation coefficients ©	240
	What have I discovered about statistics? ©	242
	R packages used in this chapter	243
	R functions used in this chapter	243
	Key terms that I've discovered	243
	Smart Alex's tasks ©	243
	Further reading	244
	Interesting real research	244
7	Regression	245
7.1.	What will this chapter tell me? ©	245
7.2.	An introduction to regression ©	246
7.2.1.	Some important information about straight lines ©	247
7.2.2.	The method of least squares ©	248
7.2.3.	Assessing the goodness of fit: sums of squares, R and R^2 ©	249
7.2.4.	Assessing individual predictors ©	252
7.3.	Packages used in this chapter ©	253
7.4.	General procedure for regression in R ©	254
7.4.1.	Doing simple regression using R Commander ©	254
7.4.2.	Regression in R ©	255
7.5.	Interpreting a simple regression ©	257
7.5.1.	Overall fit of the object model ©	258
7.5.2.	Model parameters ©	259
7.5.3.	Using the model ©	260
7.6.	Multiple regression: the basics ©	261
7.6.1.	An example of a multiple regression model ©	261
7.6.2.	Sums of squares, R and R^2 ©	262
7.6.3.	Parsimony-adjusted measures of fit ©	263
7.6.4.	Methods of regression ©	263
7.7.	How accurate is my regression model? ©	266
7.7.1.	Assessing the regression model I: diagnostics ©	266
7.7.2.	Assessing the regression model II: generalization ©	271
7.8.	How to do multiple regression using R Commander and R ©	276
7.8.1.	Some things to think about before the analysis ©	276
7.8.2.	Multiple regression: running the basic model ©	277
7.8.3.	Interpreting the basic multiple regression ©	280
7.8.4.	Comparing models ©	284
7.9.	Testing the accuracy of your regression model ©	287
7.9.1.	Diagnostic tests using R Commander ©	287
7.9.2.	Outliers and influential cases ©	288

7.9.3.	Assessing the assumption of independence ©	291
7.9.4.	Assessing the assumption of no multicollinearity ©	292
7.9.5.	Checking assumptions about the residuals ©	294
7.9.6.	What if I violate an assumption? ©	298
7.10.	Robust regression: bootstrapping (D)	298
7.11.	How to report multiple regression ©	301
7.12.	Categorical predictors and multiple regression ©	302
7.12.1.	Dummy coding®	302
7.12.2.	Regression with dummy variables ®	305
	What have I discovered about statistics? ©	308
	R packages used in this chapter	309
	R functions used in this chapter	309
	Key terms that I've discovered	309
	Smart Alex's tasks	310
	Further reading	311
	Interesting real research	311

Logistic regression

312

8.1.	What will this chapter tell me? ©	312
8.2.	Background to logistic regression ©	313
8.3.	What are the principles behind logistic regression? ©	313
8.3.1.	Assessing the model: the log-likelihood statistic ©	315
8.3.2.	Assessing the model: the deviance statistic ©	316
8.3.3.	Assessing the model: R and R^2 ©	316
8.3.4.	Assessing the model: information criteria ©	318
8.3.5.	Assessing the contribution of predictors: the z-statistic ©	318
8.3.6.	The odds ratio ©	319
8.3.7.	Methods of logistic regression ©	320
8.4.	Assumptions and things that can go wrong ©	321
8.4.1.	Assumptions©	321
8.4.2.	Incomplete information from the predictors ©	322
8.4.3.	Complete separation ©	323
8.5.	Packages used in this chapter ©	325
8.6.	Binary logistic regression: an example that will make you feel eel ©	325
8.6.1.	Preparing the data	326
8.6.2.	The main logistic regression analysis ©	327
8.6.3.	Basic logistic regression analysis using R ©	329
8.6.4.	Interpreting a basic logistic regression ©	330
8.6.5.	Model 1: Intervention only ©	330
8.6.6.	Model 2: Intervention and Duration as predictors ©	336
8.6.7.	Casewise diagnostics in logistic regression ©	338
8.6.8.	Calculating the effect size ©	341
8.7.	How to report logistic regression ©	341
8.8.	Testing assumptions: another example ©	342
8.8.1.	Testing for multicollinearity ©	343
8.8.2.	Testing for linearity of the logit ©	344
8.9.	Predicting several categories: multinomial logistic regression ©	346
8.9.1.	Running multinomial logistic regression in R ©	347
8.9.2.	Interpreting the multinomial logistic regression output ©	350

8.9.3.	Reporting the results	355
	What have I discovered about statistics? ©	355
	R packages used in this chapter	356
	R functions used in this chapter	356
	Key terms that I've discovered	356
	Smart Alex's tasks	357
	Further reading	358
	Interesting real research	358
9	Comparing two means	359
9.1.	What will this chapter tell me? ©	359
9.2.	Packages used in this chapter ©	360
9.3.	Looking at differences ©	360
9.3.1.	A problem with error bar graphs of repeated-measures designs ©	361
9.3.2.	Step 1: calculate the mean for each participant ©	364
9.3.3.	Step 2: calculate the grand mean ©	364
9.3.4.	Step 3: calculate the adjustment factor ©	364
9.3.5.	Step 4: create adjusted values for each variable ©	365
9.4.	The f-test ©	368
9.4.1.	Rationale for the f-test ©	369
9.4.2.	The f-test as a general linear model ©	370
9.4.3.	Assumptions of the f-test ©	372
9.5.	The independent f-test ©	372
9.5.1.	The independent f-test equation explained ©	372
9.5.2.	Doing the independent f-test ©	375
9.6.	The dependent f-test ©	386
9.6.1.	Sampling distributions and the standard error ©	386
9.6.2.	The dependent f-test equation explained ©	387
9.6.3.	Dependent f-tests using R ©	388
9.7.	Between groups or repeated measures? ©	394
	What have I discovered about statistics? ©	395
	R packages used in this chapter	396
	R functions used in this chapter	396
	Key terms that I've discovered	396
	Smart Alex's tasks	396
	Further reading	397
	Interesting real research	397
10	Comparing several means: ANOVA (GLM 1)	398
10.1.	What will this chapter tell me? ©	398
10.2.	The theory behind ANOVA ©	399
10.2.1	Inflated error rates ©	399
10.2.2.	Interpreting F ©	400
10.2.3.	ANOVA as regression ©	400
10.2.4.	Logic of the F -ratio ©	405
10.2.5.	Total sum of squares (SS_T) ©	407
10.2.6.	Model sum of squares (SS_J) ©	409
10.2.7.	Residual sum of squares (SS_R) ©	410
10.2.8.	Mean squares ©	411

10.2.9.	The F-ratio ©	411
10.3.	Assumptions of ANOVA ©	412
10.3.1.	Homogeneity of variance ©	412
10.3.2.	Is ANOVA robust? ©	412
10.4.	Planned contrasts ©	414
10.4.1.	Choosing which contrasts to do ©	415
10.4.2.	Defining contrasts using weights ©	419
10.4.3.	Non-orthogonal comparisons ©	425
10.4.4.	Standard contrasts ©	426
10.4.5.	Polynomial contrasts: trend analysis ©	427
10.5.	<i>Post hoc</i> procedures ©	428
10.5.1.	<i>Post hoc</i> procedures and Type I (a) and Type II error rates ©	431
10.5.2.	<i>Posthoc</i> procedures and violations of test assumptions ©	431
10.5.3.	Summary of <i>posthoc</i> procedures ©	432
10.6.	One-way ANOVA using R ©	432
10.6.1.	Packages for one-way ANOVA in R ©	433
10.6.2.	General procedure for one-way ANOVA ©	433
10.6.3.	Entering data ©	433
10.6.4.	One-way ANOVA using R Commander ©	434
10.6.5.	Exploring the data ©	436
10.6.6.	The main analysis ©	438
10.6.7.	Planned contrasts using R ©	443
10.6.8.	<i>Posthoc</i> tests using R ©	447
10.7.	Calculating the effect size ©	454
10.8.	Reporting results from one-way independent ANOVA ©	457
	What have I discovered about statistics? ©	458
	R packages used in this chapter	459
	R functions used in this chapter	459
	Key terms that I've discovered	459
	Smart Alex's tasks	459
	Further reading	461
	Interesting real research	461

Analysis of covariance, ANCOVA (GLM 2) 462

11.1.	What will this chapter tell me? ©	462
11.2.	What is ANCOVA? ©	463
11.3.	Assumptions and issues in ANCOVA ©	464
11.3.1.	Independence of the covariate and treatment effect ©	464
11.3.2.	Homogeneity of regression slopes ©	466
11.4.	ANCOVA using R ©	467
11.4.1.	Packages for ANCOVA in R ©	467
11.4.2.	General procedure for ANCOVA ©	468
11.4.3.	Entering data ©	468
11.4.4.	ANCOVA using R Commander ©	471
11.4.5.	Exploring the data ©	471
11.4.6.	Are the predictor variable and covariate independent? ©	473
11.4.7.	Fitting an ANCOVA model ©	473
11.4.8.	Interpreting the main ANCOVA model ©	477

11.4.9.	Planned contrasts in ANCOVA ©	479
11.4.10.	Interpreting the covariate (D)	480
11.4.11.	<i>Post hoc</i> tests in ANCOVA ©	481
11.4.12.	Plots in ANCOVA ©	482
11.4.13.	Some final remarks ©	482
11.4.14.	Testing for homogeneity of regression slopes ®	483
11.5.	Robust ANCOVA ©	484
11.6.	Calculating the effect size ©	491
11.7.	Reporting results ©	494
	What have I discovered about statistics? ©	495
	R packages used in this chapter	495
	R functions used in this chapter	496
	Key terms that I've discovered	496
	Smart Alex's tasks	496
	Further reading	497
	Interesting real research	497

12 Factorial ANOVA (GLM 3) 498

12.1.	What will this chapter tell me? ©	498
12.2.	Theory of factorial ANOVA (independent design) ©	499
12.2.1.	Factorial designs ©	499
12.3.	Factorial ANOVA as regression ®	501
12.3.1.	An example with two independent variables ©	501
12.3.2.	Extending the regression model ©	501
12.4.	Two-way ANOVA: behind the scenes ©	505
12.4.1.	Total sums of squares (SS_T) ©	506
12.4.2.	The model sum of squares (SS_J) ©	507
12.4.3.	The residual sum of squares (SS_R) ©	510
12.4.4.	The F -ratios ©	511
12.5.	Factorial ANOVA using R ©	511
12.5.1.	Packages for factorial ANOVA in R ©	511
12.5.2.	General procedure for factorial ANOVA ©	512
12.5.3.	Factorial ANOVA using R Commander ©	512
12.5.4.	Entering the data ©	513
12.5.5.	Exploring the data ©	516
12.5.6.	Choosing contrasts ©	518
12.5.7.	Fitting a factorial ANOVA model ©	520
12.5.8.	Interpreting factorial ANOVA ©	520
12.5.9.	Interpreting contrasts ©	524
12.5.10.	Simple effects analysis ©	525
12.5.11.	<i>Post hoc</i> analysis ©	528
12.5.12.	Overall conclusions	530
12.5.13.	Plots in factorial ANOVA ©	530
12.6.	Interpreting interaction graphs ©	530
12.7.	Robust factorial ANOVA ©	534
12.8.	Calculating effect sizes ©	542
12.9.	Reporting the results of two-way ANOVA ©	544
	What have I discovered about statistics? ©	546

R packages used in this chapter	546
R functions used in this chapter	546
Key terms that I've discovered	547
Smart Alex's tasks	547
Further reading	548
Interesting real research	548
13 Repeated-measures designs (GLM 4)	549
13.1. What will this chapter tell me? ©	549
13.2. Introduction to repeated-measures designs ©	550
13.2.1. The assumption of sphericity ©	551
13.2.2. How is sphericity measured? ©	551
13.2.3. Assessing the severity of departures from sphericity ©	552
13.2.4. What is the effect of violating the assumption of sphericity? ®	552
13.2.5. What do you do if you violate sphericity? ©	554
13.3. Theory of one-way repeated-measures ANOVA ©	554
13.3.1. The total sum of squares (SS_T) ©	557
13.3.2. The within-participant sum of squares (SS_J) ©	558
13.3.3. The model sum of squares (SS_J) ©	559
13.3.4. The residual sum of squares (SS_R) ©	560
13.3.5. The mean squares ©	560
13.3.6. The F-ratio ©	560
13.3.7. The between-participant sum of squares ©	561
13.4. One-way repeated-measures designs using R ©	561
13.4.1. Packages for repeated measures designs in R ©	561
13.4.2. General procedure for repeated-measures designs ©	562
13.4.3. Repeated-measures ANOVA using R Commander ©	563
13.4.4. Entering the data ©	563
13.4.5. Exploring the data ©	565
13.4.6. Choosing contrasts ©	568
13.4.7. Analysing repeated measures: two ways to skin a .dat ©	569
13.4.8. Robust one-way repeated-measures ANOVA ®	576
13.5. Effect sizes for repeated-measures designs ©	580
13.6. Reporting one-way repeated-measures designs ©	581
13.7. Factorial repeated-measures designs ©	583
13.7.1. Entering the data ©	584
13.7.2. Exploring the data ©	586
13.7.3. Setting contrasts ©	588
13.7.4. Factorial repeated-measures ANOVA ©	589
13.7.5. Factorial repeated-measures designs as a GLM ®	594
13.7.6. Robust factorial repeated-measures ANOVA ©	599
13.8. Effect sizes for factorial repeated-measures designs ©	599
13.9. Reporting the results from factorial repeated-measures designs ©	600
What have I discovered about statistics? ©	601
R packages used in this chapter	602
R functions used in this chapter	602
Key terms that I've discovered	602
Smart Alex's tasks	602

Further reading	603
Interesting real research	603
14 Mixed designs (GLM 5)	604
14.1. What will this chapter tell me? ©	604
14.2. Mixed designs ©	605
14.3. What do men and women look for in a partner? ©	606
14.4. Entering and exploring your data ©	606
14.4.1. Packages for mixed designs in R ©	606
14.4.2. General procedure for mixed designs ©	608
14.4.3. Entering the data (D)	608
14.4.4. Exploring the data ©	610
14.5. Mixed ANOVA©	613
14.6. Mixed designs as a GLM ®	617
14.6.1. Setting contrasts ©	617
14.6.2. Building the model ©	619
14.6.3. The main effect of gender ©	622
14.6.4. The main effect of looks ©	623
14.6.5. The main effect of personality ©	624
14.6.6. The interaction between gender and looks ©	625
14.6.7. The interaction between gender and personality ©	628
14.6.8. The interaction between looks and personality ©	630
14.6.9. The interaction between looks, personality and gender ®	635
14.6.10. Conclusions©	639
14.7. Calculating effect sizes ®	640
14.8. Reporting the results of mixed ANOVA ©	641
14.9. Robust analysis for mixed designs ®	643
What have I discovered about statistics? ©	650
R packages used in this chapter	650
R functions used in this chapter	651
Key terms that I've discovered	651
Smart Alex's tasks	651
Further reading	652
Interesting real research	652
15 Non-parametric tests	653
15.1. What will this chapter tell me? ©	653
15.2. When to use non-parametric tests ©	654
15.3. Packages used in this chapter ©	655
15.4. Comparing two independent conditions: the Wilcoxon rank-sum test ©	655
15.4.1. Theory of the Wilcoxon rank-sum test ©	655
15.4.2. Inputting data and provisional analysis ©	659
15.4.3. Running the analysis using R Commander ©	661
15.4.4. Running the analysis using R ©	662
15.4.5. Output from the Wilcoxon rank-sum test ©	664
15.4.6. Calculating an effect size ©	664
15.4.7. Writing the results ©	666

15.5.	Comparing two related conditions: the Wilcoxon signed-rank test ©	667
15.5.1.	Theory of the Wilcoxon signed-rank test ©	668
15.5.2.	Running the analysis with R Commander ©	670
15.5.3.	Running the analysis using R ©	671
15.5.4.	Wilcoxon signed-rank test output ©	672
15.5.5.	Calculating an effect size ©	673
15.5.6.	Writing the results ©	673
15.6.	Differences between several independent groups: the Kruskal-Wallis test ©	674
15.6.1.	Theory of the Kruskal-Wallis test ©	675
15.6.2.	Inputting data and provisional analysis ©	677
15.6.3.	Doing the Kruskal-Wallis test using R Commander ©	679
15.6.4.	Doing the Kruskal-Wallis test using R ©	679
15.6.5.	Output from the Kruskal-Wallis test ©	680
15.6.6.	<i>Post hoc</i> tests for the Kruskal-Wallis test ©	681
15.6.7.	Testing for trends: the Jonckheere-Terpstra test ©	684
15.6.8.	Calculating an effect size ©	685
15.6.9.	Writing and interpreting the results ©	686
15.7.	Differences between several related groups: Friedman's ANOVA ©	686
15.7.1.	Theory of Friedman's ANOVA ©	688
15.7.2.	Inputting data and provisional analysis ©	689
15.7.3.	Doing Friedman's ANOVA in R Commander ©	690
15.7.4.	Friedman's ANOVA using R ©	690
15.7.5.	Output from Friedman's ANOVA ©	691
15.7.6.	<i>Post hoc</i> tests for Friedman's ANOVA ©	691
15.7.7.	Calculating an effect size ©	692
15.7.8.	Writing and interpreting the results ©	692
	What have I discovered about statistics? ©	693
	R packages used in this chapter	693
	R functions used in this chapter	693
	Key terms that I've discovered	694
	Smart Alex's tasks	694
	Further reading	695
	Interesting real research	695

16 Multivariate analysis of variance (MANOVA) 696

16.1.	What will this chapter tell me? ©	696
16.2.	When to use MANOVA ©	697
16.3.	Introduction: similarities to and differences from ANOVA ©	697
16.3.1.	Words of warning ©	699
16.3.2.	The example for this chapter ©	699
16.4.	Theory of MANOVA ©	700
16.4.1.	Introduction to matrices ©	700
16.4.2.	Some important matrices and their functions ©	702
16.4.3.	Calculating MANOVA by hand: a worked example ©	703
16.4.4.	Principle of the MANOVA test statistic ©	710
16.5.	Practical issues when conducting MANOVA ©	717
16.5.1.	Assumptions and how to check them ©	717

16.5.2.	Choosing a test statistic ®	718
16.5.3.	Follow-up analysis ®	719
16.6.	MANOVA using R ©	719
16.6.1.	Packages for factorial ANOVA in R ©	719
16.6.2.	General procedure for MANOVA ©	720
16.6.3.	MANOVA using R Commander ©	720
16.6.4.	Entering the data ©	720
16.6.5.	Exploring the data ©	722
16.6.6.	Setting contrasts ©	728
16.6.7.	The MANOVA model ©	728
16.6.8.	Follow-up analysis: univariate test statistics ©	731
16.6.9.	Contrasts®	732
16.7.	Robust MANOVA ®	733
16.8.	Reporting results from MANOVA ©	737
16.9.	Following up MANOVA with discriminant analysis ©	738
16.10.	Reporting results from discriminant analysis ©	743
16.11.	Some final remarks ©	743
16.11.1.	The final interpretation ©	743
16.11.2.	Univariate ANOVA or discriminant analysis?	745
	What have I discovered about statistics? ©	745
	R packages used in this chapter	746
	R functions used in this chapter	746
	Key terms that I've discovered	747
	Smart Alex's tasks	747
	Further reading	748
	Interesting real research	748

17 Exploratory factor analysis

749

17.1.	What will this chapter tell me? ©	749
17.2.	When to use factor analysis ©	750
17.3.	Factors©	751
17.3.1.	Graphical representation of factors ©	752
17.3.2.	Mathematical representation of factors ©	753
17.3.3.	Factor scores ©	755
17.3.4.	Choosing a method ©	758
17.3.5.	Communality©	759
17.3.6.	Factor analysis vs. principal components analysis ©	760
17.3.7.	Theory behind principal components analysis ©	761
17.3.8.	Factor extraction: eigenvalues and the scree plot ©	762
17.3.9.	Improving interpretation: factor rotation ©	764
17.4.	Research example ©	767
17.4.1.	Sample size©	769
17.4.2.	Correlations between variables ©	770
17.4.3.	The distribution of data ©	772
17.5.	Running the analysis with R Commander	772
17.6.	Running the analysis with R	772
17.6.1.	Packages used in this chapter ©	772
17.6.2.	Initial preparation and analysis	772

17.6.3.	Factor extraction using R ©	778
17.6.4.	Rotation©	788
17.6.5.	Factor scores (D)	793
17.6.6.	Summary©	795
17.7.	How to report factor analysis ©	795
17.8.	Reliability analysis ©	797
17.8.1.	Measures of reliability ®	797
17.8.2.	Interpreting Cronbach's α (some cautionary tales ...) ©	799
17.8.3.	Reliability analysis with R Commander	800
17.8.4.	Reliability analysis using R ©	800
17.8.5.	Interpreting the output ©	801
17.9.	Reporting reliability analysis ©	806
	What have I discovered about statistics? ©	807
	R packages used in this chapter	807
	R functions used in this chapter	808
	Key terms that I've discovered	808
	Smart Alex's tasks	808
	Further reading	810
	Interesting real research	811

18 Categorical data

812

18.1.	What will this chapter tell me? ©	812
18.2.	Packages used in this chapter ©	813
18.3.	Analysing categorical data ©	813
18.4.	Theory of analysing categorical data ©	814
18.4.1.	Pearson's chi-square test ©	814
18.4.2.	Fisher's exact test ©	816
18.4.3.	The likelihood ratio ©	816
18.4.4.	Yates's correction ©	817
18.5.	Assumptions of the chi-square test ©	818
18.6.	Doing the chi-square test using R ©	818
18.6.1.	Entering data: raw scores ©	818
18.6.2.	Entering data: the contingency table ©	819
18.6.3.	Running the analysis with R Commander ©	820
18.6.4.	Running the analysis using R ©	821
18.6.5.	Output from the <i>CrossTableQ</i> function ©	822
18.6.6.	Breaking down a significant chi-square test with standardized residuals ©	825
18.6.7.	Calculating an effect size ©	826
18.6.8.	Reporting the results of chi-square ©	827
18.7.	Several categorical variables: loglinear analysis ©	829
18.7.1.	Chi-square as regression ©	829
18.7.2.	Loglinear analysis ®	835
18.8.	Assumptions in loglinear analysis ©	837
18.9.	Loglinear analysis using R ©	838
18.9.1.	Initial considerations ©	838
18.9.2.	Loglinear analysis as a chi-square test ©	840
18.9.3.	Output from loglinear analysis as a chi-square test ©	843

18.9.4.	Loglinear analysis ©	845
18.10.	Following up loglinear analysis ©	850
18.11.	Effect sizes in loglinear analysis ©	851
18.12.	Reporting the results of loglinear analysis ©	851
	What have I discovered about statistics? ©	852
	R packages used in this chapter	853
	R functions used in this chapter	853
	Key terms that I've discovered	853
	Smart Alex's tasks ©	853
	Further reading	854
	Interesting real research	854
19	Multilevel linear models	855
19.1.	What will this chapter tell me? ©	855
19.2.	Hierarchical data©	856
19.2.1.	The intraclass correlation ©	859
19.2.2.	Benefits of multilevel models ©	859
19.3.	Theory of multilevel linear models ©	860
19.3.1.	An example©	861
19.3.2.	Fixed and random coefficients ©	862
19.4.	The multilevel model ©	865
19.4.1.	Assessing the fit and comparing multilevel models©	867
19.4.2.	Types of covariance structures ©	868
19.5.	Some practical issues ©	870
19.5.1.	Assumptions©	870
19.5.2.	Sample size and power ©	870
19.5.3.	Centring variables ©	871
19.6.	Multilevel modelling in R ©	873
19.6.1.	Packages for multilevel modelling in R	873
19.6.2.	Entering the data ©	873
19.6.3.	Picturing the data©	874
19.6.4.	Ignoring the data structure: ANOVA©	874
19.6.5.	Ignoring the data structure: ANCOVA ©	876
19.6.6.	Assessing the need for a multilevel model ©	878
19.6.7.	Adding in fixed effects ©	881
19.6.8.	Introducing random slopes ©	884
19.6.9.	Adding an interaction term to the model ©	886
19.7.	Growth models©	892
19.7.1.	Growth curves (polynomials) ©	892
19.7.2.	An example: the honeymoon period ©	894
19.7.3.	Restructuring the data ©	895
19.7.4.	Setting up the basic model ©	895
19.7.5.	Adding in time as a fixed effect ©	897
19.7.6.	Introducing random slopes©	897
19.7.7.	Modelling the covariance structure ©	897
19.7.8.	Comparing models©	899
19.7.9.	Adding higher-order polynomials ©	901
19.7.10.	Further analysis ©	905

19.8. How to report a multilevel model®	906
What have I discovered about statistics? ©	907
R packages used in this chapter	908
R functions used in this chapter	908
Key terms that I've discovered	908
Smart Alex's tasks	908
Further reading	909
Interesting real research	909
Epilogue: life after discovering statistics	910
Troubleshooting R	912
Glossary	913
Appendix	929
A.1. Table of the standard normal distribution	929
A.2. Critical values of the f-distribution	935
A.3. Critical values of the F-distribution	936
A.4. Critical values of the chi-square distribution	940
References	941
Index	948
Functions in R	956
Packages in R	957