10 Moderation, mediation and more regression

FIGURE 10.1 My 10th birthday. (From left to right) My brother Paul (who still hides behind cakes rather than have his photo taken), Paul Spreckley, Alan Palsey, Clair Sparks and me

10.1. What will this chapter tell me? ①

Having successfully slayed audiences at holiday camps around the country, my next step towards global domination was my primary school. I had learnt another Chuck Berry song ('Johnny B. Goode'), but also broadened my repertoire to include songs by other artists (I have a feeling 'Over the edge' by Status Quo was one of them).[1] Needless to say, when the opportunity came to play at a school assembly I jumped at it. The headmaster tried to have me banned,[2] but the show went on. It was a huge success (I want to reiterate my earlier

[1] This would have been about 1982, so just before they became the most laughably bad band on the planet. Some would argue that they were *always* the most laughably bad band on the planet, but they were the first band that I called my favourite band.

[2] Seriously! Can you imagine a headmaster banning a 10-year-old from assembly? By this time I had an electric guitar and he used to play hymns on an acoustic guitar; I can assume only that he somehow lost all perspective on the situation and decided that a 10-year-old blasting out some Quo in a squeaky little voice was subversive or something.

point that 10-year-olds are very easily impressed). My classmates carried me around the playground on their shoulders. I was a hero. Around this time I had a childhood sweetheart called Clair Sparks. Actually, we had been sweethearts since before my newfound rock legend status. I don't think the guitar playing and singing impressed her much, but she rode a motorbike (really, a little child's one) which impressed *me* quite a lot; I was utterly convinced that we would one day get married and live happily ever after. I was utterly convinced, that is, until she ran off with Simon Hudson. Being 10, she probably literally did run off with him – across the playground. I remember telling my parents and them asking me how I felt about it. I told them I was being philosophical about it. I probably didn't know what philosophical meant at the age of 10, but I knew that it was the sort of thing you said if you were pretending not to be bothered about being dumped.

If I hadn't been philosophical, I might have wanted to look at what had lowered Clair's relationship satisfaction. We've seen in previous chapters that we could predict things like relationship satisfaction using regression. Perhaps it's predicted from your partner's love of rock bands like Status Quo (I don't recall Clair liking that sort of thing). However, life is usually more complicated than this; for example, your partner's love of rock music probably depends on your own love of rock music. For example, if you both like rock music then your love of the same music might have an additive effect, giving you huge relationship satisfaction (*moderation*), or perhaps the relationship between your partner's love of rock and your own relationship satisfaction can be explained by your own music tastes (*mediation*). In the previous chapter we also saw that regression could be done with a dichotomous predictor (e.g., rock fan or not) but what if you wanted to categorize musical taste into several categories (rock, hip-hop, R & B etc.)? Surely you can't use multiple categories as a predictor in regression? This chapter extends what we know about regression to these more complicated scenarios. First we look at two common regression-based models – moderation and mediation – before expanding what we already know about categorical predictors.

10.2. Installing custom dialog boxes in SPSS ②

Although you can do both moderation and mediation analysis in SPSS manually, it's a bit of a faff. It will require you to create new variables using the *compute* command, and in the case of mediation analysis it will limit what you can do considerably. By far the best way to tackle moderation and mediation is to use the *PROCESS* command. This is not part of SPSS; it exists only because Andrew Hayes and his colleague Kristopher Preacher have spent an enormous amount of time writing a range of tools for doing moderation and mediation analyses (e.g., Hayes & Matthes, 2009; Preacher & Hayes, 2004, 2008a). These tools were previously available only through syntax, and for inexperienced users were a bit scary and fiddly. Andrew Hayes wrote the *PROCESS* custom dialog box (Hayes, 2012) to wrap the Preacher and Hayes mediation and moderation tools in a convenient menu and dialog box interface. It's pretty much the best thing to happen to moderation and mediation analysis in a long time. While using these tools, I strongly suggest you spare a thought of gratitude that there are people like Hayes and Preacher in the world who invest their spare time doing cool stuff like this that makes it possible for you to analyse your data without having a nervous breakdown. Even if you think you are having a nervous breakdown, trust me it's not as big as the one you'd be having if *PROCESS* didn't exist.

The *PROCESS* tool is what's known as a custom dialog box. SPSS includes the ability to add your own menus and dialog boxes, which means that you can write your own functions using syntax, but then create a custom menu and dialog box for yourself so that you can access the syntax through a nice point and click menu. Of course, most of us will never use this feature, but Andrew Hayes has. Essentially, he provides a file (**process.spd**) that you download, which installs a new menu into the Analyze Regression ▸ menu.

From this menu you access a dialog box that can be used to do moderation and mediation analysis.

You install *PROCESS* in three easy steps, which are illustrated in Figure 10.2 (MacOS users can ignore step 2):

1. *Download the install file*: Download the file **process.spd** from Andrew Hayes' website: http://www.afhayes.com/spss-sas-and-mplus-macros-and-code.html. Save this file onto your computer.
2. *Start SPSS as an administrator*: To install the tool in Windows, you need to start IBM SPSS as an administrator. To do this, make sure that SPSS isn't already running, and then click on the start menu ( ). Select ▸ All Programs , which will display a list of programs installed on your machine. Within that list, there should be a folder called *IBM SPSS Statistics*. Select that folder to display its contents. You should see this icon within that folder: IBM SPSS Statistics 20 (don't be worried if the number is different from 20, it just refers to the version of SPSS that you have installed). Click on this icon with the *right mouse button* to activate the menu in Figure 10.2. Within this menu select (you're back to using the left mouse button now) Run as administrator. This action opens SPSS but allows it to make changes to your computer. A dialog box will appear that asks you whether you want to let SPSS make changes to your computer and you should select Yes.
3. Once SPSS has loaded select Utilities Custom Dialogs ▸ Install Custom Dialog..., which will open a standard dialog box for opening files (Figure 10.2). Locate the file **process.spd**, select it, and click on Open. This will install the *PROCESS* menu and dialog boxes into SPSS. If you get an error message, the most likely explanation is that you haven't opened SPSS as an administrator (see step 2). Once the installation is complete you'll find that the *PROCESS* menu has been added to the existing Analyze Regression ▸ menu (Figure 10.3).

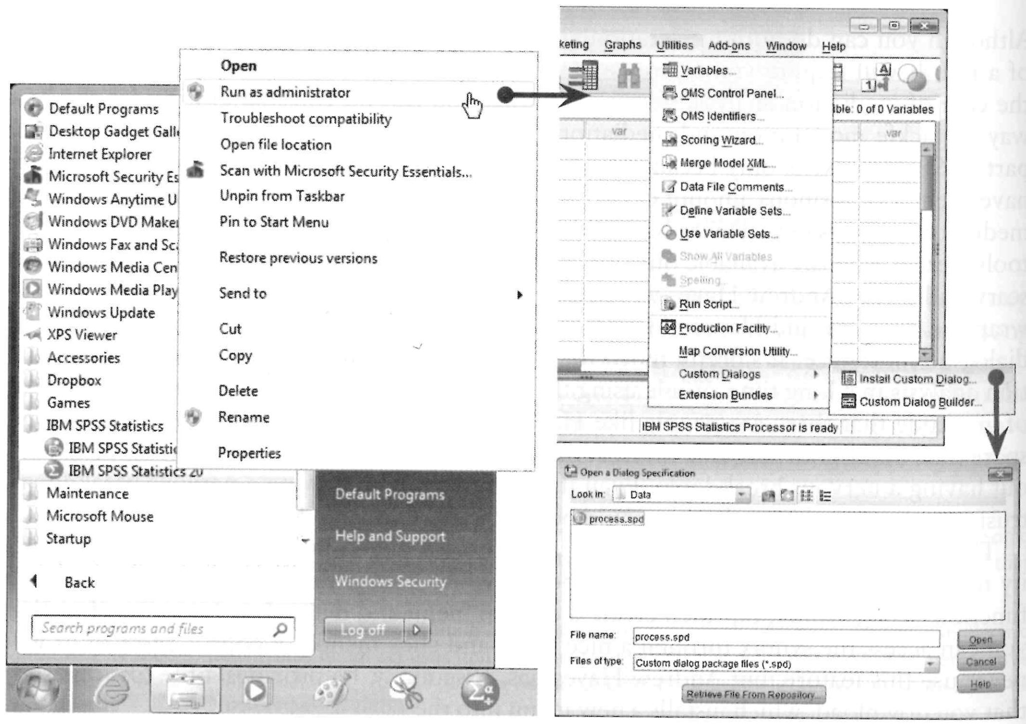

FIGURE 10.2 Installing the *PROCESS* menu

10.3. Moderation: interactions in regression ③

10.3.1. The conceptual model ③

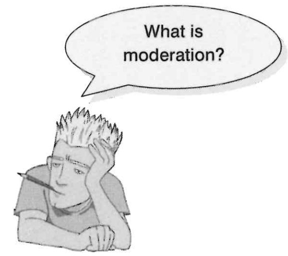

So far we have looked at individual predictors in the linear model. However, it is possible for a statistical model to include the combined effect of two or more predictor variables on an outcome. The combined effect of two variables on another is known conceptually as **moderation**, and in statistical terms as an **interaction effect**. We'll start with the conceptual and we'll use an example of whether violent video games make people antisocial. Video games are among the favourite online activities for young people: two-thirds of 5–16-year-olds have their own video games console, and 88% of boys aged 8–15 own at least one games console (Ofcom (Office of Communications), 2008). Although playing violent video games can enhance visuospatial acuity, visual memory, probabilistic inference, and mental rotation (Feng, Spence, & Pratt, 2007; Green & Bavelier, 2007; Green, Pouget, & Bavelier, 2010; Mishra, Zinni, Bavelier, & Hillyard, 2011), compared to games such as Tetris, these games have also been linked to increased aggression in youths (Anderson & Bushman, 2001). Another predictor of aggression and conduct problems is callous-unemotional traits such as lack of guilt, lack of empathy, and callous use of others for personal gain (Rowe, Costello, Angold, Copeland, & Maughan, 2010). Imagine a scientist wanted to look at the relationship between playing violent video games such as Grand Theft Auto, MadWorld and Manhunt and aggression. She gathered data from 442 youths (**Video Games.sav**). She measured their aggressive behaviour (**Aggression**), callous unemotional traits (**CaUnTs**), and the number of hours per week they play video games (**Vid_Games**).

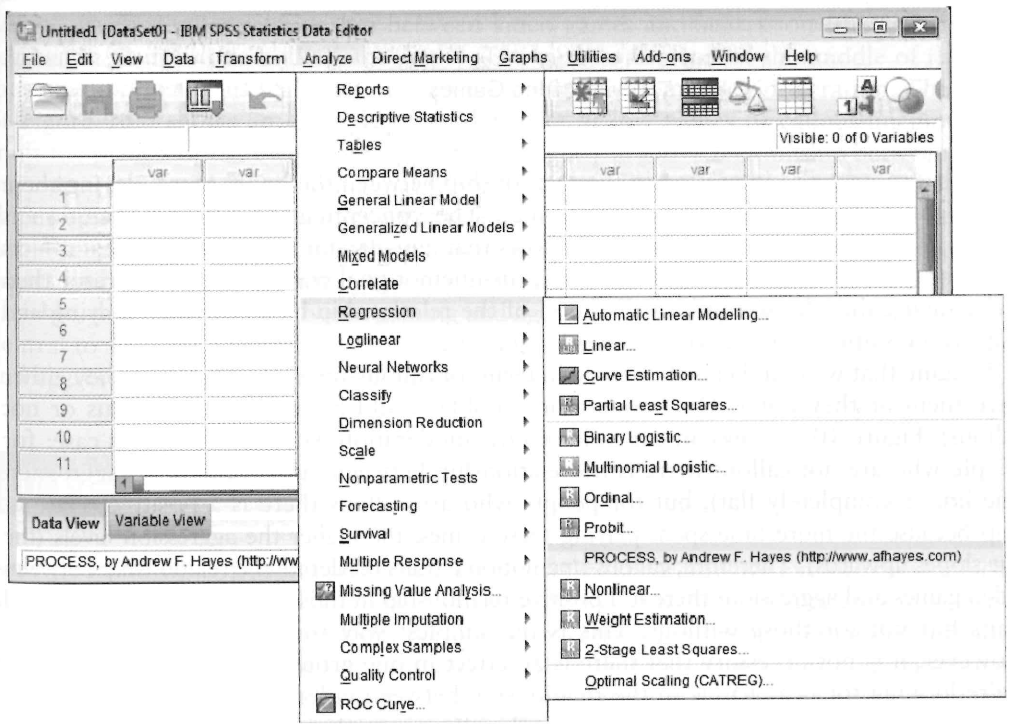

FIGURE 10.3 After installation, the *PROCESS* menu appears as part of the existing *Regression* menu

FIGURE 10.4 Diagram of the *conceptual* moderation model

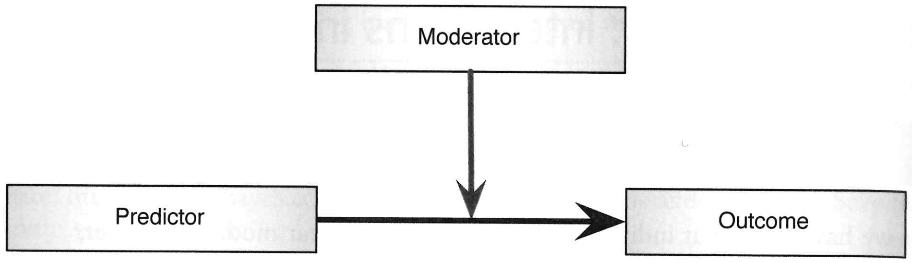

FIGURE 10.5 A categorical moderator (callous traits)

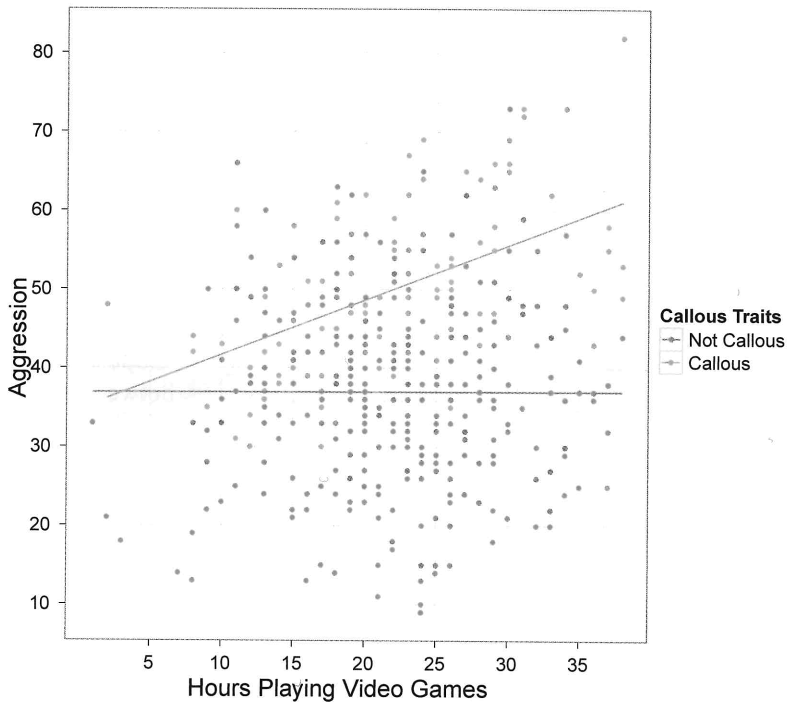

Let's assume we're interested in the relationship between the hours spent playing these games (predictor) and aggression (outcome). The conceptual model of moderation is shown in Figure 10.4, and this diagram shows that a **moderator** variable is one that affects the relationship between two others. If callous-unemotional traits were a moderator then we're saying that the strength or direction of the relationship between game playing and aggression is affected by callous-unemotional traits.

Imagine that we could classify people in terms of callous-unemotional traits: they either have them or they don't. Our moderator variable would be categorical (callous or not callous). Figure 10.5 shows an example of how moderation would work in this case: for people who are not callous there is no relationship between video games and aggression (the line is completely flat), but for people who are callous there is a positive relationship because the more time spent playing these games, the higher the aggression levels (the line slopes upwards). Therefore, callous-unemotional traits moderate the relationship between video games and aggression: there is a positive relationship in those with callous-unemotional traits but not for those without. This is the simplest way to think about moderation. However, it is not necessary that there is an effect in one group but not in the other, all we're looking for is a change in the relationship between video games and aggression in the two callousness groups. It could be that the effect is weakened or changes direction.

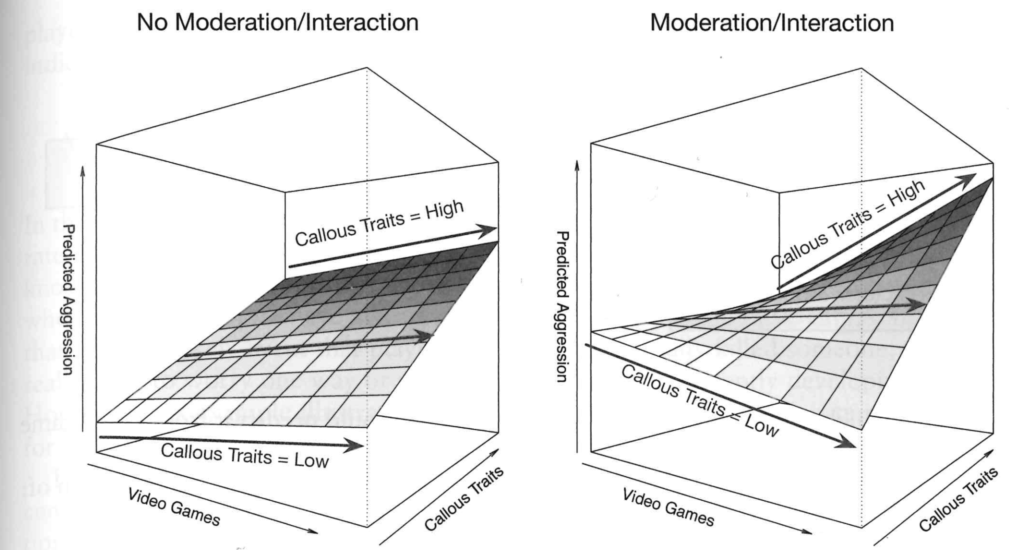

FIGURE 10.6 A continuous moderator (callous traits)

If we measure the moderator variable along a continuum it becomes a bit trickier to visualize, but the basic interpretation stays the same. Figure 10.6 shows two graphs that display the relationships between the time spent playing video games, aggression and callous-unemotional traits (measured along a continuum rather than as two groups). We're still interested in how the relationship between video games and aggression changes as a function of callous-unemotional traits. We can do this by comparing the slope of the regression plane for time spent gaming at low and high values of callous traits. To help you I have added blue arrows that show the relationship between video games and aggression. In the left of the diagram you can see that at the low end of the callous-unemotional traits scale, there is a slight positive relationship between playing video games and aggression (as time playing games increases so does aggression). At the high end of the callous-unemotional traits scale, we see a very similar relationship between video games and aggression (the ends of the regression planes slope at the same angle). The same is also true at the middle of the callous-unemotional traits scale. This is a case of no interaction or no moderation. The right of Figure 10.6 shows an example of moderation: at low values of callous-unemotional traits the plane slopes downwards, indicating a slightly negative relationship between playing video games and aggression, but at the high end of callous-unemotional traits the plane slopes upwards, indicating a strong positive relationship between gaming and aggression. At the midpoint of the callous-unemotional traits scale, the relationship between video games and aggression is relatively flat. So, as we move along the callous-unemotional traits variable, the relationship between gaming and aggression changes from slightly negative to neutral to strongly positive. We can say that the relationship between violent video games and aggression is moderated by callous-unemotional traits.

10.3.2. The statistical model ②

Now we know what moderation is conceptually, let's look at how we explore these effects within a statistical model. Figure 10.7 shows how we conceptualize moderation statistically: we predict the outcome from the predictor variable, the proposed moderator, and the interaction of the two. It is the interaction effect that tells us whether moderation has occurred, but *we must include the predictor and moderator as well for the interaction term to be valid*. This point is very important. In our example, then, we'd be looking at doing a

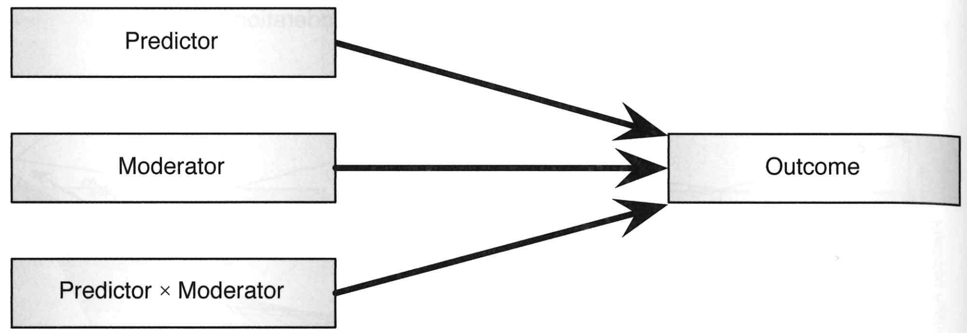

FIGURE 10.7 Diagram of the *statistical* moderation model

regression in which aggression was the outcome, and we would predict it from video game playing, callous-unemotional traits and their interaction.

All of the general linear models we've considered in this book take the general form of:

$$\text{outcome}_i = (\text{model}) + \text{error}_i$$

When we encountered multiple regression in Chapter 8 we saw that this model was written as (see equation (8.6)):

$$Y_i = (b_0 + b_1X_{1i} + b_2X_{2i} + \cdots + b_nX_{ni}) + \varepsilon_i$$

Therefore, our basic regression model for this example would be:

$$\text{Aggression}_i = (b_0 + b_1\text{Gaming}_i + b_2\text{Callous}_i) + \varepsilon_i$$

However, to test for moderation we need to consider the interaction between gaming and callous-unemotional traits. If we want to include this term too, then we have seen before that we can extend the linear model to include extra terms, and each time we do we assign them a parameter (b). A model that tests for moderation, therefore, is as follows (first expressed generally and then in terms of this specific example):

$$Y_i = (b_0 + b_1A_i + b_2B_i + b_3AB_i) + \varepsilon_i$$
$$\text{Aggression}_i = (b_0 + b_1\text{Gaming}_i + b_2\text{Callous}_i + b_3\text{Interaction}_i) + \varepsilon_i \tag{10.1}$$

10.3.3. Centring variables ②

When an interaction term is included in the model the b parameters have a specific meaning: for the individual predictors they represent the regression of the outcome on that predictor when the other predictor is zero. So, in equation (10.1), b_1 represents the relationship between aggression and gaming when callous traits are zero, and b_2 represents the relationship between aggression and callous traits when someone spends zero hours gaming per week. In our particular example this interpretation isn't problematic because zero is a meaningful score for both predictors: it's plausible that a child spends no hours playing video games, and it is plausible that a child gets a score of 0 on the continuum of callous-unemotional traits. However, there are often situations where it makes no sense for a predictor to have a score of zero. Imagine that rather than measuring how much a child played violent video games we'd measured their heart rate while playing the games as an indicator of their physiological reactivity to them:

$$\text{Aggression}_i = (b_0 + b_1\text{Heart Rate}_i + b_2\text{Callous}_i + b_3\text{Interaction}_i) + \varepsilon_i$$

In this model b_2 is the regression of aggression on callous traits when someone has a heart rate of zero while playing the games. This b makes no sense unless we're interested in knowing something about the relationship between callous traits and aggression in youths who die (and therefore have a heart rate of zero) while playing these games. It's fair to say that in the unlikely event that playing a video game actually killed someone, we wouldn't really have to worry one way or another about them subsequently developing aggression. Hopefully this example illustrates that the presence of the interaction term makes the bs for the main predictors uninterpretable in many situations.

For this reason, it is common to transform the predictors using **grand mean centring**. **Centring** refers to the process of transforming a variable into deviations around a fixed point. This fixed point can be any value that you choose, but typically it's the grand mean. When we calculated z-scores in Chapter 1 we used grand mean centring because the first step was to take each score and subtract from it the mean of all scores. This is grand mean centring. Like z-scores, the subsequent scores are centred on zero, but unlike z-scores we don't care about expressing the centred scores as standard deviations.[3] Therefore, grand mean centring for a given variable is achieved by taking each score and subtracting from it the mean of all scores (for that variable).

Centring the predictors has no effect on the b for highest-order predictor, but will affect the bs for the lower-order predictors. 'Highest-order' and 'lower-order' refer to how many variables are involved: so the gaming × callous traits interaction is a higher-order effect than the effect of gaming alone because it involves two variables rather than one. So, in our model (equation (10.1)), whether or not we centre the predictors will have no effect on b_3 (the parameter for the interaction) but it will change the values of b_1 and b_2 (the parameters for gaming and callous traits). As we have seen, if we don't centre the gaming and callous variables, then the bs represent the effect of the predictor when the other predictor is zero. However, if we centre the gaming and callous variables then the bs represent the effect of the predictor when the other predictor is its mean value. For example, b_2 represents the relationship between aggression and callous traits for someone who spends the average number of hours gaming per week.

Therefore, centring is particularly important when your model contains an interaction term because it makes the bs for lower-order effects interpretable. There are good reasons for not caring about the lower-order effects when the higher-order interaction involving those effects is significant, but when it is not, centring will make interpreting the main effects easier. For example, if the gaming × callous traits interaction is significant, then it's not clear why we would be interested in the individual effects of gaming and callous traits. In any case, with centred variables the bs for individual predictors have two interpretations: (1) they are the effect of that predictor at the mean value of the sample; and (2) they are the average effect of the predictor across the range of scores for the other predictors. To explain the second interpretation, imagine we took everyone who spent no hours gaming and computed the regression between aggression and callous traits and noted the b, then we took everyone who played games for 1 hour and did the same, then we took everyone who gamed for 2 hours per week and did the same. We continued doing this until we had computed regressions for every different value of the hours spent gaming. We'd have a lot of bs: each one representing the relationship between callous traits and aggression but for

[3] Remember that with z-scores we go a step further and divide the centred scores by the standard deviation of the original data, which changes the units of measurements to standard deviations.

different amounts of gaming. If we took an average of these *b*s then we'd get the same value as the *b* for callous traits (centred) when we use it as a predictor with gaming (centred) and their interaction.

The *PROCESS* tool will do the centring for us so we don't really need to worry too much about how it's done, but because centring is useful in other analyses Oliver Twisted has some additional material that shows you how to do it manually for this example.

OLIVER TWISTED

Please, Sir, can I have some more … centring?

'Recentgin', babbles Oliver as he stumbles drunk out of Mrs Moonshine's alcohol emporium. 'I've had some recent gin.' I think you mean *centring* Oliver, not *recentgin*. If you want to know how to centre your variables using SPSS, then the additional material for this chapter on the companion website will tell you.

10.3.4. Creating interaction variables ②

Equation (10.1) contains a variable called 'Interaction', but the data file does not. The question you might well ask is how we enter a variable into the model that doesn't exist in the data set. We can create it, and it's easier than you might think. Mathematically speaking, when we look at the combined effect of two variables (an interaction) we are literally looking at the effect of the two variables multiplied together. So the interaction variable in this case would literally be the scores on the time spent gaming multiplied by the scores for callous-unemotional traits. That's why interactions are denoted as *variable 1* × *variable 2*. The way we'll do moderation analysis in SPSS creates the interaction variable for you, but the self-help task gives you some practice at doing it manually (which might be handy for future reference).

SELF-TEST Follow Oliver Twisted's instructions to create the centred variables **CUT_Centred** and **Vid_Centred**. Then use the *compute* command to create a new variable called **Interaction** in the **Video Games.sav** file, which is **CUT_Centred** multiplied by **Vid_Centred.**

10.3.5. Following up an interaction effect ②

As we have already seen, moderation is shown by a significant interaction between variables. However, if the moderation effect is significant, then we need to delve a bit deeper to find out the nature of the moderation. In our example, we're predicting that the moderator (callous traits) will influence the relationship between playing violent video games and aggression. If the interaction of callous traits and time spent gaming is a significant predictor of aggression then we know that we have a moderation effect, but we don't know the nature of the effect. It could be that the time spent gaming always has a positive relationship with aggression, but that relationship gets stronger the more a person has callous traits. Alternatively, perhaps in people low on callous traits the time spent gaming *reduces* aggression but it *increases* aggression in those high on callous traits (i.e., the relationship reverses). To find out what is going on we need to do something known as **simple slopes analysis** (Aiken & West, 1991; Rogosa, 1981).

The idea behind simple slopes analysis is fairly straightforward and it's really no different than what was illustrated in Figure 10.6. When describing that figure I talked about comparing the relationship between the predictor (time spent gaming) and outcome (aggression) at low and high levels of the moderator (callous traits). For example, in the right panel of Figure 10.6, we saw that time spent gaming and aggression had a slightly negative relationship at low levels of callous traits, but a fairly strong positive relationship at high levels of callous traits. This is the essence of simple slopes analysis: we work out the regression equations for the predictor and outcome and low, high and average levels of the moderator. The 'high' and 'low' levels can be anything you like, but *PROCESS* uses 1 standard deviation above and below the mean value of the moderator. Therefore, in our example, we would get the regression model for aggression predicted from hours spent gaming for the average value of callous traits, for 1 standard deviation above the mean value of callous traits and for one standard deviation below the mean value of callous traits. We compare these slopes both in terms of their significance, and the value and direction of the *b* to see whether the relationship between hours spent gaming and aggression changes at different levels of callous traits.

A slightly different approach is to look at how the relationship between the predictor and outcome changes at lots of different values of the moderator (not just at high, low and mean values). One such approach implemented by *PROCESS* is based on Johnson and Neyman (1936). Essentially, it computes the regression model for the predictor and outcome at lots of different values of the moderator. For each model it computes the significance of the regression slope so you can see for which values of the moderator the relationship between the predictor and outcome is significant. It returns a 'zone of significance',[4] which consists of two values of the moderator. Typically, for values in between these two values of the moderator the predictor does not significantly predict the outcome. Values below the lower value and above the upper value are values of the moderator for which the predictor significantly predicts the outcome.

10.3.6. Running the analysis ②

Given that moderation is demonstrated through a significant interaction between the predictor and moderator in a regression, we could follow the general procedure for fitting linear models in Chapter 8 (Figure 8.11). We would first centre the predictor and moderator, then create the interaction term as discussed already, then run a forced entry regression with the centred predictor, centred moderator and the interaction of the two centred variables as predictors. The advantage of this approach is that we can inspect sources of bias in the model.

SELF-TEST Assuming you have done the other self-test, run a regression predicting **Aggression** from **CUT_Centred**, **Vid_Centred** and **Interaction**.

Using the *PROCESS* tool (if you haven't installed it yet, see Section 10.2) has several advantages over using the normal regression tools: (1) it will centre predictors for us; (2) it computes the interaction term automatically; and (3) it will do simple slopes analysis. To access the dialog boxes in Figure 10.8 select Analyze Regression ▸

[4] I have to be careful not to confuse this with my wife, who is the Zoë of significance.

PROCESS, by Andrew F. Hayes (http://www.afhayes.com). The variables in your data file will be listed in the box labelled *Data File Variables*. Select the outcome variable (in this case **Aggression**) and drag it to the box labelled *Outcome Variable (Y)*, or click on . Similarly, select the predictor variable (in this case **Vid_Games**) and drag it to the box labelled *Independent Variable (X)*. Finally, select the moderator variable (in this case **CaUnTs**) and drag it to the box labelled *M Variable(s)*, or click on . This box is where you specify any moderators (you can have more than one).

PROCESS can test 74 different types of model, and these models are listed in the drop-down box labelled *Model Number.* If you want to investigate all 74 different models then have a look at the *PROCESS* documentation (http://www.afhayes.com/public/process.pdf). Simple moderation analysis is represented by model 1, but the default model is 4 (mediation, which we'll look at next). Therefore, activate this drop-down list and select 1. The rest of the options in this dialog box are for models other than simple moderation, so we'll ignore them.

If you click on Options another dialog box will appear containing four useful options for moderation. Selecting (1) *Mean center for products* centres the predictor and moderator for you; (2) *Heteroscedasticity-consistent SEs* means we need not worry about having heteroscedasticity in the model; (3) *OLS/ML confidence intervals* produces confidence intervals for the model, and I've tried to emphasize the importance of these throughout the book; and (4) *Generate data for plotting* is helpful for interpreting and visualizing the simple slopes analysis. Talking of simple slopes analysis, if you click on Conditioning, you can change whether you want simple slopes at ±1 standard deviation of the mean of the moderator (the default, which is fine) or at percentile points (it uses the 10th, 25th, 50th, 75th and 90th percentiles). It is useful to select the *Johnson-Neyman* method to get a zone of significance for the moderator. Back in the main dialog box, click on OK to run the analysis.

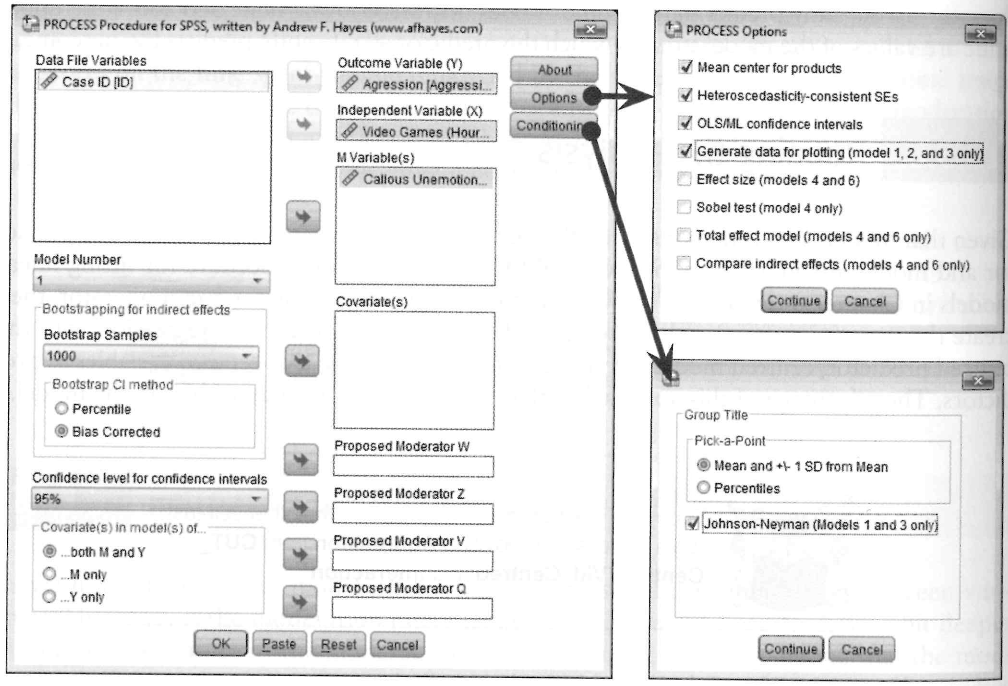

FIGURE 10.8 The dialog boxes for running moderation analysis

10.3.7. Output from moderation analysis ②

The first thing to notice about the output is it appears as text rather than being nicely formatted in tables. Try not to let this formatting disturb you. If your output looks odd

SPSS TIP 10.1 Troubleshooting *PROCESS* ②

There are a few things worth knowing about *PROCESS* that might help to prevent weird stuff happening.

- If the variable names entered into *PROCESS* are longer than 8 characters, it shortens them to 8 characters. Therefore, if you enter variables with similar long names *PROCESS* will get confused. For example, if you had two variables in the data editor called **NumberOfNephariousActs** and **NumberOfBlackSabbathAlbumsOwned** they would both be shortened to **numberof** (or possibly **number~1** and **number~2**) and *PROCESS* will get confused about which variable is which. If your output looks weird, then check your variable names.
- Don't call any of your variables **xxx** (I'm not sure why you would) because that is a reserved variable name in *PROCESS*, so naming a variable **xxx** will confuse it.
- *PROCESS* is also confused by string variables, so only enter numeric variables.

or contains warnings, or has a lot of zeros in it, it might be worth checking the variables that you input into *PROCESS* (SPSS Tip 10.1). However, assuming everything has gone smoothly, you should see Output 10.1, which is the main moderation analysis. This output is pretty much the same as the table of regression coefficients that we saw in Chapter 8. We're told the *b*-value for each predictor, and the associated standard errors (which have been adjusted for heteroscedasticity because we asked for them to be). Each *b* is compared to zero using a *t*-test, which is computed from the beta divided by its standard error. The confidence interval for the *b* is also produced (because we asked for it). Moderation is shown up by a significant interaction effect, and in this case the interaction is highly significant, $b = 0.027$, 95% CI [0.013, 0.041], $t = 3.71$, $p < .001$, indicating that the relationship between the time spent gaming and aggression is moderated by callous traits.

SELF-TEST Assuming you did the previous self-test, compare the table of coefficients that you got with those in Output 10.1.

To interpret the moderation effect we can examine the simple slopes, which are shown in Output 10.2. Essentially, the table shows us the results of three different regressions: the regression for time spent gaming as a predictor of aggression (1) when callous traits are low (to be precise when the value of callous traits is −9.6177); (2) at the mean value of callous traits (because we centred callous traits its mean value is zero as indicated in the output); and (3) when the value of callous traits is 9.6177 (i.e., high). We can interpret these three regressions as we would any other: we're interested in the value of *b* (called *Effect* in the output), and its significance. From what we have already learnt about regression we can interpret the three models as follows:

1. When callous traits are low, there is a non-significant negative relationship between time spent gaming and aggression, $b = -0.091$, 95% CI [−0.299, 0.117], $t = -0.86$, $p = .392$.
2. At the mean value of callous traits, there is a significant positive relationship between time spent gaming and aggression, $b = 0.170$, 95% CI [0.020, 0.319], $t = 2.23$, $p = .026$.
3. When callous traits are high, there is a significant positive relationship between time spent gaming and aggression, $b = 0.430$, 95% CI [0.231, 0.628], $t = 4.26$, $p < .001$.

These results tell us that the relationship between time spent playing violent video games and aggression only really emerges in people with average or greater levels of callous-unemotional traits.

OUTPUT 10.1

```
**************************************************************************
Model = 1
    Y = Aggressi
    X = Vid_Game
    M = CaUnTs

Sample size
        442
**************************************************************************
Outcome: Aggressi
Model Summary
```

R	R-sq	F	df1	df2	p
.6142	.3773	90.5311	3.0000	438.0000	.0000

Model

	coeff	se	t	p	LLCI	ULCI
constant	39.9671	.4750	84.1365	.0000	39.0335	40.9007
CaUnTs	.7601	.0466	16.3042	.0000	.6685	.8517
Vid_Game	.1696	.0759	2.2343	.0260	.0204	.3188
int_1	.0271	.0073	3.7051	.0002	.0127	.0414

```
Interactions:
 int_1    Vid_Game     X     CaUnTs
```

OUTPUT 10.2

```
**************************************************************************
Conditional effect of X on Y at values of the moderator(s)
```

CaUnTs	Effect	se	t	p	LLCI	ULCI
-9.6177	-.0907	.1058	-.8568	.3920	-.2986	.1173
.0000	.1696	.0759	2.2343	.0260	.0204	.3188
9.6177	.4299	.1010	4.2562	.0000	.2314	.6284

```
Values for quantitative moderators are the mean and plus/minus one SD from mean
```

Output 10.3 shows the output of the Johnson–Neyman method, and this gives a different approach to simple slopes. First we're told the boundaries of the zone of significance: it is between −17.1002 and −0.7232. Remember that these are the values of the centred version of the callous-unemotional traits variable, and define regions within which the relationship between the time spent gaming and aggression is significant. The table underneath gives a detailed breakdown of these regions. Essentially it's doing something quite similar to the simple slopes analysis: it takes different values of callous and unemotional traits and for each one computes the b (*Effect*) and its significance for the relationship between the time spent gaming and aggression. I have annotated the output to show the boundaries of the zone of significance. If you look at the column labelled p you can see that we start off with a significant negative relationship between time spent gaming and aggression, $b = -0.334$, 95% CI [−0.645, −0.022], $t = -2.10$, $p = .036$. As we move up to the next value of callous traits (−17.1002), the relationship between time spent gaming and aggression is still significant ($p = .0500$), but at the next value it becomes non-significant ($p = .058$). Therefore, the

OUTPUT 10.3

```
********************** JOHNSON-NEYMAN TECHNIQUE **************************

Moderator value(s) defining Johnson-Neyman significance region(s)
  -17.1002
    -.7232

Conditional effect of X on Y at values of the moderator (M)
```

CaUnTs	Effect	se	t	p	LLCI	ULCI	
-18.5950	-.3336	.1587	-2.1027	.0361	-.6454	-.0218	Significant
-17.1002	-.2931	.1492	-1.9654	.0500	-.5863	.0000	
-16.4450	-.2754	.1451	-1.8987	.0583	-.5605	.0097	Not significant
-14.2950	-.2172	.1319	-1.6467	.1003	-.4765	.0420	
-12.1450	-.1590	.1194	-1.3319	.1836	-.3937	.0756	
-9.9950	-.1009	.1077	-.9361	.3497	-.3126	.1109	
-7.8450	-.0427	.0972	-.4390	.6609	-.2338	.1484	
-5.6950	.0155	.0882	.1757	.8606	-.1579	.1889	
-3.5450	.0737	.0813	.9059	.3655	-.0862	.2336	
-1.3950	.1319	.0771	1.7111	.0878	-.0196	.2833	
-.7232	.1501	.0763	1.9654	.0500	.0000	.3001	Significant
.7550	.1901	.0759	2.5053	.0126	.0410	.3392	
2.9050	.2482	.0779	3.1878	.0015	.0952	.4013	
5.0550	.3064	.0829	3.6980	.0002	.1436	.4693	
7.2050	.3646	.0903	4.0360	.0001	.1871	.5422	
9.3550	.4228	.0997	4.2386	.0000	.2267	.6188	
11.5050	.4810	.1106	4.3490	.0000	.2636	.6983	
13.6550	.5392	.1225	4.4013	.0000	.2984	.7799	
15.8050	.5973	.1352	4.4188	.0000	.3317	.8630	
17.9550	.6555	.1484	4.4160	.0000	.3638	.9473	
20.1050	.7137	.1621	4.4017	.0000	.3950	1.0324	
22.2550	.7719	.1762	4.3814	.0000	.4256	1.1181	
24.4050	.8301	.1905	4.3580	.0000	.4557	1.2044	

```
**************************************************************************
```

threshold for significance ends at −17.1002 (which we were told at the top of the output). As we increase the value of callous-unemotional traits the relationship between time spent gaming and aggression remains non-significant until the value of callous-unemotional traits is −0.723, at which point it just crosses the threshold for significance again. For all subsequent values of callous-unemotional traits the relationship between time spent gaming and aggression is significant. Looking at the b-values themselves (in the column labelled *Effect*) we can also see that with increases in callous-unemotional traits the strength of relationship between time spent gaming and aggression goes from a small negative effect ($b = -0.334$) to a fairly strong positive one ($b = 0.830$).

The final way we can look at these effects is by graphing them. In Figure 10.8 we asked *PROCESS* to generate data for plotting and these data are at the bottom of the output (see Figure 10.9). We're given values of the variable **Vid_Games** (-6.9622, 0, 6.9622) and of **CaUnTs** (-9.6177, 0, 9.6177). These values are not important in themselves, but they correspond to low, mean and high values of the variable. The **yhat** tells us the predicted values of the outcome (aggression) for these combinations of the predictors. For example, when **Vid_Games** and **CaUnTs** are both low (−6.9622 and −9.6177, respectively) the predicted value of aggression is 33.2879, when both variables are at their mean (0 and 0), the predicted value of aggression is 39.9671, and so on. To create a simple slopes graph we need to put these values in a data file. The simplest way to create the new data file is to create coding variables that represent low, mean and high (use any codes you like). Then enter all combinations of these codes. For example, in Figure 10.9 I've created variables called **Games** and **CaUnTs** both of which are coding variables (1 = low, 2 = mean, 3 = high) and then entered the combinations of these codes that correspond to the *PROCESS* output (e.g., low–low, mean–low, high–low), then I have typed in the corresponding predicted values from the *PROCESS* output. Hopefully you can see from Figure 10.9 how the output from *PROCESS* corresponds to the new data file. You can access this file as **Video Game Graph.sav** if you can't work out how to create it yourself. Having transferred the output to a data file, we can draw line graphs using what we learnt in Chapter 4.

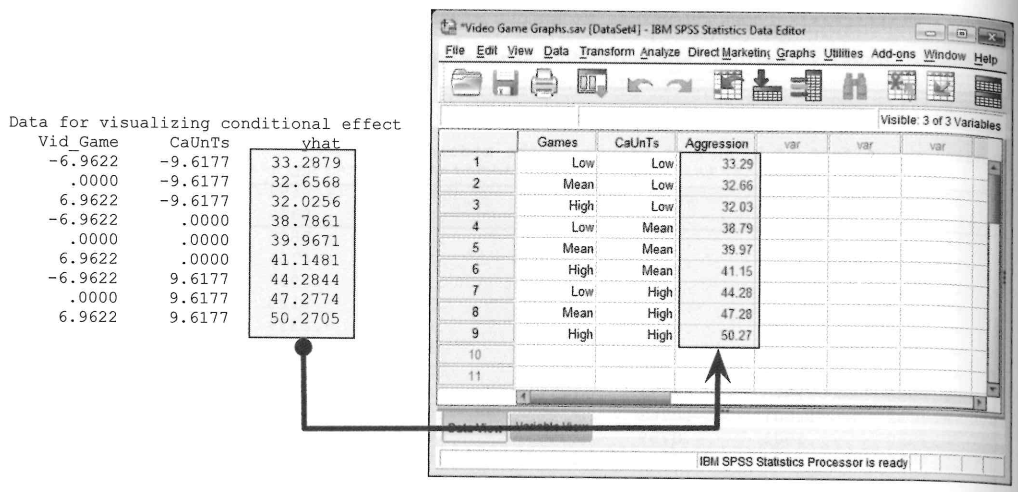

FIGURE 10.9 Entering data for graphing simple slopes

SELF-TEST Draw a multiple line graph of **Aggression** (*y*-axis) against **Games** (*x*-axis) with different-coloured lines for different values of **CaUnTs**.

The resulting graph from the self-test is shown in Figure 10.10. The graph shows what we found from the simple slopes analysis: when callous traits are low (blue line) there is a non-significant negative relationship between time spent gaming and aggression; at the mean value of callous traits (green line) there is small positive relationship between time spent gaming and aggression; and this relationship gets even stronger at high levels of callous traits (beige line).

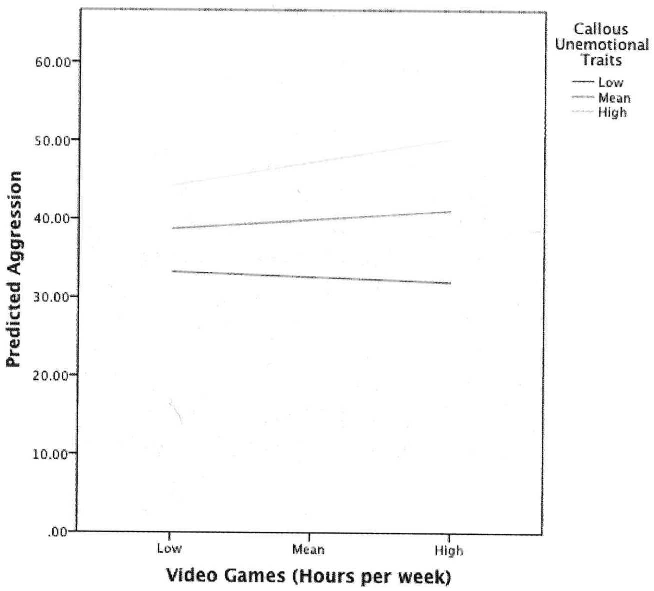

FIGURE 10.10 Simple slopes equations of the regression of aggression on video games at three levels of callous traits

SELF-TEST Now draw a multiple line graph of **Aggression** (*y*-axis) against **CaUnTs** (*x*-axis) with different-coloured lines for different values of **Games**.

10.3.8. Reporting moderation analysis ②

Moderation analysis is just regression, so we can report it in the same way as described in Section 8.9. My personal preference would be to produce a table such as Table 10.1.

TABLE 10.1 Linear model of predictors of aggression

	b	*SE B*	*t*	*p*
Constant	39.97 [39.03, 40.90]	0.475	84.13	$p < .001$
Callous Traits (centred)	0.76 [0.67, 0.85]	0.047	16.30	$p < .001$
Gaming (centred)	0.17 [0.02, 0.32]	0.076	2.23	$p = .026$
Callous Traits x Gaming	0.027 [0.01, 0.04]	0.007	3.71	$p < .001$

Note. $R^2 = .38$.

CRAMMING SAM'S TIPS Moderation

- Moderation occurs when the relationship between two variables changes as a function of a third variable. For example, the relationship between watching horror films and feeling scared at bedtime might increase as a function of how vivid an imagination a person has.
- Moderation is tested using a regression in which the outcome (fear at bedtime) is predicted from a predictor (how many horror films are watched), the moderator (imagination) and the interaction of these variables.
- Predictors should be centred before the analysis.
- The interaction of two variables is simply the scores on the two variables multiplied together.
- If the interaction is significant then moderation is present.
- If moderation is found, follow up the analysis with simple slopes analysis. This analysis looks at the relationship between the predictor and outcome at low, mean and high levels of the moderator.

10.4. Mediation ②

10.4.1. The conceptual model ②

Whereas moderation alludes to the combined effect of two variables on an outcome, **mediation** refers to a situation when the relationship between a predictor variable and an outcome variable can be explained by their relationship to a third variable (the **mediator**). The top of Figure 10.11 shows a basic relationship between a predictor and an outcome (denoted as *c*). However, the bottom of the figure shows that these variables are also related to a third variable in specific ways: (1) the predictor also predicts the mediator through the path denoted by *a*; (2) the mediator predicts the outcome through the path denoted by *b*. The relationship between the predictor and outcome will probably be different when the mediator is also included in the model and so is denoted *c*'. The letters denoting each path (*a*, *b*, *c* and *c*') represent the unstandardized regression coefficient between the variables connected by the arrow; therefore, they symbolize the strength of relationship between variables. Mediation is said to have occurred if the strength of the relationship between the predictor and outcome is reduced by including the mediator (i.e., the regression parameter for *c*' is smaller than for *c*). Perfect mediation occurs when *c*' is zero: in other words, the relationship between the predictor and outcome is completely wiped out by including the mediator in the model.

This description is all a bit abstract, so let's use an example. My wife and I often wonder what the important factors are in making a relationship last. For my part, I don't really understand why she'd want to be with a balding heavy rock fan with an oversized collection of vinyl and musical instruments and an unhealthy love of *Doctor Who* and numbers. It is important I gather as much information as possible about keeping her happy because the odds are stacked against me. For her part I have no idea why she wonders: her very existence makes me happy. Perhaps if you are in a relationship you have wondered how to make it last too.

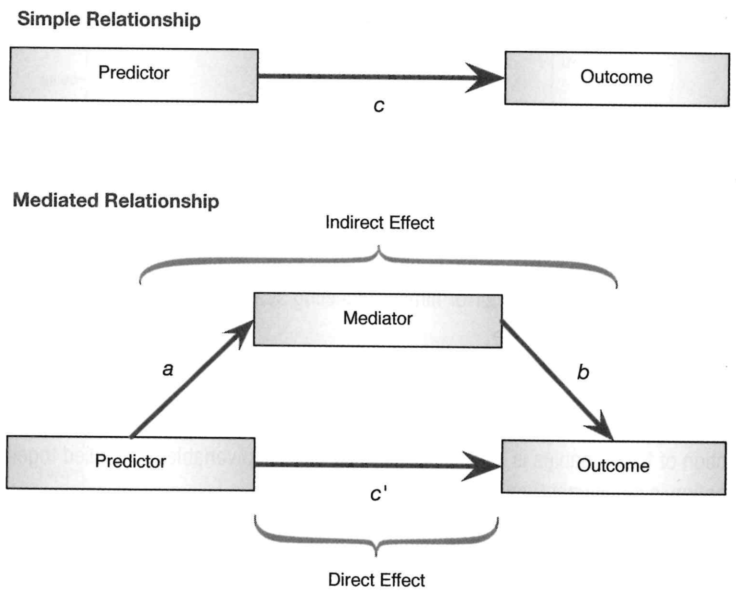

FIGURE 10.11 Diagram of a basic mediation model

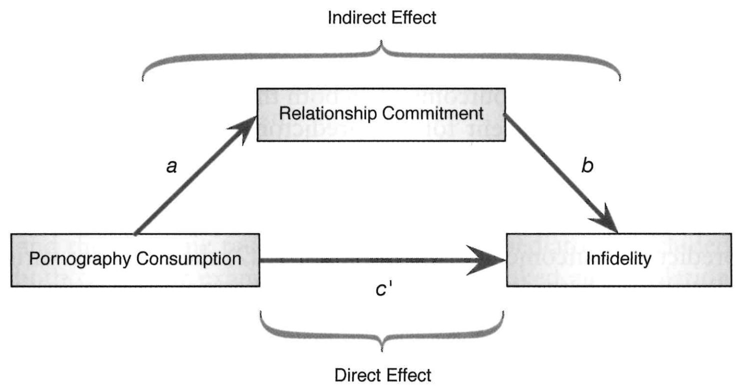

FIGURE 10.12 Diagram of a mediation model from Lambert et al. (2012)

During our cyber-travels, Mrs Field and I have discovered that physical attractiveness (McNulty, Neff, & Karney, 2008), conscientiousness and neuroticism (good for us) predict marital satisfaction (Claxton, O'Rourke, Smith, & DeLongis, 2012). Pornography use probably doesn't: it is related to infidelity (Lambert, Negash, Stillman, Olmstead, & Fincham, 2012). Mediation is really all about the variables that explain relationships like these: it's unlikely that everyone who catches a glimpse of some porn suddenly rushes out of their house to have an affair – presumably it leads to some kind of emotional or cognitive change that undermines the love glue that holds us and our partners together. Lambert et al. tested this hypothesis. Figure 10.12 shows their mediator model: the initial relationship is that between pornography consumption (the predictor) and infidelity (the outcome), and they hypothesized that this relationship is mediated by commitment (the mediator). This model suggests that the relationship between pornography consumption and infidelity isn't a direct effect but operates through a reduction in relationship commitment. For this hypothesis to be true: (1) pornography consumption must predict infidelity in the first place (path *c*); (2) pornography consumption must predict relationship commitment (path *a*); (3) relationship commitment must predict infidelity (path *b*); and (4) the relationship between pornography consumption and infidelity should be smaller when relationship commitment is included in the model than when it isn't. We can distinguish between the **direct effect** of pornography consumption on infidelity, which is the relationship between them controlling for relationship commitment, and the **indirect effect**, which is the effect of pornography consumption on infidelity through relationship commitment (Figure 10.12).

10.4.2. The statistical model ②

Unlike moderation, the statistical model for mediation is basically the same as the conceptual model: it is characterized in Figure 10.11. Historically, this model was tested through a series of regression analyses, which reflect the four conditions necessary to demonstrate mediation (Baron & Kenny, 1986). I have mentioned already that the letters denoting the paths in Figure 10.11 represent the unstandardized regression coefficients for the relationships between variables denoted by the path. Therefore, to estimate any one of these paths, we want to know the unstandardized regression coefficient for the two variables involved. For example, Baron and Kenny suggested in their seminal paper that mediation is tested through three regression models (see also Judd & Kenny, 1981):

1. A regression predicting the outcome from the predictor variable. The regression coefficient for the predictor gives us the value of *c* in Figure 10.11.

2 A regression predicting the mediator from the predictor variable. The regression coefficient for the predictor gives us the value of a in Figure 10.11.

3 A regression predicting the outcome from both the predictor variable and the mediator. The regression coefficient for the predictor gives us the value of c' in Figure 10.11, and the regression coefficient for the mediator gives us the value of b.

These models test the four conditions of mediation: (1) the predictor variable must significantly predict the outcome variable in model 1; (2) the predictor variable must significantly predict the mediator in model 2; (3) the mediator must significantly predict the outcome variable in model 3; and (4) the predictor variable must predict the outcome variable less strongly in model 3 than in model 1.

In Lambert et al.'s (2012) study, all participants had been in a relationship for at least a year. The researchers measured pornography consumption on a scale from 0 (low) to 8 (high), but this variable, as you might expect, was skewed (most people had low scores) so they analysed log-transformed values (**LnConsumption**). They also measured commitment to their current relationship (**Commitment**) on a scale from 1 (low) to 5 (high). Infidelity was measured in terms of questions asking whether the person had committed a physical act (**Infidelity**) that they or their partner would consider to be unfaithful (0 = no, 1 = one of them would consider it unfaithful, 2 = both of them would consider the act unfaithful),[5] and also in terms of the number of people they had 'hooked up' with in the previous year (**Hook_Ups**), which would mean during a time period in which they were in their current relationship.[6] The actual data from Lambert et al.'s study are in the file **Lambert et al. (2012).sav**.

SELF-TEST Run the three regressions necessary to test mediation for Lambert et al.'s data: (1) a regression predicting **Infidelity** from **LnConsumption**; (2) a regression predicting **Commitment** from **LnConsumption**; and (3) a regression predicting **Infidelity** from both **LnConsumption** and **Commitment**. Is there evidence of mediation?

Many people still use this approach to test mediation: Baron and Kenny's article has been cited over 35,000 times in scientific papers, which gives you some idea of how influential this method has been. I think it is very useful for illustrating the principles of mediation and for understanding what mediation means. However, the method of regressions has some limitations. The main one is the fourth criterion by which mediation is assessed: *the predictor variable must predict the outcome variable less strongly in model 3 than in model 1*. Although we know that perfect mediation is shown when the relationship between the predictor and outcome is reduced to zero in model 3, usually this doesn't happen. Instead, you see a reduction in the relationship between the predictor and outcome, rather than the relationship being reduced to zero. This raises the question of how much of a reduction is necessary to infer mediation.

Although Baron and Kenny advocated looking at the sizes of the regression parameters, in practice people tend to look for a change in significance; so, mediation would occur if the relationship between the predictor and outcome was significant ($p < .05$) when looked at in isolation (model 1) but not significant ($p > .05$) when the mediator is included too (model 3). This approach can lead to all sorts of silliness because of the all-or-nothing

[5] I've coded this variable differently from the original data to make interpretation of it more intuitive, but it doesn't affect the results.

[6] A 'hook-up' was defined to participants as 'when two people get together for a physical encounter and don't necessarily expect anything further (e.g., no plan or intention to do it again)'.

thinking that p-values encourage. You could have a situation in which the b-value for the relationship between the predictor and outcome changes very little in models with and without the mediator, but the p-value shifts from one side of the threshold to another (e.g., from $p = .049$ when the mediator isn't included to $p = .051$ when it is). Even though the p-values have changed from significant to not significant, the change is very small, and the size of the relationship between the predictor and outcome will not have changed very much at all. Similarly, you could have a situation where the b for the relationship between the predictor and the outcome reduces a lot when the mediator is included, but remains significant in both cases. For example, perhaps when looked at in isolation the relationship between the predictor and outcome is $b = 0.46$, $p < .001$, but when the mediator is included as a predictor as well it reduces to $b = 0.18$, $p = .042$. You'd conclude (based on significance) that no mediation had occurred despite the fact that relationship between the predictor and outcome is less than half its original value.

An alternative is to estimate the indirect effect and its significance. The indirect effect is illustrated in Figures 10.11 and 10.12: it is the combined effects of paths a and b. The significance of this effect can be assessed using the **Sobel test** (Sobel, 1982). If the Sobel test is significant it means that the predictor significantly affects the outcome variable via the mediator. In other words, there is significant mediation. This test works well in large samples, but you're better off computing confidence intervals for the indirect effect using bootstrap methods (Section 5.4.3). Now that computers make it easy for us to estimate the indirect effect (i.e., the effect of mediation) and its confidence interval, this practice is becoming increasingly common and is preferable to Baron and Kenny's regressions and the Sobel test because it's harder to get sucked into the black-and-white thinking of significance testing (Section 2.6.2.2). People tend to apply Baron and Kenny's method in a way that is intrinsically bound to looking for 'significant' relationships, whereas estimating the indirect effect and its confidence interval allows us to simply report the degree of mediation observed in the data.

10.4.3. Effect sizes of mediation ③

If we're going to look at the size of the indirect effect to judge whether mediation has occurred, then it's useful to have effect size measures to help us (see Section 2.7.1). Many effect size measures have been proposed and are discussed in detail elsewhere (MacKinnon, 2008; Preacher & Kelley, 2011). The simplest is to look at the regression coefficient for the indirect effect and its confidence interval. Figure 10.11 shows us that the indirect effect is the combined effect of paths a and b. We have also seen that a and b are unstandardized regression coefficients for the relationships between variables denoted by the path. To find the combined effect of these paths, we simply multiply these regression coefficients:

$$\text{indirect effect} = ab \tag{10.2}$$

The resulting value is an unstandardized regression coefficient like any other, and consequently is expressed in the original units of measurement. As we have seen, it is sometimes useful to look at standardized regression parameters, because these can be compared across different studies using different outcome measures (see Chapter 8). MacKinnon (2008) suggested standardizing this measure by dividing by the standard deviation of the outcome variable:

$$\text{indirect effect (partially standardized)} = \frac{ab}{s_{\text{Outcome}}} \tag{10.3}$$

This standardizes the indirect effect with respect to the outcome variable, but not the predictor or mediator. As such, it is sometimes referred to as the partially standardized indirect effect. To fully standardize the indirect effect we would need to multiply the partially standardized measures by the standard deviation of the predictor variable (Preacher & Hayes, 2008b):

$$\text{indirect effect (standardized)} = \frac{ab}{s_{\text{Outcome}}} \times s_{\text{Predictor}} \tag{10.4}$$

This measure is sometimes called the **index of mediation**. This measure is useful in that it can be compared across different mediation models that use different measures of the predictor, outcome and mediator. Reporting this measure would be particularly helpful if anyone decides to include your research in a meta-analysis.

A different approach to estimating the size of the indirect effect is to look at the size of the indirect effect relative to either the total effect of the predictor or the direct effect of the predictor. For example, if we wanted the ratio of the indirect effect (*ab*) to the total effect (*c*) we could use the regression parameters from the various regressions displayed in Figure 10.11:

$$P_M = \frac{ab}{c} \tag{10.5}$$

Similarly, if we want to express the indirect effect as a ratio of the direct effect (c'), the regressions give us everything we need:

$$R_M = \frac{ab}{c'} \tag{10.6}$$

These ratio-based measures only really re-describe the original indirect effect. Both are very unstable in small samples, and MacKinnon (2008) advises against using P_M and R_M in samples smaller than 500 and 5000, respectively. Also, although it is tempting to think of P_M as a proportion (because it is the ratio of the indirect effect compared to the *total* effect) it is not: it can exceed 1 and even take on negative values (Preacher & Kelley, 2011). For these reasons, these ratio measures are probably best avoided.

In regression we used R^2 as a measure of the proportion of variance explained by a predictor (or several predictors). We can compute a form of R^2 for the indirect effect, which tells us the proportion of variance explained by the indirect effect. MacKinnon (2008) proposes several versions, but *PROCESS* computes this one:

$$R_M^2 = R_{Y,M}^2 - \left(R_{Y,MX}^2 - R_{Y,X}^2\right) \tag{10.7}$$

This uses the proportion of variance in the outcome variables explained by the predictor ($R_{Y,X}^2$), the mediator ($R_{Y,M}^2$), and both ($R_{Y,MX}^2$). It can be interpreted as the variance in the outcome that is shared by the mediator and the predictor, but that cannot be attributed to either in isolation. Again, this measure is not bound to fall between 0 and 1, and it's possible to get negative values (which usually indicate suppression effects rather than mediation).

The final measure that I'll consider was proposed by Preacher and Kelley (2011) and is called kappa-squared (κ^2). If you read the original article, it is full of scary equations that make this measure very difficult to explain. However, at a conceptual level it is a very simple and elegant idea: kappa-squared expresses the indirect effect as a ratio to the maximum possible indirect effect that you could have found given the design of your study:

$$\kappa^2 = \frac{ab}{\max(ab)} \tag{10.8}$$

The scary maths comes into play in how the maximum possible value of the indirect effect is computed. However, we have computers to do that for us, so let's just imagine that a frog called Hugglefrall sticks his big slimy tongue out and numbers attach themselves to it. He then swirls the numbers around in his mouth, does that funny expanding throat thing that frogs sometimes do, and then belches out the value for us. Beyond that, all we need to know is that kappa is a proportion and we can interpret it as such: values of 0 mean the indirect effect is very small relative to the maximum possible value, and values close to 1 mean that it is as large as it could possibly be given the design that we have. Not that I should really encourage this sort of thing, but in terms of what constitutes a large effect, κ^2 can be equated to the values used for R^2: a small effect is .01, a medium effect would be around .09, and a large effect in the region of .25 (Preacher & Kelley, 2011).

PROCESS computes all of the effect size measures that I have discussed, but of them all probably the most useful are the unstandardized and standardized indirect effect and κ^2. All of the measures discussed have accompanying confidence intervals and are unaffected by sample sizes (although note my earlier comments about the variability of P_M and R_M in small samples). However, P_M, R_M and R_M^2 cannot be interpreted easily because they allude to being proportions but are not, and all of the measures apart from κ^2 are unbounded, which again makes interpretation tricky (Preacher & Kelley, 2011).

10.4.4. Running the analysis ②

Assuming we're going to test Lambert's mediation model (Figure 10.12) by estimating the indirect effect rather than through a Baron and Kenny style mediation analysis, then we can again use Hayes's *PROCESS* tool (see Section 10.2 if you haven't installed it yet). To access the dialog boxes in Figure 10.13 select Analyze Regression ▸ PROCESS, by Andrew F. Hayes (http://www.afhayes.com). The variables in your data file will be listed in the box labelled *Data File Variables*. Select the outcome variable (in this case **Infidelity**) and drag it to the box labelled *Outcome Variable (Y)*, or click on ➡. Similarly, select the predictor variable (in this case **LnConsumption**) and drag it to the box labelled *Independent Variable (X)*. Finally, select the mediator variable (in this case **Commitment**) and drag it to the box labelled *M Variable(s)*, or click on ➡. This box is where you specify any mediators (you can have more than one).

As I mentioned before, *PROCESS* can test many different types of model, and simple mediation analysis is represented by model 4 (this model is selected by default). Therefore, make sure that 4 is selected in the drop-down list under *Model Number*. Unlike moderation, there are other options in this dialog box that are useful: for example, to test the indirect effects we will use bootstrapping to generate a confidence interval around the indirect effect. By default *PROCESS* uses 1000 bootstrap samples, and will compute bias corrected and accelerated confidence intervals. These default options are fine, but just be aware that you can ask for percentile bootstrap confidence intervals instead (see Section 5.4.3).

If you click on Options... another dialog box will appear containing four useful options for mediation. Selecting (1) *Effect size* produces the estimates of the size of the indirect effect

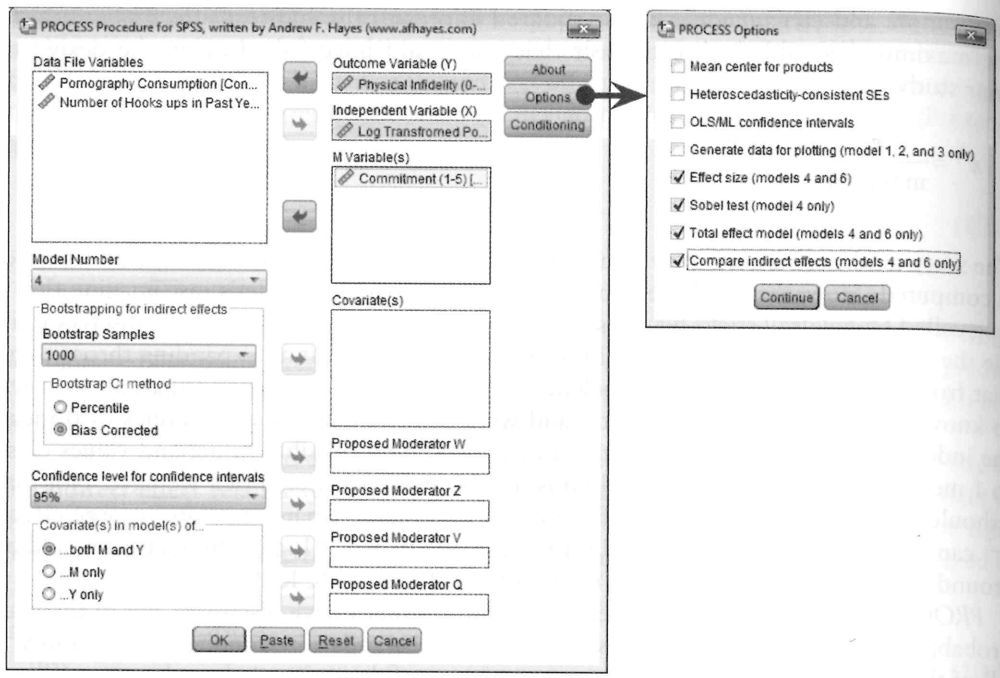

FIGURE 10.13 The dialog boxes for running mediation analysis

ODITI'S LANTERN

moderation and mediation

'I, Oditi, want you to join my cult of undiscovered numerical truths. I also want you to stare into my lantern to gain statistical enlightenment. It's possible that statistical knowledge mediates the relationship between staring into my lantern and joining my cult ... or it could be mediated by neurological changes to your brain created by the subliminal messages in the videos. Stare into my lantern to find out about mediation and moderation.'

discussed in Section 10.4.3;[7] (2) *Sobel test* produces a significance test of the indirect effect devised by Sobel; (3) *Total effect model* produces the direct effect of the predictor on the outcome (in this case the regression of infidelity predicted from pornography consumption); and (4) *Compare indirect effects* will, when you have more than one mediator in the model, estimate the effect and confidence interval for the difference between the indirect effects resulting from these mediators. This final option is useful when you have more than one mediator to compare their relative importance in explaining the relationship between the predictor and outcome. However, we have only a single mediator so we don't need to select this option (you can select it if you like, but it won't change the output produced). None of the options activated by clicking on Conditioning apply to simple mediation models, so we can ignore this button and click OK to run the analysis.

10.4.5. Output from mediation analysis ②

As with moderation, the output appears as text. Output 10.4 shows the first part of the output, which initially tells us the name of the outcome (*Y*), the predictor (*X*) and the mediator (*M*)

[7] R^2_M and κ^2 are produced only for models with a single mediator. Although I don't look at more complex models, bear this in mind if you run models including more than one mediator, or covariates.

variables, which have been shortened to 8 letters (SPSS Tip 10.1). This is useful for double-checking we have entered the variables in the correct place: the outcome is infidelity, the predictor consumption, and the mediator is commitment. The next part of the output shows us the results of the simple regression of commitment predicted from pornography consumption (i.e., path *a* in Figure 10.12). This output is interpreted just as we would interpret any regression: we can see that pornography consumption significantly predicts relationship commitment, $b = -0.47$, $t = -2.21$, $p = .028$. The R^2 value tells us that pornography consumption explains 2% of the variance in relationship commitment, and the fact that the *b* is negative tells us that the relationship is negative also: as consumption increases, commitment declines (and vice versa).

OUTPUT 10.4

```
**************************************************************************
Model = 4
    Y = Infideli
    X = LnConsum
    M = Commitme
Sample size
        239
**************************************************************************
Outcome: Commitme

Model Summary
          R       R-sq          F        df1        df2          p
      .1418      .0201     4.8633     1.0000   237.0000      .0284
Model
              coeff         se          t          p
constant     4.2027      .0545    77.1777      .0000
LnConsum     -.4697      .2130    -2.2053      .0284
```

Output 10.5 shows the results of the regression of infidelity predicted from both pornography consumption (i.e., path *c*' in Figure 10.12) and commitment (i.e., path *b* in Figure 10.12). We can see that pornography consumption significantly predicts infidelity even with relationship commitment in the model, $b = 0.46$, $t = 2.35$, $p = .02$; relationship commitment also significantly predicts infidelity, $b = -0.27$, $t = -4.61$, $p < .001$. The R^2 value tells us that the model explains 11.4% of the variance in infidelity. The negative *b* for commitment tells us that as commitment increases, infidelity declines (and vice versa), but the positive *b* for consumption indicates that as pornography consumption increases, infidelity increases also. These relationships are in the predicted direction.

OUTPUT 10.5

```
**************************************************************************
Outcome: Infideli
Model Summary
          R       R-sq          F        df1        df2          p
      .3383      .1144    15.2453     2.0000   236.0000      .0000
Model
              coeff         se          t          p
constant     1.3704      .2518     5.4433      .0000
Commitme     -.2710      .0587    -4.6128      .0000
LnConsum      .4573      .1946     2.3505      .0196
```

Output 10.6 shows the total effect of pornography consumption on infidelity (outcome). You will get this bit of the output only if you selected *Total effect model* in Figure 10.13. The total effect is the effect of the predictor on the outcome when the mediator is not present in the model – in other words, path *c* in Figure 10.11. When relationship commitment is not in the model, pornography consumption significantly predicts infidelity, $b = 0.58$, $t = 2.91$, $p = .004$. The R^2 value tells us that the model explains 3.46% of the variance in infidelity. As is the case when we include relationship commitment in the model, pornography consumption has a positive relationship with infidelity (as shown by the positive *b*-value).

OUTPUT 10.6

```
************************** TOTAL EFFECT MODEL ****************************
Outcome: Infideli

 Model Summary
          R       R-sq         F        df1        df2          p
      .1859      .0346    8.4866     1.0000   237.0000      .0039

Model
              coeff         se          t          p
constant      .2315      .0513     4.5123      .0000
LnConsum      .5846      .2007     2.9132      .0039
```

Output 10.7 is the most important part of the output because it displays the results for the indirect effect of pornography consumption on infidelity (i.e., the effect via relationship commitment). First, we're told the effect of pornography consumption on infidelity in isolation (the total effect), and these values replicate the model in Output 10.6. Next, we're told the effect of pornography consumption on infidelity when relationship commitment is included as a predictor as well (the direct effect). These values replicate those in Output 10.5. The first bit of new information is the *Indirect effect of X on Y*, which in this case is the indirect effect of pornography consumption on infidelity. We're given an estimate of this effect ($b = 0.127$) as well as a bootstrapped standard error and confidence interval. As we have seen many times before, 95% confidence intervals contain the true value of a parameter in 95% of samples. Therefore, we tend to assume that our sample isn't one of the 5% that does not contain the true value and use them to infer the population value of an effect. In this case, assuming our sample is one of the 95% that 'hits' the true value, we know that the true *b*-value for the indirect effect falls between 0.023 and 0.335.[8] This range does not include zero, and remember that $b = 0$ would mean 'no effect whatsoever'; therefore, the fact that the confidence interval does not contain zero means that there is likely to be a genuine indirect effect. Put another way, relationship commitment is a mediator of the relationship between pornography consumption and infidelity.

The rest of Output 10.7 you will see only if you selected *Effect size* in Figure 10.13; it contains various standardized forms of the indirect effect. In each case they are accompanied by a bootstrapped confidence interval. We discussed these measures of effect size in Section 10.4.3, and rather than interpret them all I'll merely note that for each one you get an estimate along with a confidence interval based on a bootstrapped standard error. As with the unstandardized indirect effect, if the confidence intervals don't contain zero then we can be confident that the true effect size is different from 'no effect'. In other words, there is mediation. All of the effect size measures have confidence intervals that don't include zero, so whatever one we look at we can be fairly confident that the indirect effect is greater than 'no effect'. Focusing on the most useful of these

[8] Remember that because of the nature of bootstrapping you will get slightly different values in your output.

effect sizes, the standardized *b* for the indirect effect, its value is $b = .041$, 95% BCa CI [.007, .103], and similarly, $\kappa^2 = .041$, 95% BCa CI [.008,.104]. κ^2 is bounded between 0 and 1, so we can interpret this as the indirect effect being about 4.1% of the maximum value that it could have been, which is a fairly small effect. We might, therefore, want to look for other potential mediators to include in the model in addition to relationship commitment.

OUTPUT 10.7

```
***************** TOTAL, DIRECT, AND INDIRECT EFFECTS ********************
Total effect of X on Y
    Effect         SE          t          p
     .5846      .2007     2.9132      .0039

Direct effect of X on Y
    Effect         SE          t          p
     .4573      .1946     2.3505      .0196

Indirect effect of X on Y
            Effect    Boot SE   BootLLCI   BootULCI
Commitme     .1273      .0716      .0232      .3350

Partially standardized indirect effect of X on Y
            Effect    Boot SE   BootLLCI   BootULCI
Commitme     .1818      .1002      .0325      .4684

Completely standardized indirect effect of X on Y
            Effect    Boot SE   BootLLCI   BootULCI
Commitme     .0405      .0220      .0073      .1032

Ratio of indirect to total effect of X on Y
            Effect    Boot SE   BootLLCI   BootULCI
Commitme     .2177     1.9048      .0348     1.4074

Ratio of indirect to direct effect of X on Y
            Effect    Boot SE   BootLLCI   BootULCI
Commitme     .2783     6.4664      .0222     6.7410

R-squared mediation effect size (R-sq_med)
            Effect    Boot SE   BootLLCI   BootULCI
Commitme     .0138      .0101      .0017      .0480

Preacher and Kelley (2011) Kappa-squared
            Effect    Boot SE   BootLLCI   BootULCI
Commitme     .0411      .0218      .0080      .1044
```

The final part of the output (Output 10.8) shows the results of the Sobel test. As I have mentioned before, it is better to interpret the bootstrap confidence intervals than formal tests of significance; however, if you selected *Sobel test* in Figure 10.13 this is what you will see. Again, we're given the size of the indirect effect ($b = 0.127$), the standard error, associated *z*-score ($z = 1.95$) and *p*-value ($p = .051$).[9] The *p*-value isn't quite under the not-at-all magic .05 threshold so technically we'd conclude that there isn't a significant indirect effect, but this just shows you how misleading these kind of tests can be: every single effect size had a confidence interval not containing zero, so there is compelling evidence that there is a small but meaningful mediation effect.

[9] You might remember in regression, we calculate a test statistic (*t*) by dividing the regression coefficient by its standard error (as in equation (8.11)). We do the same here except we get a *z* instead of a *t*: $z = 0.1273/0.0652 = 1.9526$.

OUTPUT 10.8

```
Normal theory tests for indirect effect
      Effect         se          Z          p
       .1273      .0652     1.9526      .0509
```

LABCOAT LENI'S REAL RESEARCH 10.1

I heard that Jane has a boil and kissed a tramp ②

MASSAR, K., ET AL. (2012). *PERSONALITY AND INDIVIDUAL DIFFERENCES*, 52, 106–109.

Everyone likes a good gossip from time to time, but apparently it has an evolutionary function. One school of thought is that gossip is used as a way to derogate sexual competitors – especially by questioning their appearance and sexual behaviour. For example, if you've got your eyes on a guy, but he has his eyes on Jane, then a good strategy is to spread gossip that Jane has a massive pus-oozing boil on her stomach and that she kissed a smelly vagrant called Aqualung. Apparently men rate gossiped-about women as less attractive, and they were more influenced by the gossip if it came from a woman with a high mate value (i.e., attractive and sexually desirable). Karlijn Massar and her colleagues hypothesized that if this theory is true then (1) younger women will gossip more because there is more mate competition at younger ages; and (2) this relationship will be mediated by the mate value of the person (because for those with high mate value gossiping for the purpose of sexual competition will be more effective). Eighty-three women aged from 20 to 50 (**Age**) completed questionnaire measures of their tendency to gossip (**Gossip**) and their sexual desirability (**Mate_Value**). Test Massar et al.'s mediation model using Baron and Kenny's method (as they did) but also using *PROCESS* to estimate the indirect effect (**Massar et al. (2011).sav**). Answers are on the companion website (or look at Figure 1 in the original article, which shows the parameters for the various regressions).

10.4.6. Reporting mediation analysis ②

Some people report only the indirect effect in mediation analysis, and possibly the Sobel test. However, I have repeatedly favoured using bootstrap confidence intervals, so you should report these, and preferably the effect size κ^2 and its confidence interval:

✓ There was a significant indirect effect of pornography consumption on infidelity through relationship commitment, b = 0.127, BCa CI [0.023, 0.335]. This represents a relatively small effect, κ^2 = .041, 95% BCa CI [.008, .104]

This is fine, but it can be quite useful to present a diagram of the mediation model, and indicate on it the regression coefficients, the indirect effect and its bootstrapped confidence intervals. For the current example, we might produce something like Figure 10.14.

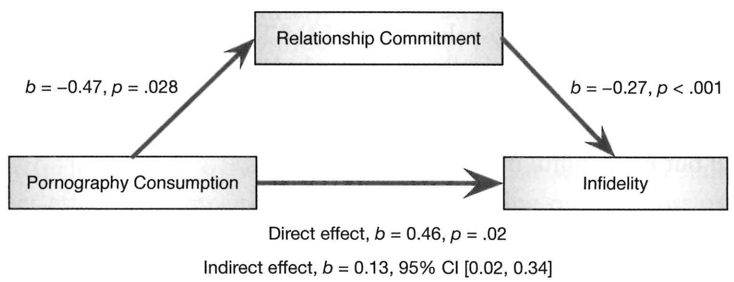

FIGURE 10.14 Model of pornography consumption as a predictor of infidelity, mediated by relationship commitment. The confidence interval for the indirect effect is a BCa bootstrapped CI based on 1000 samples

CRAMMING SAM'S TIPS Mediation

- Mediation is when the strength of the relationship between a predictor variable and outcome variable is reduced by including another variable as a predictor. Essentially, mediation equates to the relationship between two variables being 'explained' by a third. For example, the relationship between watching horror films and feeling scared at bedtime might be explained by scary images appearing in your head.
- Mediation is tested by assessing the size of the *indirect effect* and its confidence interval. If the confidence interval contains zero then we cannot be confident that a genuine mediation effect exists. If the confidence interval doesn't contain zero, then we can conclude that mediation has occurred.
- The size of the indirect effect can be expressed using kappa-squared (κ^2). Values of 0 mean that the indirect effect is very small relative to its maximum possible value, and values close to 1 mean that it is as large as it could possibly be given the research design. A small effect is .01, a medium effect would be around .09, and a large effect in the region of .25.

10.5. Categorical predictors in regression ③

We saw in the previous chapter that it is possible to include a categorical predictor in a regression model when there are only two categories: we simply code these categories with 0 and 1.[10] However, often you'll collect data about groups of people in which there are more than two categories (e.g., ethnic group, gender, socio-economic status, diagnostic category). You might want to include these groups as predictors in the regression model. Given that we have seen how to include categorical predictors with two categories into a regression model (Section 9.2.2), it shouldn't be too inconceivable that we could then extend this model to incorporate several predictors that had two categories; therefore, if we want to include a predictor with more than two categories, all we need to do is convert it to several variables each of which has two categories. This is the essence of dummy coding.

10.5.1. Dummy coding ③

10.5.1.1. What is dummy coding? ③

The obvious problem with wanting to use categorical variables as predictors is that often you'll have more than two categories. For example, if you'd collected data on religion you might have categories of Muslim, Jewish, Hindu, Catholic, Buddhist, Protestant, Jedi.[11] Clearly these groups cannot be distinguished using a single variable coded with zeros and ones. Therefore, we use what are called **dummy variables**, which is a way of representing groups of people using only zeros and ones. To do it, we have to create several variables; in fact, the number of variables we need is one less than the number of groups we're recoding.

[10] We saw in Section 9.2.2 why we use 0 and 1, and I elaborate on this issue in Section 11.2.1.

[11] For those of you not in the UK, we had a census here a few years back in which a significant portion of people put down Jedi as their religion.

There are eight basic steps:

1. Count the number of groups you want to recode and subtract 1.
2. Create as many new variables as the value you calculated in step 1. These are your dummy variables.
3. Choose one of your groups as a baseline against which all other groups will be compared. Normally you'd pick the control group, or, if you don't have a specific hypothesis, the group that represents the majority of people (because it might be interesting to compare other groups against the majority).
4. Having chosen a baseline group, assign that group values of 0 for all of your dummy variables.
5. For your first dummy variable, assign the value 1 to the first group that you want to compare against the baseline group. Assign all other groups 0 for this variable.
6. For the second dummy variable assign the value 1 to the second group that you want to compare against the baseline group. Assign all other groups 0 for this variable.
7. Repeat this process until you run out of dummy variables.
8. Place all of your dummy variables into the regression analysis in the same block.

Let's try this out using an example. In Chapter 5 we encountered a biologist who was worried about the potential health effects of music festivals. She collected some data at the Download Festival, which is a music festival specializing in heavy metal. The biologist was worried that the findings that she had were a function of the fact that she had tested only one type of person: metal fans. Perhaps it's only metal fans who get smellier at festivals (as a metal fan, I would at this point sacrifice the biologist to Odin for being so prejudiced). To find out whether the type of music a person likes predicts whether hygiene decreases over the festival, the biologist went to the Glastonbury Music Festival, which has an eclectic clientele. Again, she measured the hygiene of concertgoers over the three days of the festival using a technique that results in a score ranging between 0 (you smell like you've bathed in sewage) and 4 (you smell of freshly baked bread). The data are in the file called **GlastonburyFestivalRegression.sav**. This file contains the hygiene scores for each of three days of the festival as well as a variable called **change**, which is the change in hygiene over the three days of the festival (so it's the change from day 1 to day 3).[12] The biologist categorized people according to their musical affiliation: she used the label 'indie kid' for people who mainly like alternative music, 'metaller' for people who like heavy metal, and 'crusty' for people who like hippy/folky/ambient type of stuff. Anyone not falling into these categories was labelled 'no musical affiliation'. In the data file she coded these groups 1, 2, 3 and 4, respectively.

We have four groups, so there will be three dummy variables (one less than the number of groups). The first step is to choose a baseline group. We're interested in comparing those that have different musical affiliations against those that don't, so our baseline category will be 'no musical affiliation'. We give this group a code of 0 for all of our dummy variables. For our first dummy variable, we could look at the 'crusty' group, and to do this we give anyone who was a crusty a code of 1, and everyone else a code of 0. For our second dummy variable, we could look at the 'metaller' group, and to do this we give anyone who was a metaller a code of 1, and everyone else a code of 0. Our final dummy variable will code the 'indie kid' category. To do this, we give anyone who was an indie kid a code of 1, and everyone else a code of 0. The resulting coding scheme is shown in Table 10.2. Note that each group has a code of 1 on only one of the dummy variables (except the base category, which is always coded as 0).

TABLE 10.2 Dummy coding for the Glastonbury Festival data

	Dummy Variable 1	Dummy Variable 2	Dummy Variable 3
No Affiliation	0	0	0
Indie Kid	0	0	1
Metaller	0	1	0
Crusty	1	0	0

[12] Not everyone could be measured on day 3, so there is a change score only for a subset of the original sample.

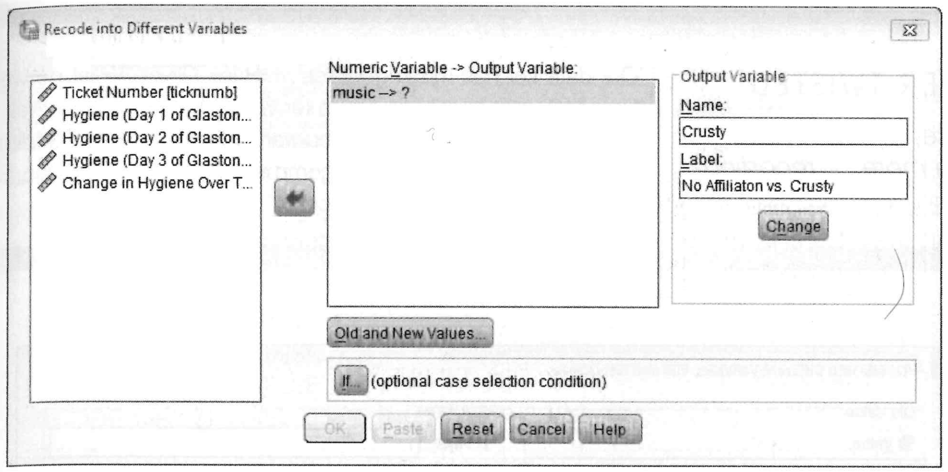

FIGURE 10.15 *Recode* dialog box

10.5.1.2. The *recode* function ③

We looked at why dummy coding works in Section 9.2.2, so let's look at how to recode our grouping variable into these dummy variables using SPSS. To recode variables you need to use the *recode* function. Select Transform ▸ Recode into Different Variables... to access the dialog box in Figure 10.15. The *Recode* dialog box lists all of the variables in the data editor, and you need to select the one you want to recode (in this case **music**) and transfer it to the box labelled *Numeric Variable → Output Variable* by clicking on ➡. You then need to name the new variable (the *Output Variable* as SPSS calls it) by going to the part labelled *Output Variable* and typing a name for your first dummy variable in the box labelled *Name* (let's call it **Crusty**). You can give this variable a more descriptive name by typing something in the box labelled *Label* (for this first dummy variable I've labelled it 'No Affiliation vs. Crusty'). Click on Change to transfer this new variable to the box labelled *Numeric Variable → Output Variable* (this box should now say *music → Crusty*).

Having defined the first dummy variable, we need to tell SPSS how to recode the values of the variable **music** into the values that we want for the new variable, **Crusty**. To do this, click on Old and New Values... to access the dialog box in Figure 10.16. This dialog box is used to change values of the original variable into different values for the new variable. For our first dummy variable, we want anyone who was a crusty to get a code of 1 and everyone else to get a code of 0. Now, crusty was coded with the value 3 in the original variable, so you need to type the value 3 in the section labelled *Old Value* in the box labelled *Value*. The new value we want is 1, so we need to type the value 1 in the section labelled *New Value* in the box labelled *Value*. When you've done this, click on Add to add this change to the list of changes (the list is displayed in the box labelled *Old → New*, which should now say *3 → 1* as in the diagram). The next thing we need to do is to change the remaining groups to have a value of 0 for the first dummy variable. To do this just select ◉ All other values and type the value 0 in the section labelled *New Value* in the box labelled *Value*.[13] When you've

[13] Using this ◉ All other values option is fine when you don't have missing values in the data, but just note that when you do (as is the case here) cases with both system-defined and user-defined missing values will be included in the recode. One way around this is to recode only cases for which there is a value (see Oliver Twisted). The alternative is to recode missing values specifically using the ◉ Range: option. It is also a good idea to use the *frequencies* or *crosstabs* commands after a recode and check that you have caught all of these missing values.

OLIVER TWISTED

Please, Sir, can I have some more … recoding?

'Our data set has missing values', worries Oliver. 'What do we do if we only want to recode cases for which we have data?' Well, we can set some other options. If you want to know more, the additional material for this chapter on the companion website will tell you. Stop worrying, Oliver, everything will be OK.

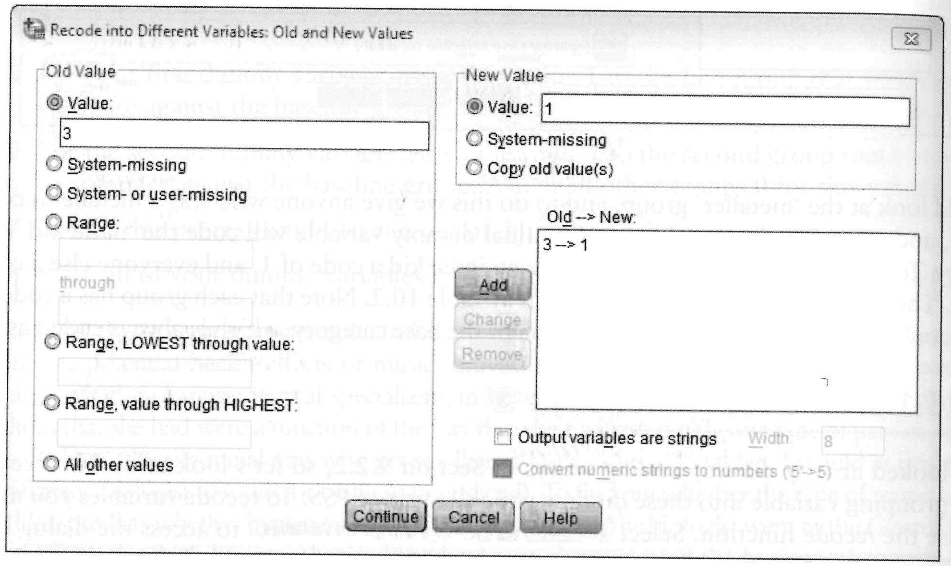

FIGURE 10.16 *Recode* dialog box for changing old values to new (see also SPSS Tip 10.2)

done this, click on Add to add this change to the list of changes (this list will now also say *ELSE* → *0*). When you've done this, click on Continue to return to the main dialog box, and then click on OK to create the first dummy variable. This variable will appear as a new column in the data editor, and you should notice that it will have a value of 1 for anyone originally classified as a crusty and a value of 0 for everyone else.

SELF-TEST Try creating the remaining two dummy variables (call them **Metaller** and **Indie_Kid**) using the same principles.

10.5.2. SPSS output for dummy variables ③

Let's assume you've created the three dummy coding variables (if you're stuck there is a data file called **GlastonburyDummy.sav** (the 'Dummy' refers to the fact it has dummy variables in it – I'm not implying that if you need to use this file you're a dummy☺). With dummy variables, you have to enter all related dummy variables in the same block (so use the *Enter* method).

SPSS TIP 10.2 Using syntax to recode ③

If you're doing a lot of recoding it soon becomes pretty tedious using the dialog boxes all of the time. I've written the syntax file, **RecodeGlastonburyData.sps**, to create all of the dummy variables we've discussed. Load this file and run the syntax, or type the following into a new syntax window (see Section 3.9):

```
DO IF(1-MISSING(change)).
RECODE music (3=1)(ELSE = 0) INTO Crusty.
RECODE music (2=1)(ELSE = 0) INTO Metaller.
RECODE music (1=1)(ELSE = 0) INTO Indie_Kid.
END IF.
VARIABLE LABELS Crusty 'No Affiliation vs. Crusty'.
VARIABLE LABELS Metaller 'No Affiliation vs. Metaller'.
VARIABLE LABELS Indie_Kid 'No Affiliation vs. Indie Kid'.
VARIABLE LEVEL Crusty Metaller Indie_Kid (Nominal).
FORMATS Crusty Metaller Indie_Kid (F1.0).
EXECUTE.
```

Each *recode* command does the equivalent of the dialog box in Figure 10.16. So, the three lines beginning *recode* ask SPSS to create three new variables (**Crusty**, **Metaller** and **Indie_Kid**), which are based on the original variable **music**. For the first variable, if **music** is 3 then it becomes 1, and every other value becomes 0. For the second, if **music** is 2 then it becomes 1, and every other value becomes 0, and so on for the third dummy variable. Note that all of these *recode* commands are within an *if* statement (beginning *do if* and ending with *end if*). This tells SPSS to carry out the *recode* commands only if a certain condition is met. The condition we have set is *(1- MISSING(change))*. *MISSING* is a built-in command that returns 'true' (i.e., the value 1) for a case that has a system- or user-defined missing value for the specified variable; it returns 'false' (i.e., the value 0) if a case has a value. Hence, *MISSING(change)* returns a value of 1 for cases that have a missing value for the variable **change** and 0 for cases that do have values. We want to recode the cases that *do* have a value for the variable change, therefore we use '*1–MISSING(change)*'. This command reverses *MISSING(change)* so that it returns 1 (true) for cases that have a value for the variable **change** and 0 (false) for system- or user-defined missing values. To sum up, the statement *DO IF (1–MISSING(change))* tells SPSS 'Do the following *recode* commands if the case has a value for the variable **change**.'

The *variable labels* command tells SPSS to assign the text in the quotations as labels for the variables **Crusty**, **Metaller**, and **Indie_Kid**, respectively. It then sets these three variables to be 'nominal', and the *formats* command changes the variables to have a width of 1 and 0 decimal places (hence the 1.0). The *execute* is essential: without it none of the commands beforehand will be executed. Note also that every line ends with a full stop.

SELF-TEST Use what you learnt in Chapter 8 to run a multiple regression using the change scores as the outcome, and the three dummy variables (entered in the same block) as predictors.

Let's have a look at the output. Output 10.9 shows the model statistics. We see that by entering the three dummy variables we can explain 7.6% of the variance in the change in hygiene scores (the R^2 value × 100%). In other words, 7.6% of the variance in the change in hygiene can be explained by the musical affiliation of the person. The ANOVA (which shows the same thing as the R^2 change statistic because there is only one step in this regression) tells us that the model is significantly better at predicting the change in hygiene scores

OUTPUT 10.9

Model Summary

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate	Change Statistics: R Square Change	F Change	df1	df2	Sig. F Change
1	.276[a]	.076	.053	.68818	.076	3.270	3	119	.024

a. Predictors: (Constant), No Affiliation vs. Indie Kid, No Affiliation vs. Crusty, No Affiliation vs. Metaller

b. Predictors: (Constant), No Affiliation vs. Indie Kid, No Affiliation vs. Metaller, No Affiliation vs. Crusty

ANOVA[a]

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	4.646	3	1.549	3.270	.024[b]
	Residual	56.358	119	.474		
	Total	61.004	122			

a. Dependent Variable: Change in Hygiene Over The Festival

b. Predictors: (Constant), No Affiliation vs. Indie Kid, No Affiliation vs. Crusty, No Affiliation vs. Metaller

c. Predictors: (Constant), No Affiliation vs. Indie Kid, No Affiliation vs. Metaller, No Affiliation vs. Crusty

than having no model (put another way, the 7.6% of variance that can be explained is a significant amount).

Output 10.10 shows a basic *Coefficients* table for the dummy variables, which is the more interesting part of the output. The first thing to notice is that each dummy variable appears in the table with a useful label (such as No Affiliation vs. Crusty) because when we recoded our variables we gave each variable a useful label; if we hadn't done this then the table would contain the less helpful variable names of Crusty, Metaller and Indie_Kid. The labels that I have used remind me of what each dummy variable represents. The first dummy variable (No Affiliation vs. Crusty) shows the difference between the change in hygiene scores for the no affiliation group and the crusty group. Remember that the beta value tells us the change in the outcome due to a unit change in the predictor. In this case, a unit change in the predictor is the change from 0 to 1. By including all three dummy variables at the same time, zero will represent our baseline category (no affiliation). For this variable 1 represents 'Crusty'. Therefore, the change from 0 to 1 represents the change from no affiliation to Crusty. Therefore, this variable represents the difference in the change in hygiene scores for a crusty, relative to someone with no musical affiliation. This difference is the difference between the two group means (see Section 9.2.2).

To illustrate this fact, I've produced a table (Output 10.11) of the group means for each of the four groups and also the difference between the means for each group and the *no affiliation* group. These means represent the average change in hygiene scores for the three groups (i.e., the mean of each group on our outcome variable). If we calculate the difference in these means for the no affiliation group and the crusty group we get, crusty − no affiliation = $(-0.966) - (-0.554) = -0.412$. In other words, the change in hygiene scores is greater for the crusty group than it is for the no affiliation group (crusties' hygiene decreases more over the festival than those with no musical affiliation). This value is the same as the *unstandardized* beta value in Output 10.10. So, the beta values tell us the relative difference between each group and the group that we chose as a baseline category. This beta value is converted to a t-statistic and the significance of this t reported. As we've seen before this t-statistic tests whether the beta value is 0; therefore, when we have two categories coded with 0 and 1, it tests whether the difference between group means is 0. If it is significant then the group coded with 1 is significantly different from the baseline category – so, it's testing the difference between two means, which is the context in which students are most familiar with the t-statistic (see Chapter 9). For our first dummy variable, the t-test is significant, and the beta value has a negative value so we could say that the change in hygiene scores goes down as a person changes from having no affiliation to being a crusty. Bear in mind that a decrease in hygiene scores represents greater change (you're becoming smellier) so what this actually means is that hygiene decreased significantly more in crusties compared to those with no musical affiliation.

OUTPUT 10.10

Coefficients[a]

Model		Unstandardized Coefficients: B	Std. Error	Standardized Coefficients: Beta	t	Sig.	95.0% Confidence Interval for B: Lower Bound	Upper Bound
1	(Constant)	-.554	.090		-6.134	.000	-.733	-.375
	No Affiliation vs. Crusty	-.412	.167	-.232	-2.464	.015	-.742	-.081
	No Affiliation vs. Metaller	.028	.160	.017	.177	.860	-.289	.346
	No Affiliation vs. Indie Kid	-.410	.205	-.185	-2.001	.048	-.816	-.004

a. Dependent Variable: Change in Hygiene Over The Festival

Bootstrap for Coefficients

Model		B	Bootstrap[a]: Bias	Std. Error	Sig. (2-tailed)	BCa 95% Confidence Interval: Lower	Upper
1	(Constant)	-.554	.005	.097	.001	-.736	-.349
	No Affiliation vs. Crusty	-.412	-.011	.179	.030	-.733	-.101
	No Affiliation vs. Metaller	.028	-.006	.149	.847	-.262	.293
	No Affiliation vs. Indie Kid	-.410	-.010	.201	.049	-.813	-.043

a. Unless otherwise noted, bootstrap results are based on 1000 bootstrap samples

OUTPUT 10.11

OLAP Cubes

Variables=Change in Hygiene Over The Festival

Musical Affiliation	Mean	Std. Deviation	N
Indie Kid	-0.964	0.670	14
Metaller	-0.526	0.576	27
Crusty	-0.966	0.760	24
No Musical Affiliation	-0.554	0.708	58
Crusty - No Musical Affiliation	-0.412	0.052	-34
Metaller - No Musical Affiliation	0.028	-0.133	-31
Indie Kid - No Musical Affiliation	-0.410	-0.038	-44
Total	-0.675	0.707	123

Our next dummy variable compares metallers to those that have no musical affiliation. The beta value again represents the difference in the change in hygiene scores for a person with no musical affiliation compared to a metaller. The difference in the group means for the no affiliation group and the metaller group is metaller − no affiliation = $(-0.526) - (-0.554) = 0.028$. This value is again the same as the unstandardized beta value in Output 10.10. For this second dummy variable, the t-test is not significant. We could conclude that the change in hygiene scores is similar if a person changes from having no affiliation to being a metaller: the change in hygiene scores is not predicted by whether someone is a metaller compared to if they have no musical affiliation.

For the final dummy variable, we're comparing indie kids to those that have no musical affiliation. The beta value again represents the shift in the change in hygiene scores if a person has no musical affiliation, compared to someone who is an indie kid. The difference in the group means for the no affiliation group and the indie kid group is indie kid − no affiliation = $(-0.964) - (-0.554) = -0.410$. It should be no surprise to you by now that this is the unstandardized beta value in Output 10.10. The t-test is significant, and the beta value has a negative value so, as with the first dummy variable, we could say that the change in hygiene scores goes down as a person changes from having no affiliation to being an indie kid. Bear in mind that a decrease in hygiene scores represents more change (you're becoming smellier) so this actually means that hygiene decreased significantly more in indie kids compared to those with no musical affiliation. We could report the results as in Table 10.3 (note I've included the bootstrap confidence intervals).

So, overall this analysis has shown that compared to having no musical affiliation, crusties and indie kids get significantly smellier across the three days of the festival, but

TABLE 10.3 Linear model of predictors of the change in hygiene scores (95% bias corrected and accelerated confidence intervals reported in parentheses). Confidence intervals and standard errors based on 1000 bootstrap samples

	b	SE B	β	p
Constant	−0.55 (−0.74, −0.35)	0.10		p = .001
No Affiliation vs. Crusty	−0.41 (−0.73, −0.10)	0.18	−.23	p = .030
No Affiliation vs. Metaller	0.03 (−0.26, 0.29)	0.15	.02	p = .847
No Affiliation vs. Indie Kid	−0.41 (−0.81, −0.04)	0.20	−.19	p = .049

Note. $R^2 = .08$ ($p = .024$).

metallers don't. This section has introduced some really complex ideas that I expand upon in Chapter 11. It might all be a bit much to take in, and so if you're confused or want to know more about why dummy coding works in this way I suggest reading Section 11.2.1 and then coming back here. Alternatively, read Hardy's (1993) excellent monograph.

10.6. Brian's attempt to woo Jane ①

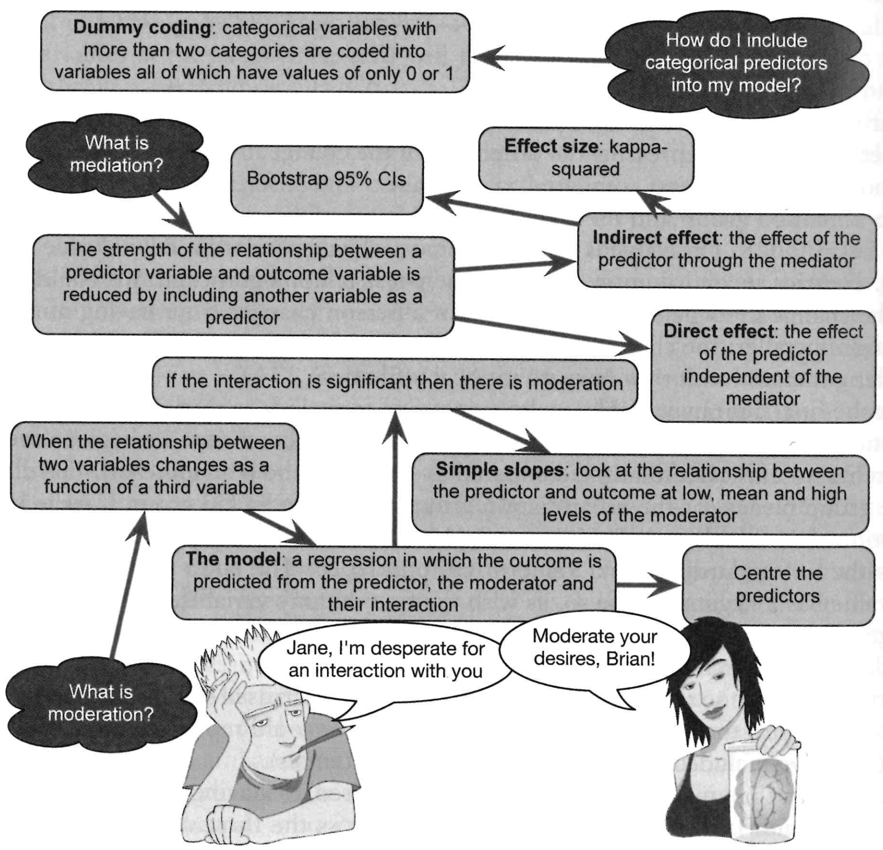

FIGURE 10.17 What Brian learnt from this chapter

10.7. What next? ①

We started this chapter by looking at my relative failures as a human being compared to Simon Hudson. I then bleated on excitedly about moderation and mediation, which could explain why Clair Sparks chose Simon Hudson all those years ago. Perhaps she could see the writing on the wall! I was true to my word to my parents, though, and I was philosophical about it. I set my sights elsewhere during the obligatory lunchtime game of kiss chase. However, my life was about to change beyond all recognition. Not that I believe in fate, but if I did I would have believed that the wrinkly and hairy hand of fate (I don't know why but I always imagine it wrinkly, hairy and in need of a manicure) had decided that I was far too young to be getting distracted by such things as girls. Waggling its finger at me, it plucked me out of primary school and cast me down into what can only be described as hell, also known as an all-boys' school. It's fair to say that my lunchtime primary school game of kiss chase was the last I would see of girls for quite some time …

10.8. Key terms that I've discovered

Grand mean centring
Direct effect
Index of mediation
Indirect effect
Interaction effect
Mediation
Mediator
Moderation
Moderator
Simple slopes analysis
Sobel test

10.9. Smart Alex's tasks

- **Task 1:** McNulty et al. (2008) found a relationship between a person's **Attractiveness** and how much **Support** they give their partner as newlyweds. Is this relationship moderated by gender (i.e., whether the data were from the husband or wife)? The data are in **McNulty et al. (2008).sav.**[14] ②
- **Task 2:** Produce the simple slopes graphs for the above example. ②
- **Task 3:** McNulty et al. (2008) also found a relationship between a person's **Attractiveness** and their relationship **Satisfaction** as newlyweds. Using the same data as the previous examples, is this relationship moderated by gender? ②
- **Task 4:** In the chapter we tested a mediation model of infidelity for Lambert et al.'s data using Baron and Kenny's regressions. Repeat this analysis, but using **Hook_Ups** as the measure of infidelity. ②
- **Task 5:** Repeat the above analysis but using the *PROCESS* tool to estimate the indirect effect and its confidence interval. ②
- **Task 6:** In Chapter 3 (Task 5) we looked at data from people who had been forced to marry goats and dogs and measured their life satisfaction as well as how much they like animals (**Goat or Dog.sav**). Run a regression predicting life satisfaction from the type of animal to which a person was married. Write out the final model. ②

[14] These are not the actual data from the study, but are simulated to mimic the findings in Table 1 of the original paper.

- **Task 7:** Repeat the analysis above but include animal liking in the first block, and type of animal in the second block. Do your conclusions about the relationship between type of animal and life satisfaction change? ②
- **Task 8:** Using the **GlastonburyDummy.sav** data, which you should've already analysed, comment on whether you think the model is reliable and generalizable. ③
- **Task 9:** Tablets like the iPad are very popular. A company owner was interested in how to make his brand of tablets more desirable. He collected data on how cool people perceived a product's advertising to be (**Advert_Cool**), how cool they thought the product was (**Product_Cool**), and how desirable they found the product (**Desirability**). Test his theory that the relationship between cool advertising and product desirability is mediated by how cool people think the product is (**Tablets.sav**). Am I showing my age by using the word 'cool'? ③

Answers can be found on the companion website.

10.10. Further reading

Cohen, J., Cohen, P., Aiken, L., & West, S. (2003). *Applied multiple regression/correlation analysis for the behavioral sciences*. Mahwah, NJ: Erlbaum.

Hardy, M. A. (1993). *Regression with dummy variables*. Sage University Paper Series on Quantitative Applications in the Social Sciences, 07-093. Newbury Park, CA: Sage.

Hayes, A. F. (2013). *An introduction to mediation, moderation, and conditional process analysis*. New York: Guilford Press.

Comparing several means: ANOVA (GLM 1)

11

FIGURE 11.1 My brother Paul (left) and I (right) in our very fetching school uniforms

11.1. What will this chapter tell me? ①

There are pivotal moments in everyone's life, and one of mine was at the age of 11. Where I grew up in England there were three choices when leaving primary school and moving on to secondary school: (1) state school (where most people go); (2) grammar school (where clever people who pass an exam called the Eleven Plus go); and (3) private school (where rich people go). My parents were not rich and I am not clever and consequently I failed my Eleven Plus, so private school and grammar school (where my clever older brother had gone) were out. This left me to join all of my friends at the local state school. I could not have been happier. Imagine everyone's shock when my parents received a letter saying that some extra spaces had become available at the grammar school; although the local authority could scarcely believe it and had checked the Eleven Plus papers several million times to confirm their findings, I was next on their list. I could not have been unhappier. So, I waved goodbye to all of my friends and trundled off to join my brother at Ilford County High School for Boys (a school that still hit students with a cane if they were particularly bad and that, for