

Chapter 13: Factorial ANOVA

Oliver Twisted

Please, Sir, can I ... customize my model?



Different types of sums of squares

In the sister book on R, I needed to explain sums of squares because, unlike SPSS, R does not routinely use Type III sums of squares. Actually, people who use R tend to turn a funny shade of pink at the mention of Type III sums of squares and start mumbling things about the developers of SPSS and SAS (another software package) being bastards. Anyway, here's what I wrote (Field, Miles, & Field, 2012):

“We can compute sums of squares in four different ways, which gives rise to what are known as Type I, II, III and IV sums of squares. To explain these, we need an example. Let's imagine that we're predicting **libido** from **partnerLibido** (the covariate), **dose** (the independent variable) and their interaction (**partnerLibido** × **dose**).

The simplest explanation of Type I sums of squares is that they are like doing a hierarchical regression in which we put one predictor into the model first, and then enter the second predictor. This second predictor will be evaluated after the first. If we entered a third predictor then this would be evaluated after the first and second, and so on. In other words the order that we enter the predictors matters. Therefore, if we entered our variables in the order **partnerLibido**, **dose** and then **partnerLibido** × **dose**, then **dose** would be evaluated after the effect of **partnerLibido** and **partnerLibido** × **dose** would be evaluated after the effects of *both* **partnerLibido** and **dose**.

Type III sums of squares differ from Type I in that all effects are evaluated taking into consideration *all other effects in the model* (not just the ones entered before). This process is comparable to doing a forced entry regression including the covariate(s) and predictor(s) in the same block. Therefore, in our example, the effect of **dose** would be evaluated after the effects of both **partnerLibido** and **partnerLibido** × **dose**, the effect of **partnerLibido** would be evaluated after the effects of both **dose** and **partnerLibido** × **dose**, finally, **partnerLibido** × **dose** would be evaluated after the effects of both **dose** and **partnerLibido**.

Type II sums of squares are somewhere in between Type I and III in that all effects are evaluated taking into consideration all other effects in the model *except for higher-order effects that include the effect being evaluated*. In our example, this would mean that the effect of **dose** would be evaluated after the effect of **partnerLibido** (note that unlike Type III sums of squares, the interaction term is not considered); similarly, the effect of **partnerLibido** would be

evaluated after only the effect of **dose**. Finally, because there is no higher-order interaction that includes **partnerLibido** \times **dose**, this effect would be evaluated after the effects of both **dose** and **partnerLibido**. In other words, for the highest-order term Type II and Type III sums of squares are the same. Type IV sums of squares are essentially the same as Type III but are designed for situations in which there are missing data.



The obvious question is which type of sums of squares should you use:

- **Type I:** Unless the variables are completely independent of each other (which is unlikely to be the case) then Type I sums of squares cannot really evaluate the true main effect of each variable. For example, if we enter **partnerLibido** first, its sums of squares are computed ignoring **dose**; therefore any variance in **libido** that is shared by **dose** and **partnerLibido** will be attributed to **partnerLibido** (i.e., variance that it shares with **dose** is attributed solely to it). The sums of squares for **dose** will then be computed excluding any variance that has already been 'given over' to **partnerLibido**. As such the sums of squares won't reflect the true effect of **dose** because variance in libido that **dose** shares with **partnerLibido** is not attributed to it because it has already been 'assigned' to **partnerLibido**. Consequently, Type I sums of squares tend not to be used to evaluate hypotheses about main effects and interactions because the order of predictors will affect the results.
- **Type II:** If you're interested in main effects then you should use Type II sums of squares. Unlike Type III sums of squares, Type IIs give you an accurate picture of a main effect because they are evaluated ignoring the effect of any interactions involving the main effect under consideration. Therefore, variance from a main effect is not 'lost' to any interaction terms containing that effect. If you are interested in main effects and do not predict an interaction between your main effects then these tests will be the most powerful. However, *if an interaction is present*, then Type II sums of squares cannot reasonably evaluate main effects (because variance from the interaction term is attributed to them). However, if there is an interaction then you shouldn't really be interested in main effects anyway. One advantage of Type II sums of squares is that they are not affected by the type of contrast coding used to specify the predictor variables.
- **Type III:** Type III sums of squares tend to get used as the default in many statistical packages. They have the advantage over Type IIs in that when an interaction is present, the main effects associated with that interaction are still meaningful (because they are computed taking the interaction into account). Perversely, this advantage is a disadvantage too because it's pretty silly to entertain 'main effects' as meaningful in the presence of an interaction. Type III sums of squares encourage people to do daft things like get excited about main effects that are superseded by a higher-order interaction. Type III sums of squares are preferable to other types when sample sizes are unequal; however, they work only when predictors are encoded with orthogonal contrasts.

Hopefully, it should be clear that the main choice in ANOVA designs is between Type II and Type III sums of squares. The choice depends on your hypotheses and which effects are important in your particular situation. If your main hypothesis is around the highest order interaction then it doesn't matter which you choose (you'll get the same results); if you don't predict an interaction and are interested in main effects then Type II will be most powerful; and if you have an unbalanced design then use Type III. This advice is, of course, a simplified version of reality; be aware that there is (often heated) debate about which sums of squares are appropriate to a given situation."


Customizing an ANOVA model

By default SPSS conducts a full factorial analysis (i.e., it includes all of the main effects and interactions of all independent variables specified in the main dialog box). However, there may be times when you want to customize the model that you use to test for certain things. To access the model dialog box, click on **Model...** in the main dialog box. You will notice that, by default, the full factorial model is selected. Even with this selected, there is an option at the bottom to change the types of sums of squares that are used in the analysis. Although we have learnt about sums of squares and what they represent, I haven't talked about different ways of calculating sums of squares. It isn't necessary to understand the computation of the different forms of sums of squares, but it is important that you know the uses of some of the different types. By default, SPSS uses Type III sums of squares, which have the advantage that they are invariant to the cell frequencies. As such, they can be used with both balanced and unbalanced (i.e., different numbers of participants in different groups) designs, which is why they are the default option. Type IV sums of squares are like Type III except that they can be used with data in which there are missing values. So, if you have any missing data in your design, you should change the sums of squares to Type IV.



To customize a model, click on **Custom** to activate the dialog box. The variables specified in the main dialog box will be listed on the left-hand side. You can select one, or several, variables from this list and transfer them to the box labelled **Model** as either main effects or interactions. By default, SPSS transfers variables as interaction terms, but there are several options that allow you to enter main effects, or all two-way, three-way or four-way interactions. These options save you the trouble of having to select lots of combinations of variables (because, for example, you can select three variables, transfer them as all two-way interactions and it will create all three combinations of variables for you). Hence, you could select **Gender** and **Alcohol** (you can select both of them at the same time by holding down *Ctrl*). Then, click on the drop-down menu and change it to **Main effects**. Having selected this, click on **+** to move the main effects of **Gender** and **Alcohol** to the box labelled **Model**. Next you could specify the interaction term. To do this, select **Gender** and **Alcohol** simultaneously (by holding down the *Ctrl* key while you click on the two variables), then select **Interaction** in the drop-down list and click on **+**. This action moves the interaction of **Gender** and **Alcohol** to the box labelled **Model**. The finished dialog box should look like that below. Having specified our

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. Although model selection has important uses, it is likely that you'd want to run the full factorial analysis on most occasions and so wouldn't customize your model.



Please, Sir, can I have some more ... contrasts?



Why do we need to use syntax?

In Chapters 12, 13 and 14 of the book we use SPSS's built-in *contrast* functions to compare various groups after conducting ANOVA. These special contrasts (described in Chapter 10, Table 10.6) cover many situations, but in more complex designs there will be times when you want to do contrasts that simply can't be done using SPSS's built-in contrasts. Unlike one-way ANOVA, there is no way in factorial designs to define contrast codes through the Windows dialog boxes. However, SPSS can do these contrasts if you define them using syntax.

An example

Imagine a clinical psychologist wanted to see the effects of a new antidepressant drug called Cheerup. He took 50 people suffering from clinical depression and randomly assigned them to one of five groups. The first group was a waiting list control group (i.e., people assigned to the waiting list who were not treated during the study), the second took a placebo tablet (i.e., they were told they were being given an antidepressant drug but actually the pills contained sugar and no active agents), the third group took a well-established SSRI antidepressant called Seroxat (Paxil to American readers), the fourth group was given a well-established SNRI


antidepressant called Effexor,¹ and the final group was given the new drug, Cheerup. Levels of depression were measured before and after two months on the various treatments, and ranged from 0 = as happy as a spring lamb to 20 = pass me the noose. The data are in the file **Depression.sav**.

This study is a two-way mixed design. There are two independent variables: treatment (no treatment, placebo, Seroxat, Effexor or Cheerup) and time (before or after treatment). Treatment is measured with different participants (and so is between-group) and time is, obviously, measured using the same participants (and so is repeated-measures). Hence, the ANOVA we want to use is a 5×2 two-way ANOVA.

Now, we want to do some contrasts. Imagine we have the following hypotheses:


1. Any treatment will be better than no treatment.
2. Drug treatments will be better than the placebo.
3. Our new drug, Cheerup, will be better than old-style antidepressants.
4. The old-style antidepressants will not differ in their effectiveness.

We have to code these various hypotheses as we did in Chapter 11. The first contrast involves comparing the no-treatment condition to all other groups. Therefore, the first step is to chunk these variables, and then assign a positive weight to one chunk and a negative weight to the other chunk.




Having done that, we need to assign a numeric value to the groups in each chunk. As I mentioned in Chapter 8, the easiest way to do this is just to assign a value equal to the number of groups in the *opposite* chunk. Therefore, the value for any group in chunk 1 will be the same as the number of groups in chunk 2 (in this case 4). Likewise, the value for any groups in chunk 2 will be the same as the number of groups in chunk 1 (in this case 1). So, we get the following codes:


¹ SSRIs, selective serotonin reuptake inhibitors, work selectively to inhibit the reuptake of the neurotransmitter serotonin in the brain, whereas SNRIs, serotonin norepinephrine reuptake inhibitors, which are newer, act not only on serotonin but also on another neurotransmitter (from the same family), norepinephrine.




The second contrast requires us to compare the placebo group to all of the drug groups. Again, we chunk our groups accordingly, assign one chunk a negative sign and the other a positive, and then assign a weight on the basis of the number of groups in the opposite chunk. We must also remember to give the no-treatment group a weight of 0 because they're not involved in the contrast.



The third contrast requires us to compare the new drug (Cheerup) to the old drugs (Seroxat and Effexor). Again, we chunk our groups accordingly, assign one chunk a negative sign and the other a positive, and then assign a weight on the basis of the number of groups in the opposite chunk. We must also remember to give the no-treatment and placebo groups a weight of 0 because they're not involved in the contrast.



The final contrast requires us to compare the two old drugs. Again, we chunk our groups accordingly, assign one chunk a negative sign and the other a positive, and then assign a weight on the basis of the number of groups in the opposite chunk. We must also give the no-treatment, placebo and Cheerup groups a weight of 0 because they're not involved in the contrast.



We can summarize these codes in the following table:

	No Treatment	Placebo	Seroxat	Effexor	Cheerup
Contrast 1	-4	1	1	1	1
Contrast 2	0	-3	1	1	1
Contrast 3	0	0	1	1	-2
Contrast 4	0	0	1	-1	0

These are the codes that we need to enter into SPSS to do the contrasts that we'd like to do.

Entering the contrasts using syntax

To enter these contrasts using syntax we have to first open a syntax window (see Chapter 2 of the book). Having done that we have to type the following commands:

MANOVA

```
before after BY treat(0 4)
```

This initializes the ANOVA command in SPSS. The second line specifies the variables in the data editor. The first two words, 'before' and 'after', are the repeated-measures variables (and these words are the words used in the data editor). Anything after **BY** is a between-group measure and so needs to be followed by brackets within which the minimum and maximum values of the coding variable are specified. I called the between-group variable **treat**, and I coded the groups as 0 = no treatment, 1 = placebo, 2 = Seroxat, 3 = Effexor, 4 = Cheerup. Therefore, the minimum and maximum codes were 0 and 4. So these two lines tell SPSS to start the ANOVA procedure, that there are two repeated-measures variables called before and after, and that there is a between-group variable called treat that has a minimum code of 0 and a maximum of 4.

```
/WSFACTORS time (2)
```

The **/WSFACTORS** command allows us to specify any repeated-measures variables. SPSS already knows that there are two variables called before and after, but it doesn't know how to treat these variables. This command tells SPSS to create a repeated-measures variable called time that has two levels (the number in brackets). SPSS then looks to the variables specified before and assigns the first one (in this case before) to be the first level of time, and then assigns the second one (in this case after) to be the second level of time.

```
/CONTRAST (time)=special(1 1, 1 -1)
```

This is used to specify the contrasts for the first variable. The **/CONTRAST** is used to specify any contrast. It's always followed by the name of the variable that you want to do a contrast on in brackets. We have two variables (time and treat) and in this first contrast we want to specify a contrast for time. Time only has two levels, and so all we want to do is to tell SPSS to compare these two levels (which actually it will do by default, but I want you to get some practice in!). What we write after the equals sign defines the contrast, so we could write the name of one of the standard contrasts such as Helmert, but because we want to specify our own contrast we use the word special. Special should always be followed by brackets, and inside those brackets are your contrast codes. Codes for different contrasts are separated using a comma, and within a contrast, codes for different groups are separated using a space. The first contrast should always be one that defines a baseline for all other contrasts and that is one that codes all groups with a 1. Therefore, because we have two levels of time, we just write 1 1, which tells SPSS that the first contrast should be one in which both before and after are given a code of 1. The comma tells SPSS that a new contrast follows and this second contrast has been defined as 1 -1, which tells SPSS that in this second contrast we want to give

before a code of 1, and after a code of -1. Note that the codes you write in the brackets are assigned to variables in the order that those variables are entered into the SPSS syntax, so because we originally wrote before after BY treat(0 4), SPSS assigns the 1 to before and -1 to after; if we'd originally written after before BY treat(0 4) then SPSS would have assigned them the opposite way round: the 1 to after and -1 to before.

```
/CONTRAST (treat)=special(1 1 1 1 1, -4 1 1 1 1, 0 -3 1 1 1, 0 0 1 1 -2, 0 0 1 -1 0)
```

This is used to specify the contrasts for the second variable. This time the /CONTRAST command is followed by the name of the second variable (treat). Treat has five levels, and we've already worked out four different contrasts that we want to do. Again we use the word special after the equals sign and specify our coding values within the brackets. As before, codes for different contrasts are separated using a comma and, within a contrast, codes for different groups are separated using a space. Also, as before, the first contrast should always be one that defines a baseline for all other contrasts, and that is one that codes all groups with a 1. Therefore, because we have five levels of treat, we just write 1 1 1 1 1, which tells SPSS that the first contrast should be one in which all five groups are given a code of 1. The comma tells SPSS that a new contrast follows and this second contrast has been defined as -4 1 1 1 1, which tells SPSS that in this second contrast we want to give the first group a code of -4 and all subsequent groups codes of 1. How does SPSS decide what the first group is? It uses the coding variable in the data editor and orders the groups in the same order as the coding variable. Therefore, because I coded the groups as 0 = no treatment, 1 = placebo, 2 = Seroxat, 3 = Effexor, 4 = Cheerup, this first contrast gives the no-treatment group a code of -4, and all subsequent groups codes of 1. The comma again tells SPSS that, having done this, there is another contrast to follow and this contrast has been defined as 0 -3 1 1 1, which tells SPSS that in this contrast we want to give the first group (no treatment) a code of 0, the second group (placebo) a code of -3 and all subsequent groups codes of 1. The comma again tells SPSS that, having done this, there is another contrast to follow and this contrast has been defined as 0 0 1 1 -2, which tells SPSS that in this contrast we want to give the first two groups (no treatment and placebo) a code of 0, the third and fourth groups (Seroxat and Effexor) a code of 1 and the final group (Cheerup) a code of -2. The comma again tells SPSS that there is yet another contrast to follow and this contrast has been defined as 0 0 1 -1 0, which tells SPSS that in this contrast we want to give the first, second and last (no treatment, placebo and Cheerup) groups a code of 0, the third group (Seroxat) a code of 1 and the fourth group (Effexor) a code of -1. As such, this one line of text has defined the four contrasts that we want to do.

```
/CINTERVAL JOINT(.95) MULTIVARIATE(BONFER)
```

This line defines the type of confidence intervals that you want to do for your contrasts. I recommend the Bonferroni option, but if you delve into the SPSS syntax guide you can find others.

```
/METHOD UNIQUE
```

```

/ERROR WITHIN+RESIDUAL

/PRINT TRANSFORM HOMOGENEITY(BARTLETT COCHRAN BOXM)

SIGNIF( UNIV MULT AVERF HF GG )

PARAM( ESTIM EFSIZE).

```

These lines of syntax specify various things (that may or may not be useful) such as a transformation matrix (TRANSFORM), which isn't at all necessary here but is useful if you've used SPSS's built-in contrasts, homogeneity tests (HOMOGENEITY(BARTLETT COCHRAN BOXM)), the main ANOVA and Huynh–Feldt and Greenhouse–Geisser corrections which we don't actually need in this example (SIGNIF(UNIV MULT AVERF HF GG)), and parameter estimates and effect size estimates for the contrasts we've specified (PARAM(ESTIM EFSIZE)).

So, the whole syntax will look like this:

```

MANOVA

before after BY treat(0 4)

/WSFACTORS time (2)

/CONTRAST (time)=special(1 1, 1 -1)

/CONTRAST (treat)=special (1 1 1 1 1, -4 1 1 1 1, 0 -3 1 1 1, 0 0 1 1 -2, 0 0 1 -1 0)

/CINTERVAL JOINT(.95) MULTIVARIATE(BONFER)

/METHOD UNIQUE

/ERROR WITHIN+RESIDUAL

/PRINT TRANSFORM HOMOGENEITY(BARTLETT COCHRAN BOXM)

SIGNIF( UNIV MULT AVERF HF GG )

PARAM( ESTIM EFSIZE).

```

It's very important to remember the full stop at the end! This syntax is in the file **DepressionSyntax.sps** as well, in case your typing goes wrong!

Output from the contrasts

The output you get is in the form of text (no nice pretty tables), and to interpret it you have to remember the contrasts you specified! I'll run you through the main highlights of this example. The first bit of the output will show the homogeneity tests (which should all be non-significant, but beware of Box's test because it tends to be inaccurate). The first important part is the main effect of the variable treat. First there's an ANOVA summary table like those you've come across before (if you've read Chapters 8–11). This tells us that there's no significant main effect

DISCOVERING STATISTICS USING SPSS

of the type of treatment, $F(4, 45) = 2.01, p = .11$. This means that if you ignore the time at which depression was measured then the levels of depression were about the same across the treatment groups. Of course, levels of depression should be the same before treatment, and so this isn't a surprising result (because it averages across scores before and after treatment). The graph shows that, in fact, levels of depression are relatively similar across groups.

* * * * * A n a l y s i s o f V a r i a n c e -- design 1 * * * * *

Tests of Between-Subjects Effects.


Tests of Significance for T1 using UNIQUE sums of squares					
Source of Variation	SS	DF	MS	F	Sig of F
WITHIN+RESIDUAL	359.95	45	8.00		
TREAT	64.30	4	16.08	2.01	.109

Estimates for T1
 --- Joint univariate .9500 BONFERRONI confidence intervals


TREAT

Parameter	Coeff.	Std. Err.	t-Value	Sig. t	Lower -95%	CL- Upper
2	-7.7781746	3.99972	-1.94468	.05808	-18.18578	2.62944
3	3.53553391	3.09817	1.14117	.25984	-4.52617	11.59723
4	3.74766594	2.19074	1.71069	.09403	-1.95282	9.44815
5	-.21213203	1.26482	-.16772	.86756	-3.50331	3.07904

Parameter	ETA Sq.
2	.07752
3	.02813
4	.06106
5	.00062



This main effect is followed by some contrasts, but we don't need to look at these because the main effect was non-significant. However, just to tell you what they are, parameter 2 is our first contrast (no treatment vs. the rest), and, as you can see, this is almost significant (p is just above 0.05); parameter 3 is our second contrast (placebo vs. the rest), and this is non-significant; parameter 4 is our third contrast (Cheerup vs. Effexor and Seroxat), and again this is almost significant; parameter 5 is our last contrast (Seroxat vs. Effexor), and this is very non-significant. However, these contrasts all ignore the effect of time and so aren't really what we're interested in.



The next part that we're interested in is the within-subject effects, and this involves the main effect of time and the interaction of time and treatment. First there's an ANOVA summary table as before. This tells us that there's a significant main effect of the time, $F(1, 45) = 43.02, p < .001$. This tells us that if you ignore the type of treatment, there was a significant difference between depression levels before and after treatment. A quick look at the means reveals that depression levels were significantly lower after treatment. Below the ANOVA table is a parameter estimate for the effect of time. As there are only two levels of time, this represents the difference in depression levels before and after treatment. No other contrasts are possible.

```

* * * * * A n a l y s i s   o f   V a r i a n c e  -- design 1 * * * * *
Tests involving 'TIME' Within-Subject Effect.

Tests of Significance for T2 using UNIQUE sums of squares
Source of Variation      SS      DF      MS      F      Sig of F
-----
WITHIN+RESIDUAL         320.35    45      7.12
TIME                     306.25     1     306.25    43.02    .000
TREAT BY TIME           125.90     4      31.47     4.42    .004
-----
Estimates for T2
    
```

DISCOVERING STATISTICS USING SPSS

--- Joint univariate .9500 BONFERRONI confidence intervals

TIME


Parameter	Coeff.	Std. Err.	t-Value	Sig.	t Lower -95%	CL- Upper
1	2.47487373	.37733	6.55891	.00000	1.71489	3.23485

Parameter	ETA Sq.
1	.48875

TREAT BY TIME

Parameter	Coeff.	Std. Err.	t-Value	Sig.	t Lower -95%	CL- Upper
2	11.3137085	3.77330	2.99836	.00441	1.49527	21.13214
3	-.56568542	2.92278	-.19354	.84740	-8.17101	7.03964
4	-5.8689863	2.06672	-2.83976	.00675	-11.24676	-.49121
5	.919238816	1.19322	.77038	.44510	-2.18562	4.02410

Parameter	ETA Sq.
2	.16651
3	.00083
4	.15197
5	.01302



The interaction term is also significant, $F(4, 45) = 4.42, p = .004$. This indicates that the change in depression over time is different in some treatments to others. We can make sense of this through an interaction graph, but we can also look at our contrasts. The key contrasts for this whole analysis are the parameter estimates for the interaction term (the bit in the output underneath the heading TREAT BY TIME) because they take into account the effect of time *and* treatment:

- ✓ Parameter 2 is our first contrast (no-treatment vs. the rest), and, as you can see, this is significant (p is below 0.05). This tells us that the change in depression levels in the no-treatment group was significantly different from the average change in all other groups, $t = 2.30$, $p = .004$. As you can see in the graph, there is no change in depression in the no-treatment group, but in all other groups there is a fall in depression. Therefore, this contrast reflects the fact that there is no change in the no-treatment group, but there is a decrease in depression levels in all other groups.
- ✓ Parameter 3 is our second contrast (placebo vs. Seroxat, Effexor and Cheerup), and this is very non-significant, $t = -.19$, $p = .847$. This shows that the decrease in depression levels seen in the placebo group is comparable to the average decrease in depression levels seen in the Seroxat, Effexor and Cheerup conditions. In other words, the combined effect of the drugs on depression is no better than a placebo.
- ✓ Parameter 4 is our third contrast (Cheerup vs. Effexor and Seroxat), and this is highly significant, $t = -2.84$, $p = .007$. This shows that the decrease in depression levels seen in the Cheerup group is significantly bigger than the decrease seen in the Effexor and Seroxat groups combined. Put another way, Cheerup has a significantly bigger effect than other established antidepressants.
- ✓ Parameter 5 is our last contrast (Seroxat vs. Effexor), and this is very non-significant, $t = .77$, $p = .445$. This tells us that the decrease in depression levels seen in the Seroxat group is comparable to the decrease in depression levels seen in the Effexor group. Put another way, Effexor and Seroxat seem to have similar effects on depression.

I hope to have shown in this example how to specify contrasts using syntax and how looking at these contrasts (especially for an interaction term) can be a very useful way to break down an interaction effect.

Please, Sir, can I have some more ... simple effects?



Calculating simple effects

A simple main effect (usually called a simple effect) is just the effect of one variable at levels of another variable. In Chapter 12 we had an example in which we'd measured the attractiveness of dates after no alcohol, 2 pints and 4 pints in both men and women. Therefore, we have two independent variables: alcohol (none, 2 pints, 4 pints) and gender (male and female). One simple effects analysis we could do would be to look at the effect of gender (i.e., compare male and female scores) at the three levels of alcohol. Let's look how we'd do this. We're partitioning the model sum of squares, and we saw in Chapter 10 that we calculate model sums of squares using this equation:

$$SS_M = \sum_{n=1}^k n_k (\bar{x}_k - \bar{x}_{\text{grand}})^2$$

For simple effects, we calculate the model sum of squares for the effect of gender at each level of alcohol. So, we'd begin with when there was no alcohol, and calculate the model sum

of squares. Thus the grand mean becomes the mean for when there was no alcohol, and the group means are the means for men (when there was no alcohol) and women (when there was no alcohol). So, we group the data by the amount of alcohol drunk. Within each of these three groups, we calculate the overall mean and also the mean of the male and female scores separately. These mean scores are all we really need. Pictorially, you can think of the data as displayed pictorially below.

No Alcohol		2 Pints		4 Pints	
Female	Male	Female	Male	Female	Male
65	50	70	45	55	30
70	55	65	60	65	30
60	80	60	85	70	30
60	65	70	65	55	55
60	70	65	70	55	35
55	75	60	70	60	20
60	75	60	80	50	45
55	65	50	60	50	40
60.625	66.875	62.50	66.875	57.500	35.625
Mean None = 63.75		Mean 2 Pints = 64.6875		Mean 4 Pints = 46.5625	

We can then apply the same equation for the model sum of squares that we used for the overall model sum of squares, but we use the grand mean of the no-alcohol data (63.75) and the means of males (66.875) and females (60.625) within this group:

$$\begin{aligned}
 SS_{\text{Gender (No Alcohol)}} &= \sum n_k (\bar{x}_k - \bar{x}_{\text{grand}})^2 \\
 &= 8(60.625 - 63.75)^2 + 8(66.875 - 63.75)^2 \\
 &= 156.25
 \end{aligned}$$

The degrees of freedom for this effect are calculated the same way as for any model sum of squares; that is, they are one less than the number of conditions being compared ($k - 1$), which in this case, where we're comparing only two conditions, will be 1.

The next step is to do the same but for the 2-pints data. Now we use the grand mean of the 2-pints data (64.6875) and the means of males (66.875) and females (62.50) within this group. The equation, however, stays the same:

$$\begin{aligned} SS_{\text{Gender (2 Pints)}} &= \sum n_k (\bar{x}_k - \bar{x}_{\text{grand}})^2 \\ &= 8(62.50 - 64.6875)^2 + 8(66.875 - 64.6875)^2 \\ &= 76.56 \end{aligned}$$

The degrees of freedom are the same as in the previous simple effect, namely $k - 1$, which is 1 for these data.

The next step is to do the same but for the 4-pints data. Now we use the grand mean of the 4-pints data (46.5625) and the means of females (57.500) and males (35.625) within this group. The equation, however, stays the same:

$$\begin{aligned} SS_{\text{Gender (4 Pints)}} &= \sum n_k (\bar{x}_k - \bar{x}_{\text{grand}})^2 \\ &= 8(57.50 - 46.5625)^2 + 8(35.625 - 46.5625)^2 \\ &= 1914.06 \end{aligned}$$

Again, the degrees of freedom are 1 (because we've compared two groups).

As with any ANOVA, we need to convert these sums of squares to mean squares by dividing by the degrees of freedom. However, because all of these sums of squares have 1 degree of freedom, the mean squares will be the same as the sum of squares because we're dividing by 1. So, the final stage is to calculate an F -ratio for each simple effect. As ever, the F -ratio is just the mean squares for the model divided by the residual mean squares. So, you might well ask, what do we use for the residual mean squares? When conducting simple effects we use the residual mean squares for the original ANOVA (the residual mean squares for the entire model). In doing so we are merely partitioning the model sums of squares and so keep control of the Type I error rate. For these data, the residual sum of squares was 83.036 (see Section 13.2.7). Therefore, we get:

$$F_{\text{Gender(No Alcohol)}} = \frac{MS_{\text{Gender(No Alcohol)}}}{MS_R} = \frac{156.25}{83.036} = 1.88$$

$$F_{\text{Gender(2 Pints)}} = \frac{MS_{\text{Gender(2 Pints)}}}{MS_R} = \frac{76.56}{83.036} = 0.92$$

$$F_{\text{Gender(4 Pints)}} = \frac{MS_{\text{Gender(4 Pints)}}}{MS_R} = \frac{1914.06}{83.036} = 23.05$$

DISCOVERING STATISTICS USING SPSS

We can evaluate these F -values in the usual way (they will have 1 and 42 degrees of freedom for these data). However, for the 2-pints data we can be sure there is not a significant effect of gender because the F -ratio is less than 1.