1

Basic Features of Statistical Analysis and the General Linear Model

INTRODUCTION

The aim of this book is to describe some of the statistical techniques which are becoming increasingly common, particularly in the social sciences. The spread of sophisticated computer packages and the machinery on which to run them has meant that procedures which were previously only available to experienced researchers with access to expensive machines and research students can now be carried out in a few seconds by almost every undergraduate. The tendency of the packages to produce items of output which are unfamiliar to most users has lead to modifications in the content of quantitative data analysis courses, but this has not always meant that students gain an understanding of what the analytic procedures do, when they should be used and what the results provided signify. Our aim has been to provide the basis for gaining such an understanding. There are many texts and Internet resources covering the material which we do, but our experience is that many of them are too advanced, starting at too high a level and including topics such as matrix algebra which leave many students baffled. What we have attempted to provide is an assistant which will help you make the transition from the simpler statistics (*t*-tests, analysis of variance) to more complex procedures; we are hoping we have translated the more technical texts into a form which matches your understanding. Each chapter provides an outline of the statistical technique, the type of question it answers, what the results produced tell you and gives examples from published literature of how the technique has been used.

In recent years there has been a considerable growth in the use of qualitative research methods in many areas of social science including psychology and nursing and this has been accompanied by a decline in the previous preponderance of quantitative research. One feature of the qualitative research movement has been an emphasis upon the ethical issues involved in carrying out research

involving people, the need to recognise that the participants own their data and that they should have an input – even perhaps a veto – over the interpretations made of it and the uses to which it is put. This concern has rejuvenated the ethical debate within quantitative research and brought back an awareness of the need to ensure that participants give informed consent to taking part, that they are not studied unknowingly or covertly, that they have the right to confidentiality and anonymity. This is not the place to discuss the ethics of research, but it is only proper that we should urge those considering quantitative research to be aware of the ethical guidelines applicable to their discipline and ensure they abide by them. Gathering the data which lends itself to quantitative analysis is not a value-free activity even if 'number crunching' may in itself appear to be so.

Before describing the more complex statistical techniques, we begin by recapitulating the basics of statistical analysis, reminding you of the analysis of variance and outlining the principles of the general linear model (GLM) which underpins many of the techniques described later.

BASIC FEATURES OF STATISTICAL ANALYSIS

Experiments or correlational research designs

In an experiment using a between-subjects design, the participants are randomly allocated to different levels of the independent variable and if all other variables are controlled by being kept constant or by the design of the experiment then it is assumed that any differences in the dependent variable measures are due to the independent variable. (This is a gross simplification of how to design an experiment!) But in many or even most fields of investigation it is impossible to carry out a true experiment because it is impossible to control the conditions, impossible to allocate participants randomly to conditions or ethically unacceptable to do so. One is then forced to consider an alternative type of investigation such as a pseudo-experiment or a correlational study in which data on independent and dependent variables is collected (often simultaneously) and the relationships between them are investigated.

Experiments typically involve analysing the data for differences: did group A score differently from group B on the dependent variable? Correlational studies usually involve analysing the data for correlations or associations: did those who score highly on measure X also score obtain high scores on measure Y?

Independent and dependent variables

Independent variables are those aspects of the respondents or cases which you anticipate will affect the output measure, the dependent variable. An independent

variable is often the 'grouping' variable which divides people, respondents or cases into separate groups. This division may be based on experimental conditions or it may be some characteristic of the participants such as their age group, sex, economic status. When the independent variable involves different participants in each group, it is referred to as a between-subjects variable. Alternatively, the independent variable may be a number of experimental conditions where all participants take part in every condition. If this is the case, the variable is a within-subjects factor and a repeated measures design is being used. A mixed design is where there are at least two independent variables, and one is between subjects while one is within subjects.

The dependent variable is usually a continuous variable such as a measure of performance on an experimental task or a score on a questionnaire which the researcher proposes is affected by the independent variables. In some types of research, the dependent variable is categorical, which means participants are divided into categories such as surviving and not surviving or relapsing and not relapsing. The data is then frequencies: how many people fall into each category? The research may be concerned with finding which factors predict category membership, and then logistic regression may be used to analyse the data.

It is important not to confuse variables with their levels. An independent variable is the experimental manipulation or the dimension upon which the participants are categorised. For example, suppose we were testing the ability of boys and girls to do mental arithmetic when they were distracted by background noise which could be loud, quiet or not present at all. We would design the experiment so that our participants carried out a mental arithmetic task with loud noise, with quiet noise or without noise. There would be two independent variables: participant sex (male or female) and noise condition (loud, quiet, absent). The first independent variable has two levels (male or female) and the second independent variable has three levels. So this would be a 2×3 (or 3×2 since it does not matter in which order the numbers here are presented) experiment. The expression 2×3 contains two digits, showing there are two independent variables. The actual digits, 2 and 3, indicate that one of the variables has two levels and the other has three levels.

Types of data

There are essentially two types of data, frequency and numerical, depending on the type of measurement scale used. One type of measurement scale is categorical or nominal, where cases are allocated to categories such as 'male' and 'female', or 'recovering after 2 months', 'recovering after 6 months', 'not recovering'. This yields frequency data which is obtained if you count the number of cases or people in each category. The other type of measurement scale is quantitative or numerical: here you are measuring not how many people or cases fall

into a particular category, but how much or how well they performed by assigning a numerical value to their performance, for example by recording the time taken to do a task or the number of questions answered.

In a nominal scale, the number is simply functioning as a label and the size of the number does not reflect the magnitude or order of the items. Telephone numbers are an example of a nominal scale, since the size of the number does not reflect anything about the size or order of the people who have those numbers. Similarly, in the example of performance under conditions of background noise described earlier, you might designate the loud noise condition as condition 1, the quiet noise condition as condition 2 and the no-noise condition as condition 3. Here the numbers are acting just as convenient labels for each condition or category and their size means nothing. When you have counted the number of cases in each category, you have frequency data which can be analysed using procedures such as chi-square, log–linear analysis or logistic regression.

Numerical data can be measured on a ratio, interval or ordinal scale. In a ratio scale there is a true zero and a number which is twice as large as another reflects the fact that the attribute being measured is twice as great. Ratio scales are rare in the social sciences unless one is using a measure of some physical feature such as height or time: someone who is 2 metres tall is twice as tall as someone who is 1 metre tall; someone who took 30 seconds to do a task took twice as long as someone who did it in 15 seconds. In an interval scale, the difference between two values at one point on the scale is the same as the difference between two equidistant values at another point on the scale: the usual example cited is the Fahrenheit scale where the difference between 15 and 20 degrees is the same as the difference between 5 and 10 degrees. There are few cases of interval scales in the social sciences (e.g. IQ is not an interval scale because the difference between an IQ of 100 and 105 is not the same as the difference between 70 and 75), although many examples of data being treated as though it were an interval scale. In an ordinal or rank scale, the size of the numbers reflects the order of the items as in a race where first came before second and second before third. But this information tells you nothing about the intervals between the scale points: first may have been just in front of second with third trailing far behind. In practice, the distinction between ratio and interval scales is widely ignored but ordinal or rank data is treated differently by using non-parametric procedures.

Non-parametric and parametric analyses

There are two groups of statistical techniques: parametric and non-parametric. Non-parametric techniques are considered distribution free, meaning that they do not involve any assumptions about the distribution of the population from which the sample of dependent variable measures is drawn. (It does not, for example, have to be normally distributed.) Non-parametric techniques are used

with frequency data and when the dependent variable has been measured on an ordinal (rank) scale. If the dependent variable has been measured on an interval scale but does not fulfil certain assumptions described below, it can be transformed into ordinal data and the non-parametric techniques can then be applied.

The parametric tests are generally considered more powerful, offer greater flexibility in analysis and address a greater number of research questions. The majority of the statistical techniques outlined in this book require that the dependent variable measures meet the requirements for being parametric data, which are:

1 the dependent variable is measured on either an interval or a ratio scale;
2 scores on the dependent variable approximate to a normal distribution or are drawn from a population where the variable can be assumed to be normally distributed;
3 scores on the dependent variable show homogeneity of variance between groups of participants.

Regarding point 1, strictly speaking parametric analysis should only be performed on continuous interval or ratio data but in practice many types of data are taken to be interval even if one could reasonably argue that they are not. An example is where a Likert scale has been used to measure attitudes. This is where participants are presented with a statement such as 'The death penalty is morally acceptable' and indicate their response by indicating how far they agree or disagree with it using a five- or seven-point scale with one end of the scale indicating 'strongly agree', and the other indicating 'strongly disagree'. If the data is ordinal (ranks), then non-parametric analysis is needed.

Concerning point 2, parametric statistical analysis is based on the assumption that the scores come from a normal distribution, meaning that if one could obtain the scores from the population then they would be normally distributed. Of course one does not know the distribution in the population, only the distribution in the sample one has. So it is necessary to check whether these approximate to a normal distribution. This can be done by plotting the distribution and examining it to see if it is more or less normal. The shape of the distribution can be evaluated in terms of skewness and kurtosis. Skewness reflects the positioning of the peak of the curve (is it in the centre?) and kurtosis refers to the height of the tails of the curve (is the curve too flat?). Statistical packages may give indices of skew and kurtosis; the values should be close to zero.

On point 3, homogeneity of variance is the assumption that the amount of variance is the same in the different sets of scores of the participants. It can be assessed using Levene's test for homogeneity of variance which gives a t value: if t is significant, the groups differ in their variances, that is there is heterogeneity of variance. If you are comparing groups of equal size, heterogeneity of variance is not important, but for groups of unequal sizes it needs to be dealt with. This

can be done in a number of ways including transforming the scores, using a more stringent significance level (perhaps 0.01 rather than 0.05), applying a non-parametric procedure.

Statistical significance

Probability testing is at the centre of statistical analysis and is essentially concerned with deciding how probable it is that the results observed could have been due to chance or error variation in the scores. To make the explanation simpler, we shall take the case of testing to see whether there is a difference between two groups of respondents. Suppose we have measured the amount of concern people have with their body image using a questionnaire in which a high score indicates a high level of concern, and done this for a group of women and for a group of men. The null hypothesis states that there is no difference between the scores of the two groups. The research (or alternative) hypothesis states that there is a difference between the scores of the two groups. The research hypothesis may predict the direction of the outcome (e.g. women will have a higher score than the men) in which case it is a directional or one-tailed hypothesis. Or the research hypothesis may just predict a difference between the two groups, without specifying which direction that difference will take (e.g. women and men will score differently from one another) in which case it is referred to as a non-directional or two-tailed hypothesis.

In our example, we want to know how likely it is that the difference in the mean scores of the women and men was the result of chance variation in the scores. You will probably recognise this as a situation in which you would turn to the t-test, and may remember that in the t-test you calculate the difference between the means of the two groups and express it as a ratio of the standard error of the difference which is calculated from the variance in the scores. If this ratio is greater than a certain amount, which you can find from the table for t, you can conclude that the difference between the means is unlikely to have arisen from the chance variability in the data and that there is a 'significant' difference between the means.

It is conventional to accept that 'unlikely' means having a 5% (0.05) probability or less. So if the probability of the difference arising by chance is 0.05 or less, you conclude it did not arise by chance. There are occasions when one uses a more stringent probability or significance level and only accepts the difference as significant if the probability of its arising by chance is 1% (0.01) or less. Much more rarely, one may accept a less stringent probability level such as 10% (0.1).

In considering the level of significance which you are going to use, there are two types of errors which need to be borne in mind. A Type I error occurs when a researcher accepts the research hypothesis and incorrectly rejects the null hypothesis. A Type II error occurs when the null hypothesis is accepted and the research

hypothesis is incorrectly rejected. When you use the 5% (0.05) significance level, you have a 5% chance of making a Type I error. You can reduce this by using a more stringent level such as 1% (0.01), but this increases the probability of making a Type II error.

When a number of significance tests are applied to a set of data it is generally considered necessary to apply some method of correcting for multiple testing. (If you carried out 100 *t*-tests, 5% of them are expected to come out 'significant' just by chance. So multiple significance testing can lead to accepting outcomes as significant when they are not, a Type I error.) To prevent the occurrence of a Type I error some researchers simply set the significance level needed to be reached at the 1% level, but this does seem rather arbitrary. A more precise correction for multiple testing is the Bonferroni correction in which the 0.05 probability level is divided by the number of times the same test is being used on the data set. For example, if four *t*-tests are being calculated on the same data set then 0.05 would be divided by 4 which would give a probability level of 0.0125 which would have to be met to achieve statistical significance.

RECAPITULATION OF ANALYSIS OF VARIANCE (ANOVA)

ANOVAs, like *t*-tests, examine the differences between group means. However, an ANOVA has greater flexibility than a *t*-test since the researcher is able to use more than two groups in the analysis. Additionally ANOVAs are able to consider the effect of more than one independent variable and to reveal whether the effect of one of them is influenced by the other: whether they interact.

Variance summarises the degree to which each score differs from the mean, and as implied by the name ANOVAs consider the amount of variance in a data set. Variance potentially arises from three sources: individual differences, error and the effect of the independent variable. The sources of variance in a set of data are expressed in the sum of squares value in an ANOVA. In a between-groups analysis of variance, the between-groups sum of squares represents the amount of variance accounted for by the effect of the independent variable with the estimate of error removed. The sums of squares are divided by the corresponding degrees of freedom to produce the mean square between groups and the mean square within groups. The ratio between these two figures produces a value for *F*. The further away from one the *F* value is, the lower the probability that the differences between the groups arose by chance and the higher the level of significance of the effect of the independent variable.

Repeated measures ANOVAs occur when the participants complete the same set of tasks under different conditions or in longitudinal studies where the same tasks are completed by the same participants on more than one occasion. There are additional requirements which need to be met in order to ensure that the

Table 1.1 Hypothetical data from an experiment comparing the effects of distracting noise on mental arithmetic performance for males and females

Noise	Sex of participants	
	Males	Females
Loud	9.46	10.34
Quiet	12.32	12.11
None	12.82	17.23

results from a repeated measures ANOVA are reliable. One of these is known as sphericity and tests for it are given in the output from the computerised statistical packages used to analyse data. If the test is significant, sphericity is a problem within the data. Sphericity is concerned with the similarity of the relationship between the dependent and independent variables in a repeated measures design. If the relationship between them changes over repeated measures of the dependent variable, the assumption of sphericity is violated and this will increase the chances of a Type I error occurring. To overcome violation of this assumption a stricter significance value can be set, or one can use corrections such as the Greenhouse–Geisser, Huynh–Feldt and lower bound epsilon formulae which adjust the degrees of freedom. (Statistical packages will print these adjusted degrees of freedom.)

Main effects, simple effects, Interaction

Our hypothetical experiment on boys and girls doing mental arithmetic with loud background noise, quiet noise or no noise would yield six means as illustrated in Table 1.1, where the numbers represent the number of mental arithmetic problems solved and have just been made up for the purpose of helping us explain some of the concepts. One would anticipate using analysis of variance to analyse a set of data like this.

When using ANOVA models, terms such as main effects, interactions and simple effects are encountered. It is worth reminding ourselves what these mean (especially as the concept of interactions seems to baffle many students).

A significant main effect demonstrates that an independent variable influences the dependent variable. In this example, a significant main effect of noise would establish that there was a difference between the overall means for the three noise conditions of loud, quiet and none. However, it does not establish which of the independent variable groups are significantly different from one another. Is the loud condition significantly different from the quiet condition? Is it significantly different from the none condition? Is the quiet condition significantly different
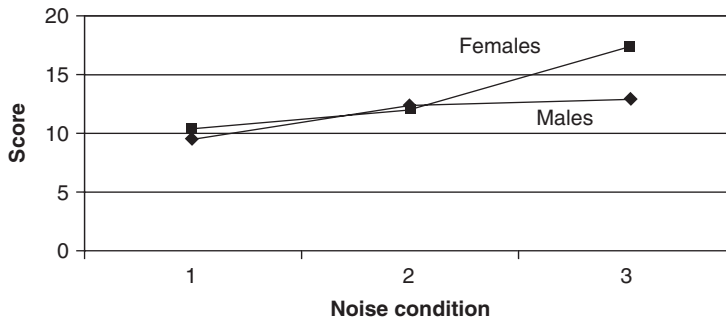
Figure 1.1   Graph of the (hypothetical) data from Table 1.1 showing an interaction

from the none condition? These questions are answered by testing the differences between the particular pairs of means of the groups using a priori or *post hoc* tests.

Whether to use a priori or *post hoc* tests depends on whether the researcher has previously stated the hypotheses to test. If you have honestly stated beforehand the comparisons between individual pairs of means which you intend to make, then you are entitled to use an a priori test such as a *t*-test. If you have looked at the data and then decided it looks worth comparing one mean with another, you have to use a *post hoc* test such as Newman–Keuls, Tukey or Scheffé.

An interaction occurs when the effect of one independent variable is affected by another independent variable. In our example, the difference between males and females seems to be greater in the none noise condition than it is in the loud or quiet conditions. This would be an interaction: the magnitude of the differences between levels on one variable (sex) is influenced by the level on the other variable (noise). (An interaction can be shown graphically if you plot the data as shown in Figure 1.1. The lines representing the data for the two sexes are not parallel, and this suggests an interaction.) Interactions can be two way, three way or more complex depending on the number of independent variables combining to have an effect on the dependent variable. A two-way interaction will contain two independent variables (as in our example), a three-way interaction will contain three, etc.

A simple effect refers to the effect of one independent variable at one particular level of the other independent variable. For example, in the data shown in Table 1.1, it looks as though there is a difference between the sexes under the none noise condition but not under the loud or quiet conditions. You could test the differences between the sexes for each noise condition separately, and would then be testing simple effects. Usually you only look at simple effects if you have found a significant interaction.

An example of ANOVA, which is frequently used in research, is provided by Burgess et al. (2003) who measured attachment style between young children

| | Main Effect of Gender: | | Main Effect of Attachment Style: | | | | Interaction between Attachment Style and Gender | | | |

	Male	Female
Inhibition		

	A	B	C
Inhibition			

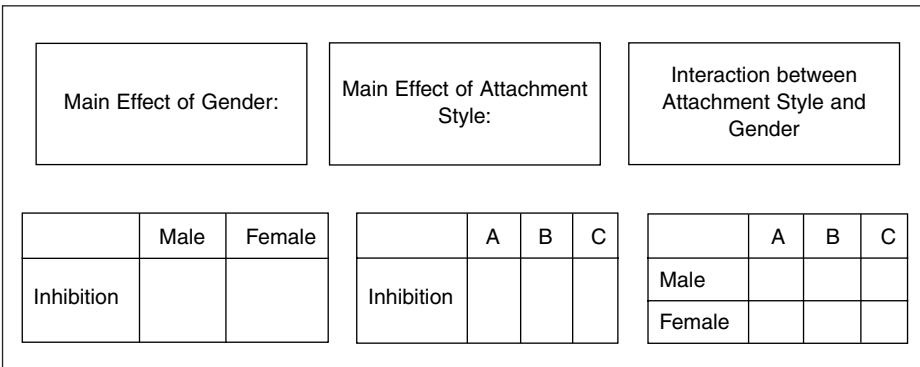|	A	B	C
Male			
Female			

Figure 1.2   The ANOVA considering the effect of attachment style and gender on oral inhibition in a group of children considered in Burgess et al. (2003)

and their parents at 14 months of age and classified the children into three attachment style groups: insecure–avoidant (A), securely attached (B) and insecure–ambivalent (C). ANOVA was used to examine the effects of attachment style on oral inhibition measured at 24 months. Burgess et al. also entered gender into the model as they thought that it may be an additional factor which could influence oral inhibition. (Including gender in the analysis demonstrates a useful aspect of ANOVA: it allows a consideration of extraneous variables, namely uncontrollable variables which may influence the outcome measure.) So this was a 2×3 between-groups ANOVA. The aspects of the data considered by this ANOVA are demonstrated in Figure 1.2.

The authors reported that gender did not have a significant main effect on oral inhibition nor did it interact with attachment style, so the analysis was repeated with gender removed. They reported a trend for a main effect of attachment group on oral inhibition. Since the hypothesis which they made regarding the effect of attachment style on oral inhibition was one tailed (i.e. they had predicted the direction of the relationship), they performed *post hoc* analysis to determine whether the directional nature of the relationship was found even though the main effect was only significant at a trend level. When considering their *post hoc* analysis, the authors reported the mean difference between the groups which proved to be significantly different from one another.

THE GENERAL LINEAR MODEL

The general linear model (GLM) is not a discrete statistical technique itself but the statistical theory which underpins many parametric techniques. The general

aim of methods underpinned by the GLM is to determine whether the independent variable(s) affect or relate to the dependent variable(s). The GLM is summarized by the regression equation which is:

$$Y = c + bX$$

In the equation c is the value of Y when X is zero, and b is the slope of the curve or the amount by which X must be multiplied to give the value of Y.

The equation basically means that the dependent variable, Y, is related to the independent variable X. The fundamental notion that the dependent variable is related to the independent variable also applies to the ANOVA. When you are comparing the means of groups of respondents, as in the example of arithmetic performance under different levels of background noise mentioned earlier, you are actually proposing that the levels of the independent variable affect the dependent variable. So analysis of variance can be seen as a form of regression, and regression, correlation and the comparison of means in the ANOVA share the same underlying concepts. This is why we can say we have a general model. Whichever GLM technique is being used, it is assumed that the relationship between the independent and dependent variables is linear in nature – and so we have the notion of a GLM.

When variables are entered into one of the statistical techniques which use the GLM, it is assumed they have an additive effect, which means they each con-tribute to the prediction of the dependent variable. For example, when three vari-ables are placed into a regression model (or equation) the second variable adds to the predictive value of the first, and the third adds to the predictive value of the second and first combined. (The independent contribution of each variable is assessed by considering the additional variance accounted for or explained by the model when each variable is added.)

When do you need the GLM?

The GLM underlies a number of commonly used statistical techniques. Those which are explained in this book are:

Analysis of variance (ANOVA)
Analysis of covariance (ANCOVA)
Multivariate analysis of variance (MANOVA)
Multivariate analysis of covariance (MANCOVA)
Regression
Multiple regression
Log–linear analysis
Logistic regression

Factor analysis
Structural equation modelling
Survival analysis

ANOVA examines differences between three or more groups of participants or conditions. There is at least one independent variable, which consists of different categories, and a numerical (continuous) dependent variable. (In the example used earlier, mental arithmetic scores formed the dependent variable. Sex of participant and level of background noise formed the independent variables.) The ANOVA determines whether the amount of variance between the groups is greater than the variance within the groups; that is, whether the effect of the independent variable is greater than the effect of individual differences. In an independent measures (between-subjects) design, different participants take part in each condition. ANOVA can be used in repeated measures (within-subjects) studies where the same subjects take part in all of the experimental conditions and when a study has used a mixed design, where some of the independent variables are within subjects and some are between subjects.

ANCOVA is an extension of ANOVA in which a covariate, another relevant variable for which data is available, is classed as an additional independent variable. ANCOVA allows the researcher to 'control' for the effect of the additional variable(s) which may have had an impact on the outcome. For example, Bautmans et al. (2004) used an ANCOVA to examine whether the age of the participants was a potential covariate for the relationship they had found between the health category of elderly respondents and distance walked in 6 minutes. Even after correcting for age using the ANCOVA, the distance walked in 6 minutes decreased with health category so the authors concluded that the existence of pathology rather than increasing age was responsible for decreasing the exercise capacity of the elderly.

MANOVA and MANCOVA are extensions of ANOVA and ANCOVA which are used when there are multiple continuous dependent variables in the analysis.

Regression involves correlations, which are concerned with the relationship between pairs of variables. The type of data determines the correlation which it is appropriate to use. Pearson's product moment coefficient is a parametric correlation which expresses the relationship between two numerical continuous variables. If there is at least one dichotomous variable, the point biserial is appropriate. One would use this, for example, to correlate height with sex. When both variables are dichotomous, such as correlating sex with whether people are or are not smokers, the phi-correlation is required.

In multiple regression there is a single dependent variable predicted from a number of independent variables. Two-group logistic regression is an extension of multiple regression which is used when the dependent variable is dichotomous. Polychotomous logistic regression is necessary if the dependent variable consists of more than two categories.

Log–linear analysis is used to analyse contingency tables or cross-tabulations where more than two variables are included. (If there are just two variables, the chi-square test is often appropriate.)

Logistic regression analysis predicts the values of one dependent variable from one or more independent (predicting) variables when the dependent variable is dichotomous (meaning that it divides the respondents or cases into two exclusive groups such as having a particular illness or not having it). It is used to identify which predictor variables do predict the outcome.

Factor analysis and principal components analysis are a group of techniques which seek to reduce a large number of variables (e.g. questionnaire items or the diagnostic criteria for a psychiatric illness) to a few factors or components.

Structural equation modelling includes a number of statistical techniques and is referred to by a variety of names: causal modelling, causal analysis, simultaneous equation modelling, analysis of covariance structures. It is the most flexible statistical approach using the GLM and is able to deal with multiple independent and dependent variables of categorical or continuous data. Two analytic techniques which use structural equation modelling, path analysis and confirmatory factor analysis, are considered discrete types of this analysis.

Survival analysis is also known (rather perversely) as failure analysis. In medicine, it may be used to examine how long it takes before cancer patients go into remission when receiving one type of treatment versus another. In industry it is used to examine the length of time taken for a component to fail under particular conditions. It can take into account many independent variables such as demographics, environmental conditions, treatment received, etc., depending on the research question.

SUMMARY

In an experiment, participants are randomly allocated to different conditions which form the different levels of the independent variable. In a correlational study, participants are not randomly allocated to different conditions; data on the independent and dependent variables is collected for the respondents and the relationships between them are investigated.

There are two types of data: frequency and numerical. Frequency data involves counting how many people fall into a particular category. Numerical data involves measuring how well people performed by assigning a numerical value to their performance.

ANOVAs examine the differences between group means when there are more than two groups, are able to consider the effect of more than one independent variable and can reveal whether the effect of one of them is influenced by the other.

The general linear model (GLM) is a statistical theory which underpins many parametric analytic techniques. The general aim of methods underpinned by the GLM is to determine whether the independent variable(s) affect or relate to the dependent variable(s).

GLOSSARY

Dependent variable   the output measure of performance.

Frequency data   a count of the number of cases or people in each category.

Homogeneity of variance   when the amount of variance is the same in the different sets of scores of the participants.

Independent variable   those aspects of the respondents, which may be the experimental condition they undergo, which are anticipated will affect the dependent variable (output measure).

Interaction   when the effect of one independent variable is influenced by another independent variable.

Levels of an independent variable   the number of categories into which an independent variable has been divided.

Numerical data   obtained by measuring how well people performed by assigning a numerical value to their performance.

REFERENCES

Bautmans, I., Lambert, M. and Mets, T. (2004). The six-minute walk test in community dwelling elderly: influence of health status. *BMC Geriatrics*, 23, 6.
Burgess, K.B., Marshall, P.J., Rubin, K.H. and Fox, N.A. (2003). Infant attachment and temperament as predictors of subsequent externalizing problems and cardiac physiology. *Journal of Child Psychology and Psychiatry*, 44, 819–831.

FURTHER READING

Harris, R.J. (2001). *A Primer of Multivariate Statistics*. London: Lawrence Erlbaum.
Horton, R.L. (1978). *The General Linear Model: Data Analysis in the Social and Behavioral Sciences*. New York: McGraw-Hill.
Jamieson, J. (2004). Analysis of covariance (ANCOVA) with difference scores. *International Journal of Psychophysiology*, 52, 277–282.

Thompson, B. (1984). *Canonical Correlation Analysis: Uses and Interpretation.* Beverly Hills, CA: Sage.

Timm, N.H. (2002). *Applied Multivariate Analysis.* New York: Springer.

INTERNET SOURCES

http://www.statsoftinc.com/textbook/stglm.html

http://ibgwww.colorado.edu/~carey/p7291dir/handouts/glmtheory.pdf

http://name.math.univ-rennes1.fr/bernard.delyon/textbook/stglm.html