but there are other things too such as other medication that suppresses libido (such as antidepressants or the contraceptive pill, both of which lower libido) and fatigue. If these variables (called covariates) are measured, then it is possible to control for the influence they have on the dependent variable by including them in the regression model. From what we know of hierarchical regression (see Chapter 5) it should be clear that if we enter the covariate into the regression model first, and then enter the dummy variables representing the experimental manipulation, we can see what effect an independent variable has *after* the effect of the covariate. As such, we control for (or partial out) the effect of the covariate. There are two reasons for including covariates in ANOVA:

- **To reduce within-group error variance**: In the discussion of ANOVA and *t*-tests we got used to the idea that we assess the effect of an experiment by comparing the amount of variability in the data that the experiment can explain against the variability that it cannot explain. If we can explain some of this 'unexplained' variance (SS_R) in terms of other variables (covariates), then we reduce the error variance, allowing us to assess more accurately the effect of the independent variable (SS_M).
- **Elimination of confounds**: In any experiment, there may be unmeasured variables that confound the results (i.e. variables that vary systematically with the experimental manipulation). If any variables are known to influence the dependent variable being measured, then ANCOVA is ideally suited to remove the bias of these variables. Once a possible confounding variable has been identified, it can be measured and entered into the analysis as a covariate.

There are other reasons for including covariates in ANOVA but because I do not intend to describe the computation of ANCOVA I recommend that the interested reader consult Wildt & Ahtola (1978) or Stevens (1992, Chapter 9).

Imagine that the researcher who conducted the Viagra study in the previous chapter suddenly realized that the libido of the participants' sexual partners would affect the participants' own libido (especially because the measure of libido was behavioural). Therefore, they repeated the study on a different set of participants, but this time took a measure of the partner's libido. The partner's libido was measured in terms of how often they tried to initiate sexual contact. In the previous chapter, we saw that this experimental scenario could be characterized in terms of equation (8.2). Think back to what we know about multiple regression (Chapter 5) and you can hopefully see that this equation can be extended to include this covariate (see equation (9.1)):

$$\begin{aligned}\text{Libido}_i &= b_0 + b_3\text{Covariate}_i + b_2\text{High}_i + b_1\text{Low}_i + \varepsilon_i \\ \text{Libido}_i &= b_0 + b_3\text{Partner's Libido}_i + b_2\text{High}_i + b_1\text{Low}_i + \varepsilon_i\end{aligned} \tag{9.1}$$

9.3. CONDUCTING ANCOVA ON SPSS ②

The data for this example are in Table 9.1 and can be found in the file **ViagraCovariate. sav**. Table 9.1 shows the participant's libido, their partner's libido, and the means (and standard deviations in brackets) of the various scores. I recommend putting these data into the data

Table 9.1 Data from **ViagraCovariate.sav**

Dose	Participant's Libido	Partner's Libido
Placebo	3	5
	2	2
	5	6
	2	2
	2	3
	2	3
	7	7
	2	4
$\bar{X}$	4	5
S	3.22	4.11
	(1.79)	(1.76)
Low Dose	7	10
	5	8
	3	6
	4	7
	4	7
	7	11
	5	9
	4	7
$\bar{X}$	4.88	8.13
S	(1.46)	(1.73)
High Dose	9	2
	2	3
	6	5
	3	4
	4	3
	4	3
	4	2
	6	0
	4	1
	6	3

(Continued)

Table 9.1 (Continued)

2	0
8	1
5	0
4.85	2.08
(2.12)	(1.61)

editor by hand. This can be done in much the same way as the Viagra data from the previous chapter except that an extra variable must be created in which to place the values of the covariate.

9.3.1. Inputting data ①

In essence, the data should be laid out in the data editor as they are in Table 9.1 (excluding the rows for the means and standard deviations). So, create a coding variable called **dose** and use the *Labels* option to define value labels (as in Chapter 8 I recommend 1 = placebo, 2 = low dose, 3 = high dose). There were different numbers of participants in each condition, so you need to enter nine values of 1 into this column (so that the first nine rows contain the value 1), followed by eight rows containing the value 2, and followed by 14 rows containing the value of 3. At this point, you should have one column with 30 rows of data entered. Next, create a second variable called **libido** and enter the 30 scores that correspond to the person's libido. Finally, create a third variable called **partner** and use the *Labels* option to give this variable a more descriptive title of 'Partner's libido'. Then, enter the 30 scores that correspond to the partner's libido. Remember that the means (and standard deviations) that I've included in the table are not required by SPSS!

9.3.2. Main analysis ②

Most of the factorial ANOVA procedures in SPSS include the facility to include one or more covariates. However, for simpler designs (most designs that don't involve repeated measures) it is probably best to conduct ANCOVA via the general factorial procedure. To access the main dialog box follow the menu path **Analyze⇒General Linear Model ⇒Univariate...** (see Figure 9.1).[1] The main dialog box is similar to that for one-way ANOVA, except that there is a space to specify covariates. Select **libido** and place this in the box labelled *Dependent Variable* by clicking on ▸. Select **dose** and transfer it to the box labelled *Fixed Factor(s)* and then select **partner** and transfer it to the box labelled *Covariate(s)*.

1 **Statistics⇒General Linear Model⇒GLM—General Factorial...** in version 8.0 and earlier.

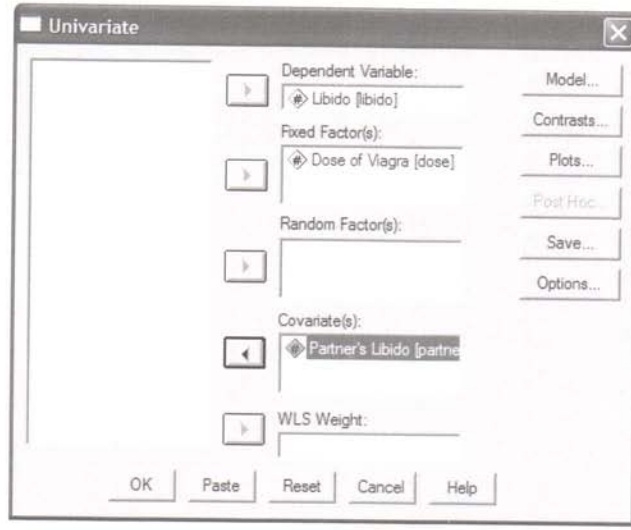

Figure 9.1 Main dialog box for GLM univariate

9.3.3. Contrasts and other options ②

There are various dialog boxes that can be accessed from the main dialog box. The first thing to notice is that if a covariate is selected, the *post hoc* tests are disabled (you cannot access this dialog box). *Post hoc* tests are not designed for situations in which a covariate is specified; however, some comparisons can still be done using contrasts.

Click on Contrasts... to access the *contrasts* dialog box. This dialog box is different to the one we met in Chapter 8 in that you cannot enter codes to specify particular contrasts. Instead, you can specify one of several standard contrasts. These standard contrasts were listed in Table 8.6. In this example, there was a placebo control condition (coded as the first group), so a sensible set of contrasts would be simple contrasts comparing each experimental group with the control. The default contrast in SPSS is a deviation contrast and to change this we must first click on ▾ next to the box labelled *Contrast*. A list of contrasts will drop down and you should select a type of contrast (in this case *Simple*) from this list and the list will automatically disappear. For simple contrasts you have the option of specifying a reference category (which is the category against which all other groups are compared). By default the reference category is the last category: because in this case the control group was the first category (assuming that you coded placebo as 1) we need to change this option by selecting ◉ First. When you have selected a new contrast, you must click on Change to register this change. The final dialog box should look like Figure 9.2 Click on Continue to return to the main dialog box.

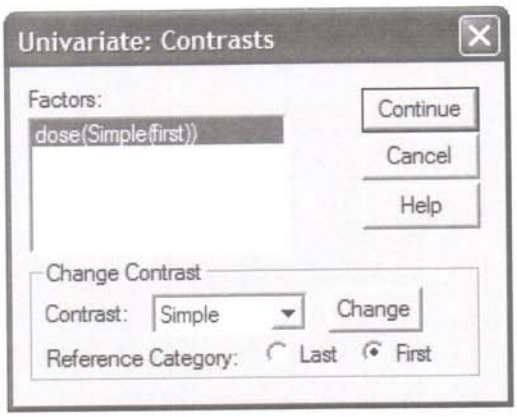

Figure 9.2 Options for standard contrasts in GLM univariate

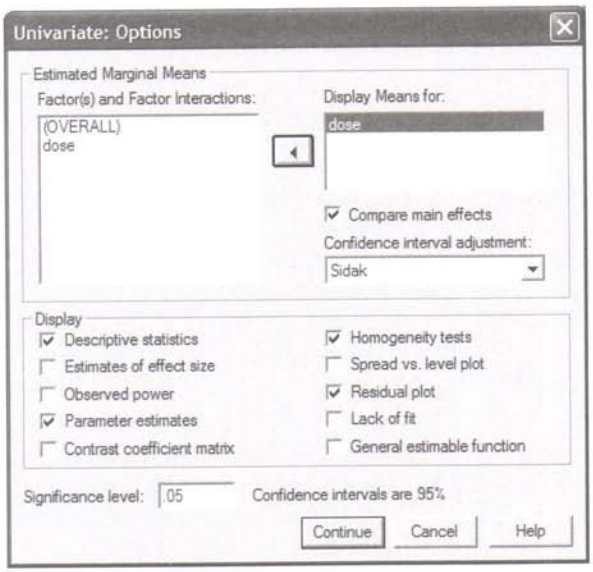

Figure 9.3 *Options* dialog box for GLM univariate

Another way to get *post hoc* tests is by clicking on Options... to access the *options* dialog box (see Figure 9.3). To specify *post hoc* tests, select the independent variable (in this case **dose**) from the box labelled *Estimated Marginal Means: Factor(s) and Factor Interactions* and transfer it to the box labelled *Display Means for* by clicking on ▶. Once a variable has been transferred, the box labelled *Compare main effects* becomes active and you should select this option (Compare main effects). If this option is selected, the box labelled

Confidence interval adjustment becomes active and you can click on ▼ to see a choice of three adjustment levels. The default is to have no adjustment and simply perform a Tukey LSD *post hoc* test (this option is not recommended); the second is to ask for a Bonferroni correction (recommended); the final option is to have a Sidak correction. The Sidak correction is similar to the Bonferroni correction but is less conservative and so should be selected if you are concerned about the loss of power associated with Bonferroni corrected values. For this example use the Sidak correction (we will use Bonferroni later in the chapter). As well as producing *post hoc* tests for the **dose** variable, placing **dose** in the *Display Means for* box will result in a table of estimated marginal means for this variable. These means provide an estimate of the *adjusted* group means (i.e. the means after the covariate has been accounted for). When you have selected the options required (see Box 9.1), click on Continue to return to the main dialog box. There are other options available from the main dialog box. For example, if you have several independent variables you can plot them against each other (which is useful for interpreting interaction effects—see section 10.3.3). For this analysis, there is only one independent variable and so we can click on OK to run the analysis.

Box 9.1

Options for ANCOVA ②

The remaining options in this dialog box are as follows:

- **Descriptive statistics**: This option produces a table of means and standard deviations for each group.
- **Estimates of effect size**: This option produces the value of eta squared (η^2) described in section 8.5, which is a measure of the size of experimental effect. In fact, eta squared is the regression coefficient (R^2) for a non-linear regression line (i.e. a curve) assumed to pass through all group means. In a population this assumption is true, but in samples it is not: therefore, eta squared is usually biased (see section 8.5 and Howell, 2002, section 11.11). For this reason there is little to recommend this option, and the effect size may be estimated more productively using omega squared (ω^2)—see section 9.7.
- **Observed power**: This option provides an estimate of the probability that the statistical test could detect the difference between the observed group means (see section 1.8.5). This measure is of little use because if the *F*-test is significant then the probability that the effect was detected will, of course, be high. Likewise, if group differences were small, the observed power will be low. Observed power is of little use and I would advise that power calculations (with regard to sample size) are made before the experiment is conducted (see Cohen, 1988, 1992; Howell, 2002, for ideas on how to do this by hand; Field, 1998b, for ideas on doing it using a computer; or use the free software G*Power available from the CD-ROM of this book and http://www.psycho.uni-duesseldorf.de/aap/projects/gpower/).

(Continued)

Box 9.1 (Continued)

- **Parameter estimates**: This option produces a table of regression coefficients and their tests of significance for the variables in the regression model (see section 9.5).
- **Contrast coefficient matrix**: This option produces matrices of the coding values used for any contrasts in the analysis. This option is useful only for checking which groups are being compared in which contrast.
- **Homogeneity tests**: This option produces Levene's test of the homogeneity of variance assumption (see sections 3.6 and 8.4.1).
- **Spread vs. level plot**: This option produces a chart that plots the mean of each group of a factor (*X*-axis) against the standard deviation of that group (*Y*-axis). This is a useful plot to check that there is no relationship between the mean and standard deviation. If a relationship exists then the data may need to be stabilized using a logarithmic transformation (see Chapter 3).
- **Residual plot**: This option produces plots of observed-by-predicted-by-standardized residual values. These plots can be used to assess the assumption of equality of variance.

9.4. INTERPRETING THE OUTPUT FROM ANCOVA ②

9.4.1. Main analysis ②

SPSS Output 9.1 shows (for illustrative purposes) the ANOVA table for these data when the covariate is not included. It is clear from the significance value, which is greater than .05, that there are no differences in libido between the three groups; therefore Viagra seems to have no significant effect on libido. It should also be noted that the total amount of variation to be explained (SS_T) was 110.97 (*Corrected Total*), of which the experimental manipulation accounted for 16.84 units (SS_M), whilst 94.12 were unexplained (SS_R).

Tests of Between-Subjects Effects

Dependent Variable: Libido

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	16.844[a]	2	8.422	2.416	.108
Intercept	535.184	1	535.184	153.522	.000
DOSE	16.844	2	8.422	2.416	.108
Error	94.123	27	3.486		
Total	683.000	30			
Corrected Total	110.967	29			

a R Squared = .152 (Adjusted R Squared = .089)

SPSS Output 9.1

SPSS Output 9.2 shows the results of Levene's test and the ANOVA table when partner's libido is included in the model as a covariate. Levene's test is significant, indicating that the group variances are not equal (hence the assumption of homogeneity of variance has been violated). However, as I've mentioned in section 3.6, Levene's test is not necessarily the best way to judge whether variances are unequal enough to cause problems. A good double-check is to look at the highest and lowest variances. For our three groups we have standard deviations of 1.79 (placebo), 1.46 (low dose) and 2.12 (high dose)—see Table 9.1. If we square these values we get variances of 3.20 (placebo), 2.13 (low dose) and 4.49 (high dose). We then take the largest variance and divide it by the smallest: in this case 4.49/2.13 = 2.11. If the resulting value is less than 2 then we probably don't need to worry too much; if it's greater than 2 (as it is here) then we probably do! However, for the time being don't worry too much about the differences in variances.

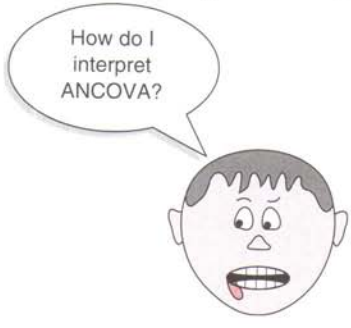

The format of the ANOVA table is largely the same as without the covariate, except that there is an additional row of information about the covariate (**partner**). Looking first at the significance values, it is clear that the covariate significantly predicts the dependent variable, because the significance value is less than .05. Therefore, the person's libido is influenced by their partner's libido. What's more interesting is that when the effect of partner's libido is removed, the effect of Viagra becomes significant (p is .016 which is less than .05). The amount of variation accounted for by the model (SS_M) has increased to 34.75 units (*corrected model*) of which Viagra accounts for 28.34 units. Most important, the large amount of variation in libido that is accounted for by the covariate has meant that the unexplained variance (SS_R) has been reduced to 76.22 units. Notice that SS_T has not changed; all that has changed is how that total variation is explained.

This example illustrates how ANCOVA can help us to exert stricter experimental control by taking account of confounding variables to give us a 'purer' measure of effect of the experimental manipulation. Without taking account of the libido of the participants' partners we would have concluded that Viagra had no effect on libido, yet clearly it does. However, the effect of the partner's libido seems stronger than that of Viagra. Looking back at the group means from Table 9.1 for the libido data, it seems pretty clear that the significant ANOVA reflects a difference between the placebo group and the two experimental groups (because the low- and high-dose groups have very similar means—4.88 and 4.85—whereas the placebo group mean is much lower at 3.22). However, we'll need to check some contrasts to verify this.

SPSS Output 9.3 shows the parameter estimates selected in the *options* dialog box. These estimates are calculated using a regression analysis with **dose** split into two dummy coding variables (see section 8.2.2 and section 9.5). SPSS codes the two dummy variables such that the last category (the category coded with the highest value in the data editor—in this case the high dose group) is the reference category. This reference category (labelled DOSE=3 in the output) is coded with a 0 for both dummy variables (see section 8.2.2 for a reminder of how dummy coding works). DOSE=2, therefore, represents the difference between the group coded as 2 (low dose) and the reference category

Levene's Test of Equality of Error Variances[a]

Dependent Variable: Libido

F	df1	df2	Sig.
5.525	2	27	.010

Tests the null hypothesis that the error variance of the dependent variable is equal across groups.
a Design: Intercept+ PARTNER+DOSE

Tests of Between-Subjects Effects

Dependent Variable: Libido

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	34.750[a]	3	11.583	3.952	.019
Intercept	12.171	1	12.171	4.152	.052
PARTNER	17.906	1	17.906	6.109	.020
DOSE	28.337	2	14.169	4.833	.016
Error	76.216	26	2.931		
Total	683.000	30			
Correlated Total	110.967	29			

a R Squared = .313 (Adjusted R Squared = .234)

SPSS Output 9.2

(high dose), and DOSE=1 represents the difference between the group coded as 1 (placebo) and the reference category (high dose). The *b*-values literally represent the differences between the means of these groups and so the significances of the *t*-tests tell us whether the group means differ significantly. The degrees of freedom for these *t*-tests can be calculated as in normal regression (see section 5.2.4) as $N - p - 1$ in which N is the total sample size (in this case 30), and p is the number of predictors (in this case 3, the two dummy variables and the covariate—see equation (9.1)). For these data we get $df = 30 - 3 - 1 = 26$.

Therefore, from these estimates we could conclude that the high dose differs significantly from the placebo group (DOSE=1 in the table) but that the high dose group also differs significantly from the low dose groups (DOSE = 2 in the table). This last conclusion is slightly odd because it contradicts what we initially concluded from the ANOVA (remember the means of the low- and high-dose groups were virtually identical)—can you think why? (All will be revealed in due course!). The final thing to notice is the value of *b* for the covariate (0.483). This value tells us that, other things being equal, if a partner's libido increases by one unit, then the person's libido should increase by just under half a unit (although there is nothing to suggest a causal link between the two). The sign of this coefficient tells us the direction of the relationship between the covariate and the outcome. So, in this example, because the coefficient is positive it means that the partner's libido has

Parameter Estimates

Dependent Variable: Libido

Parameter	B	Std. Error	t	Sig.	95% Confidence Interval	
					Lower Bound	Upper Bound
Intercept	3.843	.625	6.150	.000	2.558	5.127
PARTNER	.483	.196	2.472	.020	.081	.885
[DOSE = 1]	−2.607	.842	−3.095	.005	−4.338	−.876
[DOSE = 2]	−2.894	1.411	−2.051	.050	−5.794	.006
[DOSE = 3]	0[a]	.	.	.	.	.

a This parameter is set to zero because it is redundant

SPSS Output 9.3

a positive relationship with the participant's libido: as one increases so does the other. A negative coefficient would mean the opposite: as one increases, the other decreases.

9.4.2. Contrasts ②

SPSS Output 9.4 shows the result of the contrast analysis specified in Figure 9.2 and compares level 2 (low dose) against level 1 (placebo) as a first comparison, and level 3 (high dose)

Contrast Results (K Matrix)

Dose of Viagra Simple Contrast[a]			Dependent Variable
			Libido
Level 2 vs. Level 1	Contrast Estimate		−.287
	Hypothesized Value		0
	Difference (Estimate—Hypothesized)		−.287
	Std. Error		1.144
	Sig.		**.804**
	95% Confidence Interval for Difference	Lower Bound	−2.638
		Upper Bound	2.064
Level 3 vs. Level 1	Contrast Estimate		2.607
	Hypothesized Value		0
	Difference (Estimate—Hypothesized)		2.607
	Std. Error		.842
	Sig.		**.005**
	95% Confidence Interval for Difference	Lower Bound	.876
		Upper Bound	4.338

a Reference category = 1

SPSS Output 9.4

Estimates

Dependent Variable: Libido

Dose of Viagra	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
Placebo	3.313[a]	.572	2.138	4.489
Low Dose	3.027[a]	.962	1.049	5.004
High Dose	5.920[a]	.644	4.597	7.244

a Covariates appearing in the model are evaluated at the following values: Partner's Libido = 4.30

SPSS Output 9.5

against level 1 (placebo) as a second comparison. These contrasts are consistent with what was specified: all groups are compared to the first group. The group differences are displayed: a difference value, standard error, significance value and 95% confidence interval. These results show that the low-dose group did not have a significantly different libido than the placebo group (contrast 1, $p = .804$), but that the high-dose group did differ significantly from the placebo group ($p = .005$). These results are consistent with the regression parameter estimates (in fact, note that contrast 2 is identical to the regression parameters for DOSE=1 in the previous section).

Again, this all seems very odd because, at face value, the significant effect of libido seemed to reflect a difference between the placebo group and the two Viagra groups (which have similar means), yet the contrasts so far contradict these conclusions. The reason for this inconsistency is that the initial conclusion was based on group means that had not been adjusted for the effect of the covariate. These values tell us nothing about the group differences reflected by the significant ANCOVA. SPSS Output 9.5 gives the adjusted values of the group means and it is these values that should be used for interpretation (this is the main reason for selecting the *Display Means for* option). The adjusted means show a very different pattern of responses: it looks as though the significant ANCOVA reflects a difference between the high-dose group and both the low-dose group and the placebo group. The low-dose and placebo groups appear to have fairly similar adjusted means indicating that not enough Viagra does not increase libido above normal levels—a high dose is required! These conclusions support what we know from the contrasts and regression parameters but can be verified with the *post hoc* tests specified in the options menu.

SPSS Output 9.6 shows the results of the Sidak corrected *post hoc* comparisons that were requested as part of the *options* dialog box. The significant difference between the high-dose and placebo groups remains, but interestingly the significant difference between the high- and low-dose groups shown by the regression parameters (SPSS Output 9.3) is gone (p is only .14). This contradiction might result from a loss of power in the *post hoc* tests (remember that planned comparisons have greater power to detect effects than *post hoc* procedures). However, there could be other reasons why these comparisons are non-significant and we should be very cautious in our interpretation of the significant ANCOVA and subsequent comparisons.

Pairwise Comparisons

Dependent Variable: Libido

(I) Dose of Viagra	(J) Dose of Viagra	Mean Difference (I–J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
Placebo	Low Dose	.287	1.144	.992	-2.631	3.205
	High Dose	-2.607*	.842	.014	-4.756	-.458
Low Dose	Placebo	-.287	1.144	.992	-3.205	2.631
	High Dose	-2.894	1.411	.144	-6.493	.706
High Dose	Placebo	2.607*	.842	.014	.458	4.756
	Low Dose	2.894	1.411	.144	-.706	6.493

Based on estimated marginal means
*The mean difference is significant at the .05 level.
a Adjustment for multiple comparisons: Sidak

SPSS Output 9.6

9.4.3. Interpreting the covariate ②

I've already mentioned that the parameter estimates (SPSS Output 9.3) tell us how to interpret the covariate. If the *b*-value for the covariate is positive then it means that the covariate and the outcome variable have a positive relationship (as the covariate increases, so does the outcome). If the *b*-value is negative it means the opposite: that the covariate and the outcome variable have a negative relationship (as the covariate increases, the outcome decreases). For these data the *b*-value was positive, indicating that as the partner's libido increases, so does the participant's libido. Another way to discover the same thing is simply to draw a scatterplot of the covariate against the outcome. We came across scatterplots in section 4.4 so have a look back there to find out how to do one. Figure 9.4 shows the resulting scatterplot for these data and confirms what we already know: the effect of the covariate is that as the partner's libido increases, so does the participant's libido (as shown by the slope of the regression line).

9.5. ANCOVA RUN AS A MULTIPLE REGRESSION ②

Although the ANCOVA is essentially done, it may be of interest to rerun the analysis as a hierarchical multiple regression. As an exercise, enter these data and run the analysis yourself by adding two dummy variables to the file **ViagraCovariate.sav** that we've used in this chapter (see section 8.2.2 for help with the dummy coding). If you get stuck then I've included a completed data set called **ViagraCovariate Dummy.sav** on the CD-ROM. To run the analysis, we use the regression procedure (see Chapter 5) with **libido** as the

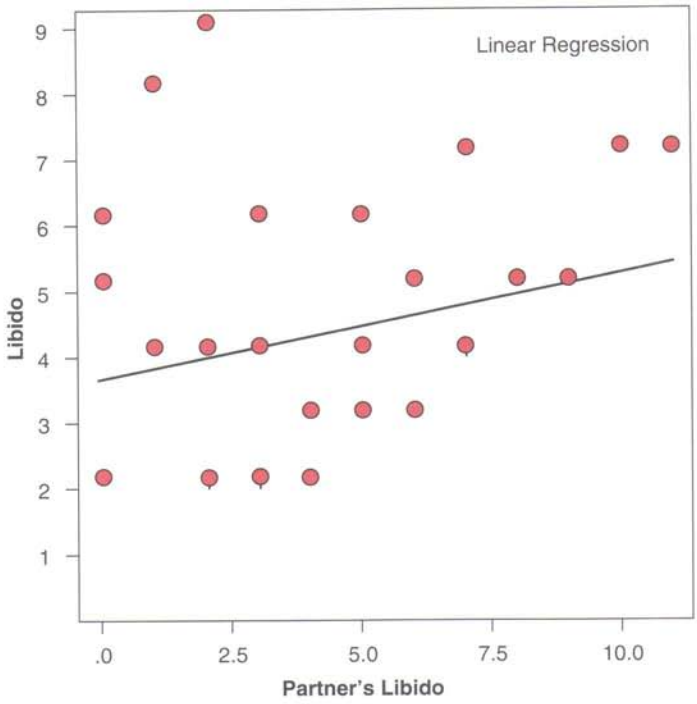

Figure 9.4 Scatterplot of libido against partner's libido

outcome (dependent variable in SPSS terminology) and then in the first block we enter partner's libido (**partner**) as a predictor, and then in a second block we enter both dummy variables. In all cases the entry method should be *Enter* (see section 5.7 for how to enter variables for a hierarchical regression but remember to use the *Enter* method for both blocks). The summary of the resulting regression model (SPSS Output 9.7) shows us the goodness-of-fit of the model first when only the covariate is used in the model, and second when both the covariate and the dummy variables are used.

Therefore, the difference between the values of R^2 (.313 – .058 = .255) represents the individual contribution of the dose of Viagra. Therefore, we can say that the dose of Viagra accounted for 25.5% of the variation in libido, whereas the partner's libido accounted for only 5.8%. This additional information provides some insight into the substantive importance of Viagra. The next table is the ANOVA table, which is again divided into two sections. The top half represents the effect of the covariate alone, whereas the bottom half represents the whole model (i.e. covariate and dose of Viagra included). Notice at the bottom of the ANOVA table (the bit for Model 2) that the entire model (partner's libido and the dummy variables) accounts for 34.57 units of variance (SS_M), there are 110.97 units in total (SS_T) and the unexplained variance (SS_R) is 76.22. The bottom half, therefore,

Model Summary

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.240[a]	.058	.024	1.932
2	.560[b]	.313	.234	1.712

a Predictors: (Constant), Partner's Libido
b Predictors: (Constant), Partner's Libido, Dummy Variables 1 (Placebo vs. High), Dummy Variable 1 c (Placebo vs. Low)

ANOVA[c]

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	6.413	1	6.413	1.717	.201[a]
	Residual	104.554	28	3.734		
	Total	110.967	29			
2	Regression	34.750	3	11.583	3.952	.019[b]
	Residual	76.216	26	2.931		
	Total	110.967	29			

a Predictors: (Constant), Partner's Libido
b Predictors: (Constant), Partner's Libido, Dummy Variables 1 (Placebo vs. High), Dummy Variable 1 (Placebo vs. Low)
c Dependent Variable: Libido

SPSS Output 9.7

contains the same values as the ANCOVA summary table in SPSS Output 9.2. These values are the same as the row labelled 'corrected model' row of the ANCOVA summary table we encountered when run the analysis as ANCOVA (see SPSS Output 9.2).

SPSS Output 9.8 shows the remainder of the regression analysis. This table of regression coefficients is more interesting. Again, this table is split into two and so the bottom of the table looks at the whole model. When the dose of Viagra is considered with the covariate, the value of *b* for the covariate is .483, which corresponds to the value in the ANCOVA parameter estimates (SPSS Output 9.3). The *b*-values for the dummy variables represent the difference between the means of the low-dose group and the placebo group (**dummy1**) and the high-dose group and the placebo group (**dummy2**)—see section 8.2.2 for an explanation of why. The means of the low- and high-dose groups were 4.88 and 4.85 respectively, and the mean of the placebo group was 3.22. Therefore, the *b*-values for the two dummy variables should be roughly the same (4.88 – 3.22 = 1.66 for dummy 1 and 4.85 – 3.22 = 1.63 for dummy 2). The astute among you might notice from the SPSS output that, in fact, the *b*-values are not only very different from each other (which shouldn't be the case because the high- and low-dose group means are virtually the same) but are different from the values I've just calculated. So, does this mean I've been lying to you for the past 50 pages about what the beta values represent? Well, even I'm not that horrible:

the reason for this apparent anomaly is because the b-values in this regression represent the differences between the adjusted means, not the original means; that is, the difference between the mean of each group and the placebo when these means have been adjusted for the partner's libido. The adjusted values were given in SPSS Output 9.5 and from this table we can see that:

$$b_{\text{Dummy 1}} = \overline{X}_{\text{Low(adjusted)}} - \overline{X}_{\text{Placebo(adjusted)}} = 3.027 - 3.313 = -0.286$$
$$b_{\text{Dummy 2}} = \overline{X}_{\text{High(adjusted)}} - \overline{X}_{\text{Placebo(adjusted)}} = 5.920 - 3.313 = -2.607 \quad (9.12)$$

These are the values you can see in the SPSS table. The t-tests conducted on these values show that the significant ANCOVA reflected a significant difference between the high-dose and placebo groups.[2] There was no significant difference between the low-dose and placebo groups. You should also notice that the significances of the t-values are the same as we saw in the contrasts table of the original ANCOVA (see SPSS Output 9.4). As a final point, we obviously don't know whether there was a difference between the low-dose and high-dose groups: to find this out we would need to use different dummy coding (perhaps comparing the high and low to the placebo and then comparing high to low just like we used for the planned comparisons in Chapter 8—see Box 9.2).

Coefficients[a]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	3.689	.626		5.894	.000
	Partner's Libido	.158	.120	.240	1.311	.201
2	(Constant)	1.236	.986		1.253	.221
	Partner's Libido	.483	.196	.737	2.472	.020
	Dummy Variable 1 (Placebo vs. Low)	−.287	1.144	−.066	−.251	.804
	Dummy Variable 2 (Placebo vs. High)	2.607	.842	.672	3.095	.005

a Dependent Variable: Libido

SPSS Output 9.8

2 As I mentioned earlier in this chapter the degrees of freedom for these t-tests are $N - p - 1$, as in any regression analysis (see section 5.2.4). N is the total sample size (in this case 30), and p is the number of predictors (in this case 3, the two dummy variables and the covariate—see equation (9.1)). For these data we get $df = 30 - 3 - 1 = 26$.

Box 9.2

Planned contrasts for ANCOVA ③

You may have noticed that although we can ask SPSS to do certain standard contrasts, there is no option for specifying planned contrasts as there was with one-way independent ANOVA (see section 8.3.1). However, these contrasts can be done if we run the ANCOVA through the regression menu. Imagine you chose some planned contrasts as in Chapter 8, in which the first contrast compared the placebo group with all doses of Viagra, and the second contrast then compared the high and low doses (see section 8.2.10). We saw in sections 8.2.10 and 8.3.1 that to do this in SPSS we had to enter certain numbers to code these contrasts. For the first contrast we discovered an appropriate set of codes would be –2 for the placebo group and then 1 for both the high- and low-dose groups. For the second contrast the codes would be 0 for the placebo group, –1 for the low-dose group and 1 for the high-dose group (see Table 8.4). If you want to do these contrasts for ANCOVA, then you enter these values as two dummy variables. So, taking the data in this example, we'd add a column called **Dummy1** and in that column we'd put the value –2 for every person who was in the placebo group, and the value 1 for all other participants. We'd then add a second column called **Dummy2**, in which we'd place a 0 for everyone in the placebo group, a –1 for everyone in the low-dose group and the value 1 for those in the high dose group. The completed data would be like the file **ViagraCovariateContrasts.sav** on the CD-ROM.

Run the analysis as described in section 9.5. The resulting output will begin with a model summary and ANOVA table that should be identical to those in SPSS Output 9.7 (because we've done the same thing as before, the only difference is how the model variance is subsequently broken down with the contrasts). The regression coefficients for the dummy variables will be different, though, because we've now specified different codes.

Coefficients[a]

Model		Unstandardized Coefficients		Standardized Coefficients	t	Sig.
		B	Std. Error	Beta		
1	(Constant)	3.689	.626		5.894	.000
	Partner's Libido	.158	.120	.240	1.311	.201
2	(Constant)	2.009	.986		2.038	.052
	Partner's Libido	.483	.196	.737	2.472	.020
	Dummy Variable 1 (Placebo vs. Low & High)	.387	.238	.276	1.623	.117
	Dummy Variable 2 (Low vs. High)	1.447	.705	.617	2.051	.050

a Dependent Variable: Libido

(Continued)

Box 9.2 (Continued)

The first dummy variable compares the placebo group with the low- and high-dose groups. As such, it compares the adjusted mean of the placebo group (3.313) with the average of the adjusted means for the low- and high-dose groups ((3.027 + 5.920)/2 = 4.474). The *b*-value for the first dummy variable should therefore be the difference between these values: 4.474 – 3.313 = 1.16. However, we also discovered in a rather complex and boring bit of section 8.2.10.2 that this value gets divided by the number of groups within the contrast (i.e. 3) and so will be 1.16/3 = .387 (as it is in the output). The associated *t*-statistic is not significant, indicating that the placebo group was not significantly different from the combined mean of the Viagra groups.

The second dummy variable compares the low- and high-dose groups, and so the *b*-value should be the difference between the adjusted means of these groups: 5.920 – 3.027 = 2.89. We again discovered in section 8.2.10.2 that this value also gets divided by the number of groups within the contrast (i.e. 2) and so will be 2.89/2 = 1.447 (as in the output). The associated *t*-statistic is significant (its significance is exactly .05), indicating that the high-dose group produced a significantly higher libido than the low-dose group after controlling for the effect of the partner's libido.

This illustrates how you can apply the principles from section 8.2.10 to ANCOVA: although SPSS doesn't provide an easy interface to do planned contrasts, they can be done if you use the regression menus rather than the ANCOVA ones!

9.6. ADDITIONAL ASSUMPTIONS IN ANCOVA ③

9.6.1. Homogeneity of regression slopes ③

When an ANCOVA is conducted we look at the overall relationship between the outcome (dependent variable) and the covariate: we fit a regression line to the entire data set, ignoring to which group a person belongs. In fitting this model we therefore assume that this overall relationship is true for all groups of participants. For example, if there's a positive relationship between the covariate and the outcome in one group, we assume that there is a positive relationship in all of the other groups too. If, however, the relationship between the outcome (dependent variable) and covariate differs across the groups then the overall regression model is inaccurate (it does not represent all of the groups). This assumption is very important and is called the assumption of homogeneity of regression slopes. The best way to think of this assumption is to imagine plotting a scatterplot for each experimental condition with the covariate on one axis and the outcome on the other. If you then calculated, and drew, the regression line for each of these scatterplots you should find that the regression lines look more or less the same (i.e. the values of *b* in each group should be equal).

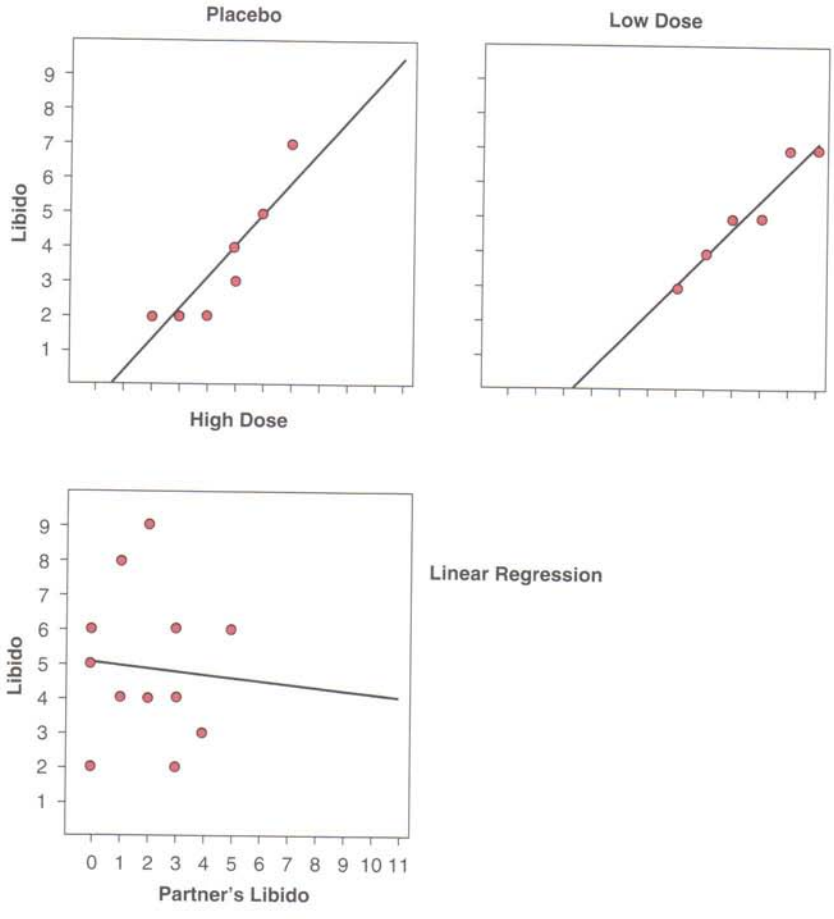

Figure 9.5 Scatterplots and regression lines of libido against partner's libido for each of the experimental conditions

Figure 9.5 shows scatterplots that display the relationship between partner's libido (the covariate) and the outcome (participant's libido) for each of the three experimental conditions. In each scatterplot a dot represents the data from a particular participant and the lines are the regression slopes for the particular group (i.e. they summarize the relationship between libido and partner's libido shown by the dots). It should be clear that there is a positive relationship (the regression line slopes upwards from left to right) between partner's libido and participant's libido in both the placebo and low-dose conditions. However, in the high-dose condition there appears to be no relationship at all between participant's libido and that of their partner (the dots are fairly randomly scattered and, in fact, the regression line slopes downwards from left to right, indicating a slightly negative

relationship). This observation gives us cause to doubt whether there is homogeneity of regression slopes (because the relationship between participant's libido and that of their partner is not consistent across the three experimental groups).

9.6.2. Testing for homogeneity of regression slopes in SPSS ③

To test the assumption of homogeneity of regression slopes we need to rerun the ANCOVA but this time use a customized model. Access the main dialog box as before and place the variables in the same boxes as before (the finished box should look like Figure 9.1). To customize the model we need to access the *model* dialog box (Figure 9.6) by clicking on Model....

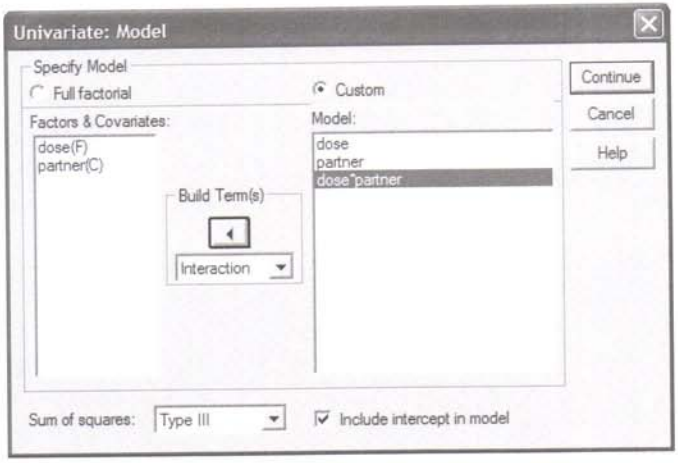

Figure 9.6 GLM univariate *model* dialog box

To customize your model, click on the circle labelled *Custom* to activate the dialog box (Figure 9.6). The variables specified in the main dialog box are listed on the left-hand side and are followed by a letter indicating the variable type (F = fixed factor, C = covariate). To test the assumption of homogeneity of regression slopes, we need to specify a model that includes the interaction between the covariate and dependent variable. Ordinarily, ANCOVA includes only the main effect of dose and partner and does not include this interaction term. To test this interaction term it's important still to include the main effects of dose and partner so that the interaction term is tested controlling for these main effects.

Hence, to begin with you should select **dose** and **partner** (you can select both of them at the same time). Then, where it says *Build Term(s)* there is a drop-down menu. Click on ▾ to access this drop-down menu and then click on *Main effects*. Having selected this, click on ▸ to move the main effects of **dose** and **partner** to the box labelled *Model*. Next we need to specify the interaction term. To do this, select **dose** and **partner** simultaneously and then click on ▾ to access the drop-down menu again, but this time select *Interaction*. Having selected this, click on ▸ to move the interaction of **dose** and **partner** to the box labelled *Model*. The

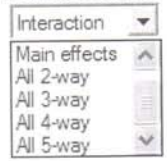

Tests of Between-Subjects Effects

Dependent Variable: Libido

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	52.000[a]	5	10.400	4.233	.007
Intercept	1.706	1	1.706	.694	.413
DOSE	40.068	2	20.034	8.154	.002
PARTNER	22.430	1	22.430	9.129	.006
DOSE * PARTNER	17.250	2	8.625	3.510	.046
Error	58.967	24	2.457		
Total	683.000	30			
Corrected Total	110.967	29			

a R Squared = .469 (Adjusted R Squared = .358)

SPSS Output 9.9

finished dialog box should look like Figure 9.6. You can find out more about specifying effects in section 10.3.2. Having specified the two main effects and the interaction term, click on Continue to return to the main dialog box and then click on OK to run the analysis.

SPSS Output 9.9 shows the main summary table for the ANCOVA including the interaction term. The effects of the dose of Viagra and the partner's libido are still significant, but the main thing in which we're interested is the interaction term, so look at the significance value of the covariate by outcome interaction (**dose*partner**): if this effect is significant then the assumption of homogeneity of regression slopes has been broken. The effect here is significant ($p < .05$); therefore the assumption is not tenable. Although this finding is not surprising given the pattern of relationships shown in Figure 9.5 it does raise concern about the main analysis, especially in the light of the contradictory findings of the multiple comparisons. This example illustrates why it is important to test assumptions and not just blindly to accept the results of an analysis.

Cramming Samantha's Tips

- Analysis of covariance (ANCOVA) compares several means, but controlling for the effect of one or more other variables (called covariates); for example, if you have several experimental conditions and want to control for the age of the participants.
- In the table labelled *Tests of between-subjects effects*, look at the column labelled *Sig.* for both the covariate and the independent variable (the grouping variable); if the value is less than .05 then for the covariate it means that this variable has a significant relationship to the outcome variable (the dependent variable); for the grouping variable it means that the means of the groups are significantly different after controlling for the effect that the covariate has on the outcome.

(Continued)

Cramming Samantha's Tips (Continued)

- As with ANOVA, if you have generated specific hypotheses before the experiment use *planned comparisons*, but if you don't have specific hypotheses use *post hoc* tests. Although SPSS will let you specify certain standard contrasts, other planned comparisons will have to be done by analysing the data using the regression procedure in SPSS.
- For contrasts and *post hoc* tests, again look to the columns labelled *Sig.* to discover if your comparisons are significant (they will be if the significance value is less than .05).
- Test the same assumptions as for ANOVA, but in addition you should test the assumption of *homogeneity of regression slopes*. This has to be done by customizing the ANCOVA model in SPSS.

9.7. CALCULATING THE EFFECT SIZE ②

We saw in Box 9.1 that we can get SPSS to produce eta squared (η^2), which is just r^2 calculated from the between-group effect, SS_M, divided by the total amount of variance in the data, SS_T. However, this measure of effect size is slightly biased and I recommended not using it. Therefore, as with ANOVA you're well advised to use omega squared (ω^2), which is a less biased version of eta squared (see section 8.5). However, as we saw in section 8.5 this can only be measured when you have equal numbers of participants in each group (and we don't have that here!). So, we're a bit stumped!

However, not all is lost because as I've said many times already, the overall effect size is not nearly as interesting as the effect size for more focused comparisons. These are easy to calculate because we selected regression parameters (see SPSS Output 9.3) and so we have *t*-statistics for the covariate and comparisons between the low- and high-dose groups and the placebo and high-dose groups. These *t*-statistics have $N - 2$ degrees of freedom (see Chapter 5), where N is the total sample size (in this case 30). We can use the same equation as in section 7.5.5:[3]

$$r_{\text{contrast}} = \sqrt{\frac{t^2}{t^2 + df}}$$

3 Strictly speaking, we have to use a slightly more elaborate procedure when groups are unequal. It's a bit beyond the scope of theis book but Rosnow, Rosenthal & Rubin (2000) give a very clear account.

Therefore we get (*t*s from SPSS output 9.3):

$$r_{\text{Covariate}} = \sqrt{\frac{2.47^2}{2.47^2 + 28}} = \sqrt{\frac{6.10}{34.10}} = .42$$

$$r_{\text{High Dose vs. Placebo}} = \sqrt{\frac{-3.095^2}{-3.095^2 + 28}} = \sqrt{\frac{9.58}{37.57}} = .50$$

$$r_{\text{High vs. Low Dose}} = \sqrt{\frac{-2.051^2}{-2.051^2 + 28}} = \sqrt{\frac{4.21}{32.21}} = .36$$

If you think back to our benchmarks for effect sizes these all represent medium to large effect sizes (they're all between .3 and .5). Therefore, as well as being statistically significant, these effects are substantive findings.

9.8. REPORTING RESULTS ②

Reporting ANCOVA is much the same as reporting an ANOVA except we now have to report the effect of the covariate as well. For the covariate and the experimental effect we give details of the *F*-ratio and the degrees of freedom from which it was calculated. In both cases, the *F*-ratio was derived from dividing the mean squares for the effect by the mean squares for the residual. Therefore, the degrees of freedom used to assess the *F*-ratio are the degrees of freedom for the effect of the model ($df_M = 1$ for the covariate and 2 for the experimental effect) and the degrees of freedom for the residuals of the model ($df_R = 26$ for both the covariate and the experimental effect)—see SPSS Output 9.2. Therefore, the correct way to report the main findings would be:

- The covariate, partner's libido, was significantly related to the participant's libido, $F(1, 26) = 6.11$, $p < .05$, $r = .42$. There was also significant effect of Viagra on levels of libido after controlling for the effect of partner's libido, $F(2, 26) = 4.83$, $p < .05$.

We can also report some contrasts:

- Planned contrasts revealed that having a high dose of Viagra significantly increased libido compared to having both a placebo, $t(26) = -3.10$, $p < .05$, $r = .50$, and a low dose, $t(26) = -2.05$, $p < .05$, $r = .36$.

9.9. WHAT HAVE WE DISCOVERED ABOUT STATISTICS? ②

This chapter has shown you how the general linear model (GLM) that was described in Chapter 8 can be extended to include additional variables. The advantages of doing so are that we can control for factors other than our experimental manipulation that might influence our outcome measure. This gives us tighter experimental control, and may also help us to explain some of our error variance, and, therefore, give us a purer measure of the experimental manipulation. We didn't go into too much theory about ANCOVA, just looked conceptually at how the regression model can be expanded to include these additional variables (*covariates*). Instead we jumped straight into an example, which was to look at the effect of Viagra on libido (as in Chapter 8) but including partner's libido as a covariate. I explained how to do the analysis on SPSS and interpret the results but also showed how the same output could be obtained by running the analysis as a regression. This was to try to get the message home that ANOVA and ANCOVA are merely forms of regression! Anyway, we finished off by looking at an additional assumption that has to be considered when doing ANCOVA: the assumption of homogeneity of regression slopes. This just means that the relationship between the covariate and the outcome variable should be the same in all of your experimental groups. We also had a look at how to test this assumption on SPSS. We'll now move on to look at situations in which you've got more than one experimental manipulation.

9.10. KEY TERMS THAT WE'VE DISCOVERED

- Adjusted means
- Analysis of covariance (ANCOVA)
- Homogeneity of regression slopes
- Partial out
- Sidak correction
- Covariates

9.11. SMART ALEX'S TASKS

- **Task 1**: Stalking is a very disruptive and upsetting (for the person being stalked) experience in which someone (the stalker) constantly harasses or obsesses about another person. It can take many forms, from sending intensely disturbing letters

threatening to boil your cat if you don't reciprocate the stalker's undeniable love for you, to literally following you around your local area in a desperate attempt to see which CD you buy on a Saturday (as if it would be anything other than Fugazi!). A psychologist, who'd had enough of being stalked by people, decided to try two different therapies on different groups of stalkers (25 stalkers in each group—this variable is called **Group**). The first group of stalkers he gave what he termed cruel-to-be-kind therapy. This therapy was based on punishment for stalking behaviours; in short every time the stalker followed him around, or sent him a letter, the psychologist attacked them with a cattle prod until they stopped their stalking behaviour. It was hoped that the stalkers would learn an aversive reaction to anything resembling stalking. The second therapy was psychodyshamic therapy, which was a recent development on Freud's psychodynamic therapy that acknowledges what a sham this kind of treatment is (so, you could say it's based on Fraudian theory!). The stalkers were hypnotized and regressed into their childhood, the therapist would also discuss their penis (unless it was a woman in which case they discussed their lack of penis), the penis of their father, their dog's penis, the penis of the cat down the road, and anyone else's penis that sprang to mind. At the end of therapy, the psychologist measured the number of hours in the week that the stalker spent stalking their prey (this variable is called **stalk2**). Now, the psychologist believed that the success of therapy might well depend on how bad the problem was to begin with, so before therapy he measured the number of hours that the patient spent stalking as an indicator of how much of a stalker the person was (this variable is called **stalk1**). The data are in the file **Stalker.sav.** Analyse the effect of therapy on stalking behaviour after therapy, controlling for the amount of stalking behaviour before therapy.②

- **Task 2**: A marketing manager for a certain well-known drinks manufacturer was interested in the therapeutic benefit of certain soft drinks for curing hangovers. He took 15 people out on the town one night and got them drunk. The next morning as they awoke, dehydrated and feeling as though they'd licked a camel's sandy feet clean with their tongue, he gave five of them water to drink, five of them Lucozade (in case this isn't sold outside of the UK, it's a very nice glucose-based drink), and the remaining five a leading brand of cola (this variable is called **drink**). He then measured how well they felt (on a scale from 0 = I feel like death to 10 = I feel really full of beans and healthy) 2 hours later (this variable is called **well**). He wanted to know which drink produced the greatest level of wellness. However, he realized it was important to control for how drunk the person got the night before, and so he measured this on a scale of 0 = as sober as a nun to 10 = flapping about like a haddock out of water on the floor in a puddle of their own vomit. The data are in the file **HangoverCure.sav**. Conduct an ANCOVA to see whether people felt better after different drinks when controlling for how drunk they were the night before.②

The answers are in the file **Answers(Chapter9).pdf** and task 1 has a full interpretation in Field & Hole (2003).

9.12. FURTHER READING

Howell, D. C. (2002). *Statistical methods for psychology* (5th edition). Belmont, CA: Duxbury. Chapter 16.

Rutherford, A. (2000). *Introducing ANOVA and ANCOVA: A GLM Approach*. London: Sage.

Wildt, A. R. & Ahtola, O. (1978). *Analysis of covariance*. Sage university paper series on quantitative applications in the social sciences, 07–012. Newbury Park, CA: Sage. This text is pretty high level but very comprehensive if you want to know the maths behind ANCOVA.