

Chapter 1

Statistics? I Thought This Was Psychology!

In This Chapter

- ▶ Understanding variables
 - ▶ Introducing SPSS
 - ▶ Outlining descriptive and inferential statistics
 - ▶ Differentiating between parametric and non-parametric statistics
 - ▶ Explaining research designs
-

When we tell our initially fresh-faced and enthusiastic first year students that statistics is a substantial component of their course approximately half of them are genuinely shocked. ‘We came to study psychology, not statistics’, they shout. Presumably they thought they would be spending the next three years ordering troubled individuals to ‘lie down on the couch and tell me about your mother’. We tell them there is no point running for the exits as they will quickly learn that statistics is part of all undergraduate psychology courses, and that if they plan to undertake post-graduate studies or work in this area they will be using these techniques for a long time to come (besides, we were expecting this reaction and have locked the exits). Then we hear the cry ‘But I’m not a mathematician. I am interested in people and behaviour’. We don’t expect students to be mathematicians. If you have a quick scan through this book you won’t be confronted with pages of scary looking equations. These days we use computer-based software packages such as SPSS to do all the complex calculations for us. We tell them that psychology is a scientific discipline. If they want to learn about people they have to objectively collect information, summarise it and analyse it. Summarising and analysing allows you to interpret the information and give it meaning in terms of theories and real world problems. Summarising and analysing information is statistics; it is a fundamental and integrated component of psychology.

The aim of this chapter is to give you a roadmap of the main statistical concepts you will encounter during your undergraduate psychology studies and to signpost you to relevant chapters on topics where you can learn how to become a statistics superhero (or at least scrape by).

Know Your Variables

All quantitative research in psychology involves collecting information (called *data*) that can be represented by numbers. For example, levels of depression can be represented by depression scores obtained from a questionnaire, or a person's gender can be represented by a number (1 for male and 2 for female). The characteristics you are measuring are known as *variables* because they vary! They can vary over time within the same person (depression scores can vary over a person's life time) or vary between different individuals (individuals can be classified as male or female, but once a person is classified this variable doesn't tend to change!).

Several names and properties exist, associated with variables in any data set, which you must become familiar with. Variables can be continuous or discrete, have different levels of measurement and can be independent or dependent. We cover all this information in Chapter 2. Initially these terms may seem a little bamboozling, but it is important you ensure you have a good understanding of them, as they dictate the statistical analyses that are available and appropriate for your data. For example, it helps to report a mean depression score of 32.4 for a particular group of participants, but a mean gender score of 1.6 for the same group doesn't much make sense (we discuss the mean in Chapter 4)!

Variables can be classified as *discrete*, where you specify discrete categories (for example, male and female), or *continuous*, where scores can lie anywhere along a continuum (for example, depression scores may lie anywhere between 0 and 63 if measured by the Beck Depression Inventory).

Variables also differ in their measurement properties. Four levels of measurement exist:

- ✓ **Nominal:** This contains the least amount of information of the levels. At the nominal level, a numerical value is applied arbitrarily. Gender is an example of *nominal* level of measurement (for example, 1 for male and 2 for female), and it makes no sense to say one is greater or less than the other.
- ✓ **Ordinal:** Rankings on a class test are an example of an *ordinal* level of measurement; we can order participants from the highest to the lowest score but we don't know how much better the first person did compared to the second person (it could be 1 mark or it could be 20 marks!).
- ✓ **Interval:** IQ scores are measured at the *interval* level, which means we can order the scores but the difference between each point is equal. That is, the difference between 95 and 100 is the same as the difference between 115 and 120.

✓ **Ratio:** In a ratio level of measurement, the scores can be ordered, the difference between each point on the scale is equal and the scale also has a true absolute zero. Weight, for example, is measured at the *ratio* level; having a true zero means a weight of zero signifies an absence of any weight and it also allows you to make proportional statements, such as '10 kg is half the weight of 20 kg'.

You will also need to classify the variables in your data as *independent* or *dependent* and the classification will depend on the research question you are asking. For example, if you are investigating the difference in depression scores between males and females, the independent variable is gender (this is the variable you think that is predicting a change), and depression scores are the dependent variable (this is the outcome variable where the scores depend on the independent variable).

What is SPSS?

The initials SPSS stand for Statistical Package for the Social Sciences and it provides you with a program capable of storing, manipulating and analysing your data. This book assumes you will be using SPSS to analyse your data. SPSS is probably the most commonly used statistics package in the social sciences, but of course other similar packages exist as well as those designed to conduct more specialised analysis.

There are three main 'views' or windows you will be using in SPSS. The first is the 'variable view' and this is where you label and code the variables you are working with, for example, if you wanted to specify the two variables 'gender' and 'depression'. The 'data view' is the spreadsheet where you enter all your data. The normal format when entering data is that each column will represent a variable (for example, gender or depression) and each row will represent one individual or participant. Therefore if you collected and entered information on the gender and depression scores of 10 participants you would have 2 columns and 10 rows in your SPSS data view. SPSS allows you to enter numeric data, string data (which is non-numeric information such as names) and also assign codes (for example, 1 for male and 2 for female).



Unlike other programs you may have used (for example, Microsoft Office Excel) cells do not contain formulae or equations.

Once your data is entered, SPSS allows you to run a wide variety of analyses by using drop down menus. There are literally hundreds of different analyses and options you can choose; in this book we will only explain the statistical procedures necessary for your course. When you have selected the analyses you want to conduct, your results will appear in a separate 'output window'; your job then is to read and interpret the relevant information.



In addition to using the pull-down menus you can also program SPSS by using a simple syntax language. This can be useful if you need to repeat the same analyses on many different data sets, but explaining how to use it is beyond the scope of an introductory text.

SPSS was first released in 1968 and has been through many versions and upgrades; at the time of writing this chapter, the most recent version was SPSS 20.0 which was released in August 2011. Between 2009 and 2010 SPSS briefly was known as Predictive Analytics SoftWare and could be found on your computer under the name *PASW*. In 2010 it was purchased by IBM and now appears in your computer's menu under the name *IBM SPSS statistics* (and no, we don't know why the last 'statistics' is necessary either!).

Descriptive Statistics

When you collect your data you need to communicate your findings to other people (tutor, boss, colleagues or whoever it may be). Let's imagine you collect data from 100 people on their levels of coulrophobia (fear of clowns); if you simply produce a list of 100 scores in SPSS this won't be very useful or easy to comprehend for your audience. Instead you need a way to describe your data set in a concise and repeatable format. The standard way to do this is to present two pieces of information, a measure of central tendency and a measure of dispersion.

Central tendency

There are several different types of central tendency, but they all attempt to give a single number that represents your variable. The most common measure is sometimes known as *average* but it is more correctly called the *arithmetic mean*, and you are probably familiar with it. To obtain the mean you simply add all the scores on a variable and divide by the number of participants or cases you had. One of the strengths of the mean as a measure of central tendency is that it represents all your data; however, this means it has a weakness in that it can be influenced by extreme scores. The mean isn't always an appropriate number to represent your data. The *median* (the middle value when the scores are ranked) is more appropriate when your variable is measured at the ordinal level of measurement and the *mode* (the most frequently occurring value) is appropriate when your variable is measured at the nominal level. Measures of central tendency are covered in Chapter 4.

Dispersion

There are also several measures of dispersion, and each aims to give a single number that represents the spread or variability of your variable. The larger

the dispersion value for your variable, the larger its variability (participants vary on the scores they obtain), whereas a small value of dispersion indicates that the scores vary less (participants tend to score similarly). The most common measure of is the standard deviation, an estimate of the average variability of spread of your variable. Chapter 5 describes the standard deviation along with the other important measures of dispersion which include the variance, range and interquartile range.

Graphs

Another way of displaying your data is to provide a visual representation in the form of a graph. Graphs are important for another reason; the type of statistical analysis you can conduct with variables will depend on the distribution of your variables, which you will need to assess by using graphs. Chapter 6 outlines the common types of graphs used in psychology (the histogram, bar chart, cumulative frequency plot, and box and whisker plot) and how to generate each of them in SPSS.

Standardised scores

Imagine you measured a friend's extraversion level with the Revised NEO Personality Inventory and told them they obtained a score of 164; it is likely they will want to know how this score compares to other people's scores. Is it high or low? They also might want to know how it compares to the psychoticism score of 34 you gave them last week from the Eysenck Personality Questionnaire. Simply reporting raw scores often isn't that informative. You need to be able to compare these scores to other people's scores and importantly you'll need to compare scores that are measured on different scales. The good news is that it is quite easy to convert your raw score into a standardised score, which means it is possible to make these comparisons. A standardised score is measured in terms of standard deviations (so it can be compared against standardised scores from different variables) and allows instant comparisons to the mean score on a variable (which means you can tell if an individual's score is greater or less than the mean). We cover standardisation in more detail in Chapter 10.

Inferential Statistics

Descriptive statistics are useful in illustrating the properties of your sample (that is, the participants you have collected data from), but the majority of the time you will be more interested in the properties of the population (that is, all possible participants of interest). For example, if you are interested in differences in attitudes to sectarianism between boys and girls enrolled in

schools in Northern Ireland, then your population is all Northern Irish school children. As it is unrealistic to recruit all the children in Northern Ireland (in terms of time, money and consent) you would measure sectarianism in a small subset or *sample* of the children (we examine the differences between samples and populations in Chapter 7).

Inferential statistics allows you to make inferences from your sample about the larger population. For example, if you found a difference in sectarianism between boys and girls in your sample you could infer this difference to all schoolchildren in Northern Ireland (inferential statistics are explained in Chapter 7). You can't be completely sure this difference exists in the population as you haven't tested every child in the population, but you should be 95 per cent confident that the difference you found in your sample exists in the population (under certain conditions), and this is what an inferential statistic assesses (this is discussed in more detail in Chapter 7).

The inferential statistic you conduct will tell you about the probability of your result occurring in the population (that is, whether the difference in your sample really exists in the population, or whether you obtained this result by chance) but it does not tell you anything about the size of the difference. For instance, you may find that males are more likely to show sectarian attitudes than females, but this isn't very interesting if the difference between the attitudes is really tiny. *Effect sizes* indicate the strength of the relationship between your variables (we cover effect sizes in Chapter 11) and should always be reported in conjunction with the probability level associated with any inferential statistic.

Hypotheses

Before you commence any study it is important to have a hypothesis or a specific testable statement that reflects the aim of your study. In the example above you would specify the hypothesis: 'there is no difference between levels of sectarianism between boys and girls enrolled in Northern Ireland schools.' We outline hypothesis testing in Chapter 8 and explain why we always start with the assumption that your data demonstrates no effect, difference or relationship.

Parametric and non-parametric variables



When you are addressing a hypothesis there are two main types of statistical analysis you can conduct; a *parametric* test or the *non-parametric* equivalent. Parametric statistics assume that the data approximates a certain distribution, such as the normal distribution explained in Chapter 9. This allows us to make inferences which make these types of statistics powerful (see Chapter 11 for a discussion of power) and capable of producing accurate results.

However, because parametric statistics are based on certain assumptions, you must check your data to ensure it adheres to these assumptions (we explain how to do this for each individual statistic in the book). Failure to check the assumptions means you run the risk of performing inappropriate analyses which means your results, and therefore conclusions, may be incorrect.

By comparison, non-parametric statistics make fewer assumptions about your data, which mean they can be used to analyse a more diverse range of data. Non-parametric tests tend to be less powerful than their parametric equivalents, so you should always attempt to use the parametric version unless the data violates the assumptions of that test.

Research Designs

The type of statistical analyses you should conduct depends on several things, but to decide on the correct ‘family’ of tests, you must first decide on the design of your study. The choice of design is influenced by the question you want answered or your hypothesis. Research design can be broadly classified into *correlational design* and *experimental design*.

Correlational design

Correlational design is when you are interested in the relationships or associations between two or more variables. Correlational design is distinguished from experimental design as there is no attempt to manipulate the variables; instead you are investigating existing relationships between the variables. For example, you may be conducting a study to look at the relationship between the use of illegal recreational drugs and visual hallucinations; in this case you need to recruit participants with varying levels of existing drug use and measure their experience of hallucinations. The ethics panel of your department may have some serious misgivings if you try to conduct an experimental study in this area, manipulating your variables by handing out various amounts of illegal drugs or attempting to induce hallucinations in your participants. Part III of the book deals with inferential statistics that assess relationships or associations between variables which normally relate to correlational designs (please note our use of normally! There are always exceptions!).

Correlation coefficients provide you with a number that represents the strength of a linear (straight line) relationship and also its direction; if there is a strong positive correlation, high scores on one variable tend to be correlated with high scores on the other variable (for example, participants who report high drug use tend to report high levels of hallucinations) and a strong negative correlation indicates that high scores on one variable tend to be correlated with low scores on the other variable (for example participants

who report high drug use tend to report low levels of hallucinations). There are several different types of correlation coefficients and you will need to decide which one is appropriate for your data. We explain how to do this, as well as obtaining and interpreting correlation coefficients, in Chapter 12.

Regression takes the concept of looking at relationships further, as it allows you to test whether one or more variables can predict an outcome variable or criterion. For example, you could use regression to test whether use of illegal drugs, neuroticism and age can predict visual hallucinations. In comparison to the correlation, this technique gives you more information, such as which variables are significant predictors of hallucinations and the relative strength of each predicting variable. Linear regression is covered in Chapter 13 of this book.

There may be times where you want to examine the association between two discrete variables, for example, if whether or not someone ever took illegal drugs was associated with whether or not they ever experienced a visual hallucination. This type of data is examined in a specific way and in Chapter 14 we outline how to use contingency tables and inferential statistics (for example, the chi-square test and the McNemar test) to analyse discrete data.

Experimental design

Experimental designs differ from correlational designs as they may involve manipulating the independent variable (see Chapter 2 for a discussion of independent variables). Correlational studies focus on the relationship between existing variables, whereas in experimental designs a variable is changed (directly or indirectly) and you assess whether this has an effect on your outcome variable. For example, you may hypothesise that ergophobia (fear of work) in psychology students increases throughout their courses. There are two experimental designs you could use to test this hypothesis.

In the first design, you could test levels of ergophobia in students at different stages of their courses. This example, where you are comparing separate independent groups of people, is known as *independent groups design* (the statistical analyses relevant for this design is covered in Part IV of the book). The second design involves measuring the ergophobia levels of all the students in their first year and then measuring the same participants several times throughout their course. This type of study, where you are interested in changes within the same group of participants, is known as *repeated measures design* (the statistical analyses relevant for this design is covered in Part V of the book).

Independent groups design

When you employ an independent groups design you are looking for differences on a variable between separate groups of people. If you are investigating

differences between two separate groups on a variable (for example, first year and second year psychology students on ergophobia levels) it is likely the two sets of scores will differ to some extent. Using an inferential statistic estimates whether this difference is likely to exist in the population or if you could have obtained the result by chance (assuming certain conditions). In this scenario you can employ either the parametric independent t -test or non-parametric Mann–Whitney test. We explain these tests in Chapter 15.

If you want to investigate the difference between more than two groups (for example, first, second and third year psychology students on ergophobia levels) the parametric between-groups ANOVA is most appropriate. (The non-parametric equivalent is the Kruskal–Wallis test). These analyses are covered in Chapter 16. If you find there is a statistically significant difference in ergophobia levels between the first, second and third year psychology students you will need to find out where these differences exist (is it between first years and second years, between first years and third years, and so on). The post-hoc tests and planned comparisons you need to do this are covered in Chapter 17. You could employ a slightly more sophisticated design where you examine the effect of two independent variables on a dependent variable. For example, does year of study and gender affect ergophobia levels? The benefit of this design is that you can examine the interaction between your two independent variables, for example, females' ergophobia levels may remain constant, but male scores increase in the three year groups. This two-way between-groups ANOVA design is covered in Chapter 16.

Repeated measures design

When you employ a repeated measures design you are looking for differences on a variable within the same group of people. For example, you could measure ergophobia levels when the students first start their psychology course and then 12 months later test the same group to see if the scores have changed. If you have tested your participants twice (that is, the independent variable has two levels), you can use the paired t -test or non-parametric Wilcoxon test to see if the differences in scores are statistically significant. We discuss these tests in Chapter 18. If you have tested the same group of participants more than twice the parametric within-groups ANOVA or the non-parametric Friedman test is most appropriate. These analyses are covered in Chapter 19. If you find there is a statistically significant difference in ergophobia levels between the testing sessions, you will need to find out where these differences exist (the students' first year or second year, or did the change occur between the students' second year and third year, and so on). The post-hoc tests and planned comparisons you need to do this are covered in Chapter 20. If you want to construct a more sophisticated design you can examine the effect of two repeated measures variables on a dependent variable (this is known as a two-way within-groups ANOVA and is covered in Chapter 19) or one repeated-measures and one independent-groups design on a dependent variable (this is a mixed ANOVA design and is addressed in Chapter 21).

Getting Started

The critical stage of any research study is always the start. It is important to specify a hypothesis that is testable and actually addresses the question you are interested in (see Chapter 8 for more on hypotheses); your hypothesis must be informed by theory and previous research. At this early stage you also need to consider how you will analyse the data by deciding on the appropriate statistic; this will help you decide how to measure your variables, so you don't find yourself unable to address your hypothesis because you cannot perform the analysis you want with the data you have collected. Deciding on the appropriate statistical analysis also allows you to calculate the sample size you will need (see Chapter 11). If you do not recruit enough participants you are unlikely to discover a significant effect in your data even if one exists in the population, and your efforts will be a waste of time. When you are preparing your SPSS file, take time to label your data and assign values that are easy to read and will make sense when you re-visit them months later. It helps to keep a paper record of the labels and values you used so you can refer to them when you are running analyses. Remember to save your SPSS file regularly when you are entering your data. When the fateful day arrives and you begin to analyse your data, take a deep breath and relax. Give yourself plenty of time, take notes as you go along, save your appropriately named output files and allow yourself to make mistakes. SPSS will run statistics almost instantly so if you make a mistake (and you will) you can simply start again.

We advise you that the best time to consult a statistical advisor is when you are designing your study. They will be able to offer you advice on the type of data you should collect, the analyses you will have to conduct and the sample size you will require. Asking for help after the data has been collected may be too late!