Repeated Measures ANOVA

Introduction

Repeated measures is a term used when the same participants participate in all conditions of an experiment. So, for example, you might want to test the effects of alcohol on enjoyment of a party. In this type of experiment it is important to control for individual differences in tolerance to alcohol: some people can drink a lot of alcohol without really feeling the consequences, whereas others, like me, only have to sniff a pint of lager and they fall to the floor and pretend to be a fish. To control for these individual differences we can test the same people in all conditions of the experiment: so we would test each subject after they had consumed one pint, two pints, three pints and four pints of lager. After each drink the subject might be given a questionnaire assessing their enjoyment of the party. Therefore, every subject provides a score representing their enjoyment before the study (no alcohol consumed), after one pint, after two pints, and so on. This design is said to use repeated measures.

What is Sphericity?

This type of design has one several advantages. Most important it reduces the unsystematic variability in the data (see Field, 2005, Chapter 7) and so provides greater power to detect effects. Repeated measures are also more economical because fewer participants are required in total. However, there is a disadvantage too. In week 1 (exploring data) we saw that tests based on parametric data assume that data points are independent. This is not the case in a repeated measures design because data for different conditions have come from the same people. This means that data from different experimental conditions will be related; because of this we have to make an additional assumption to those of the independent ANOVAs you have so far studied. Put simply (and not entirely accurately), we assume that the relationship between pairs of experimental conditions is similar (i.e. the level of dependence between pairs of groups is roughly equal). This assumption is known as the assumption of *sphericity*. (If you want the less simple but entirely accurate explanation then see Field, 2005, section 11.2.1).

The assumption of sphericity can be likened to the assumption of homogeneity of variance (see your handout on exploring data): if you were to take each pair of treatment levels, and calculate the differences between each pair of scores, then it is necessary that these differences have equal variances (see Field, 2005).

What is the Effect of Violating the Assumption of Sphericity?

The effect of violating sphericity is a loss of power (i.e. an increased probability of a Type II error) and a test statistic (F-ratio) that simply cannot be compared to tabulated values of the F-distribution (for more details see Field, 1998, 2005).

Assessing the Severity of Departures from Sphericity

SPSS produces a test known as Mauchly's test, which tests the hypothesis that the variances of the differences between conditions are equal.

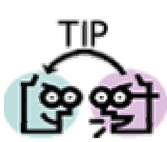

- ✓ If Mauchly's test statistic is significant (i.e. has a probability value less than .05) we conclude that there are significant differences between the variance of differences: the condition of sphericity has not been met.
- ✓ If, Mauchly's test statistic is nonsignificant (i.e. $p > .05$) then it is reasonable to conclude that the variances of differences are not

significantly different (i.e. they are roughly equal).

- ✓ If Mauchly's test is significant then we cannot trust the *F*-ratios produced by SPSS.

Correcting for Violations of Sphericity

Fortunately, if data violate the sphericity assumption there are several corrections that can be applied to produce a valid *F*-ratio. All of these corrections involve adjusting the degrees of freedom associated with the *F*-value. In all cases the degrees of freedom are reduced based on an estimate of how 'spherical' the data are; by reducing the degrees of freedom we make the *F*-ratio more conservative (i.e. it has to be bigger to be deemed significant). There are three different estimates of sphericity used to correct the degrees of freedom:

1. Greenhouse and Geisser's(1958)
2. Huynh and Feldt's (1976)
3. The Lower Bound estimate

For more details on these estimates see Field (2005) or Girden (1992).

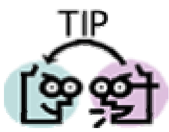

Which correction should I use?

- ✓ Look at the estimates of sphericity (ε) in the SPSS handout.
- ✓ When ε > .75 then use the Huynh-Feldt correction.
- ✓ When ε < 0.75, or nothing is known about sphericity at all, then use the Greenhouse-Geisser correction.

One-Way Repeated Measures ANOVA using SPSS

All is not well in Lapland. The organisation 'Statisticians Hate Interesting Things' have executed their long planned Campaign Against Christmas by abducting Santa. Spearheaded by their evil leader Professor N. O. Life, a chubby bearded man with a penchant for red jumpers who gets really envious of other chubby bearded men that people actually like, and his crack S.W.A.T. team (Statisticians With Autistic Tendencies), they have taken Santa from his home and have bricked him up behind copies of a rather heavy and immoveable Stats textbook written by some brainless gibbon called 'Field'.

It's Christmas Eve and the elves are worried. The elf leader, twallybliddle (don't blame me, I didn't name him …) has rallied his elf troops and using Rudolph's incredibly powerful nose, they tracked down the base of S.H.I.T. and planned to break down the door. They then realised they didn't know what a door was and went down the chimney instead, much to the surprise of a room full of sweaty men with abacuses. Armed with a proof of Reimann's Hypothesis[1] they overcame the bemused huddle of statisticians and located Santa. They slowly peeled away the tower of books. One by one, the barrier came down until they could see he tip of a red hat, they could hear a hearty chuckle. Imagine their surprise as the last book was removed revealing three identical Santas …

[1] There is currently a $1 million prize on offer to prove Reimann's hypothesis, which for reasons I won't bore you with is a very important hypothesis—the proof of which has eluded the greatest mathematicians (including Reimann) of the last 150 years or so. An elf popping out of a chimney with proof of this theory would definitely send a room full of mathematicians into apoplexy.

Somehow, they had to identify the 'real' Santa and return him to Lapland in time to deliver the Christmas presents, what should they do? They decided that each elf in turn would approach all three of the Santas (in counterbalanced order obviously …) and would stroke his beard, sniff him and tickle him under the arms. Having done this the elf would give the Santa a rating from 0 (definitely not the real Santa) to 10 (definitely the real Santa).

Table 1: Data for the Santa example

Ratings of Santa 1	Ratings of Santa 2	Ratings of Santa 3
1	3	1
2	5	3
4	6	6
5	7	4
5	9	1
6	10	3

Entering the Data

The independent variable was the Santa that was being assessed by the elf (Santa 1 2 or 3) and the dependent variable was the elf's rating out of 10.

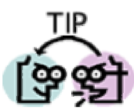

✓ Levels of repeated measures variables go in different columns of the SPSS data editor.

Therefore, separate columns should represent each level of a repeated measures variable. As such, there is no need for a coding variable (as with between-group designs). The data can, therefore, be entered as they are in Table 1.

✓ Save these data in a file called **TheRealSanta.sav**

To conduct an ANOVA using a repeated measures design, select the *define factors* dialog box by following the menu path **Analyze⇒General Linear Model⇒GLM-Repeated Measures …**.

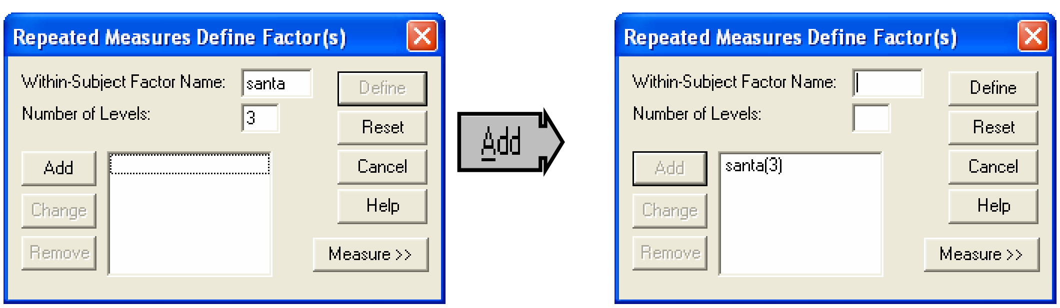

Figure 1: *Define Factors* dialog box for repeated measures ANOVA

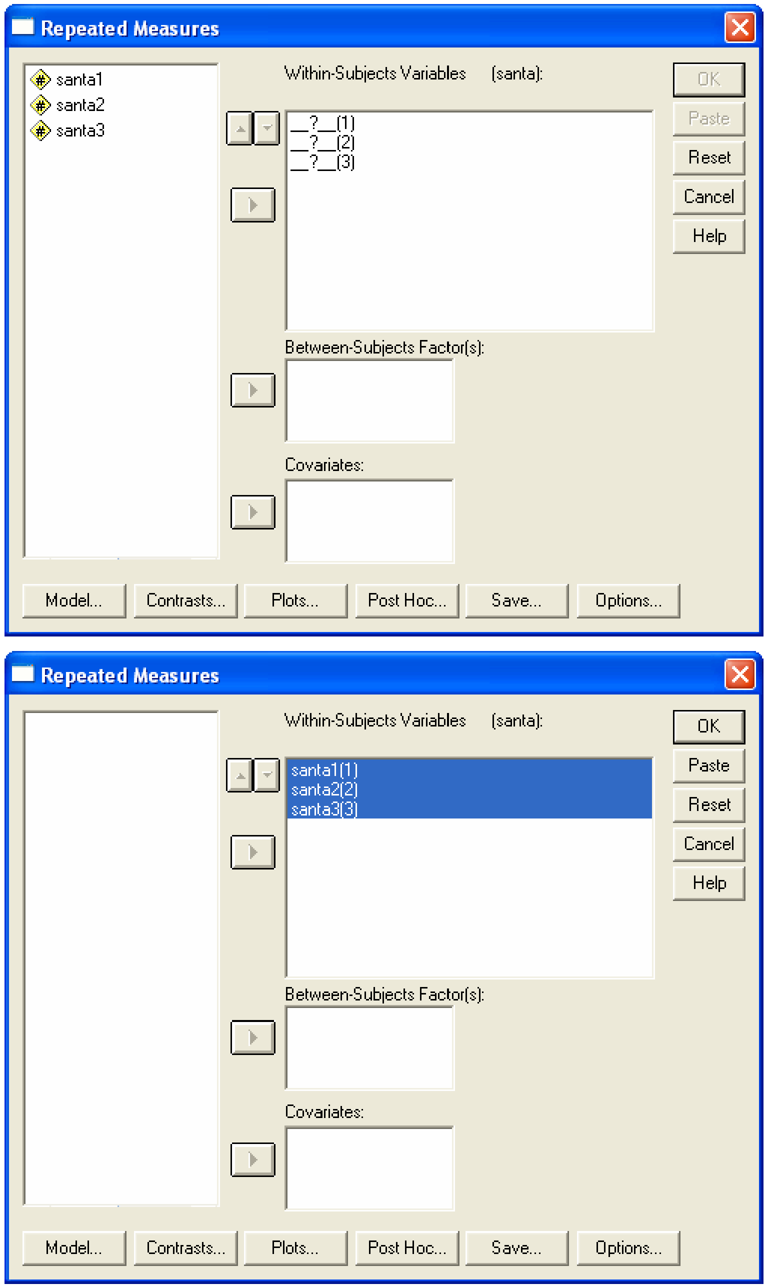

Figure 2: Main dialog box for repeated measures ANOVA

In the main dialog box, you are asked to supply a name for the within-subject (repeated-measures) variable. In this case the repeated measures variable was the Santa that the Elves tested, so replace the word *factor1* with the word *Santa*. The name you give to the repeated measures variable is restricted to 8 characters. When you have given the repeated measures factor a name, you have to tell the computer how many levels there were to that variable (i.e. how many experimental conditions there were). In this case, there were 3 different Santas that the Elves had to rate, so we have to enter the number 3 into the box labelled *Number of Levels*. Click on Add to add this variable to the list of repeated measures variables. This variable will now appear in the white box at the bottom of the dialog box and appears as *Santa(3)*. If your design has several repeated measures variables then you can add more factors to the list. When you have entered all of the repeated measures factors that were measured click on Define to go to the *Main Dialog Box*.

The Main dialog box has a space labelled *within subjects variable list* that contains a list of 3 question marks proceeded by a number. These question marks are for the variables representing the 3 levels of the independent variable. The variables corresponding to these levels should be selected and placed in the appropriate space. We have only 3 variables in the data editor, so it is possible to select all three variables at once (by clicking on the variable at the top, holding the mouse button down and dragging down over the other variables). The selected variables can then be transferred by clicking on ▸.

When all three variables have been transferred, you can select various options for the analysis. There are several options that can be accessed with the buttons at the bottom of the main dialog box. These options are similar to the ones we have already encountered.

Post Hoc Tests

There is no proper facility for producing post hoc tests for repeated measures variables in SPSS (but see below)! However, you can use the paired *t*-test procedure to compare all pairs of levels of the independent variable, and then apply a Bonferroni correction to the probability at which you accept any of these tests. The resulting probability value should be used as the criterion for statistical significance.

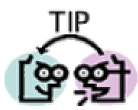

✓ A 'Bonferroni correction' is achieved by dividing the probability value (usually .05) by the number of tests conducted.

For example, if we compared all levels of the independent variable of these data, we would make 3 comparisons in all and so the appropriate significance level would be .05/3 = .0167. Therefore, we would accept *t*-tests as being significant only if they have a *p* value that is less than .0167. One way to salvage what power you can from this procedure is to compare only the pairs of groups between which you expect differences to arise (rather than comparing all pairs of treatment levels). The fewer tests you perform, the less you have to correct the significance level, and the more power you retain.

Additional Options

The final options, that haven't previously been described, can be accessed by clicking Options... in the main dialog box. The *options* dialog box (Figure 3) has various useful options. You can ask for descriptive statistics, which will provide the means, standard deviations and number of participants for each level of the independent variable. The option for homogeneity of variance tests will be active only when there is a between group factor as well (Mixed designs, which are covered next term).

Perhaps the most useful feature is that you can get some *post hoc* tests via this dialog box. To specify *post hoc* tests, select the repeated measures variable (in this case **Santa**) 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 is not recommended). The second option is a Bonferroni correction (recommended for the reasons mentioned above), and the final option is a Sidak correction, which should be selected if you are concerned about the loss of power associated with Bonferroni corrected values.

When you have selected the options of interest, click on Continue to return to the main dialog box, and then click on OK to run the analysis.

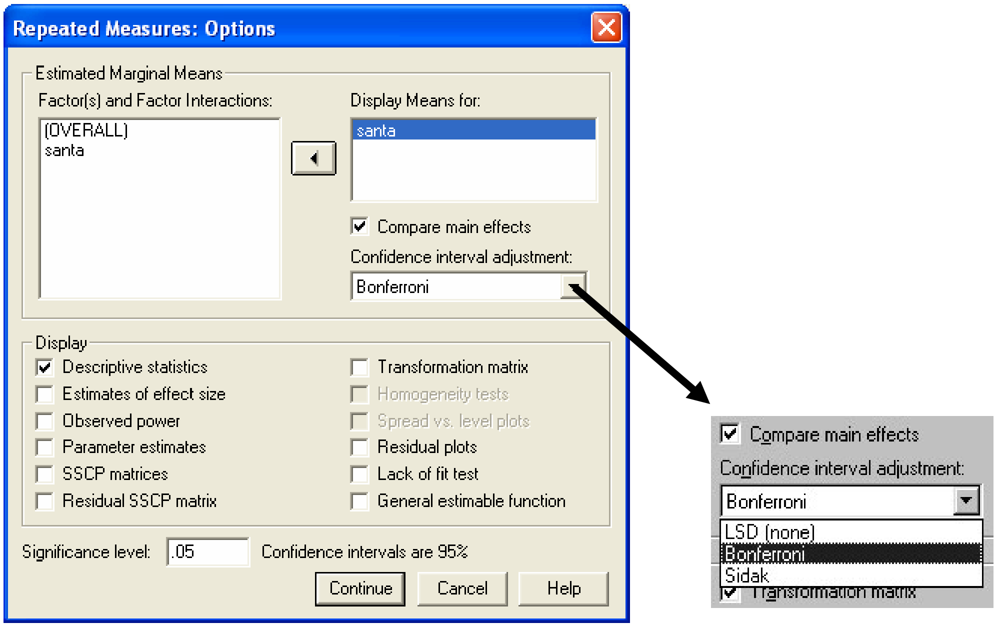

Figure3**: Options dialog box**

Output for Repeated Measures ANOVA

Descriptive statistics and other Diagnostics

Within-Subjects Factors

Measure: MEASURE_1

SANTA	Dependent Variable
1	SANTA1
2	SANTA2
3	SANTA3

Descriptive Statistics

	Mean	Std. Deviation	N
SANTA1	3.83	1.941	6
SANTA2	6.67	2.582	6
SANTA3	3.00	1.897	6

SPSS Output 1

SPSS Output 1 shows the initial diagnostics statistics. First, we are told the variables that represent each level of the independent variable. This box is useful mainly to check that the variables were entered in the correct order. The following table provides basic descriptive statistics for the four levels of the independent variable. From this table we can see that, on average, the Elves rated Santa 2 highest (i.e. most likely to be the real Santa) and rated Santa 3 as being the least likely to be the real Santa (he received the lowest ratings from the elves).

Assessing Sphericity

Earlier you were told that SPSS produces a test that looks at whether the data have violated the assumption of sphericity. The next part of the output contains information about this test.

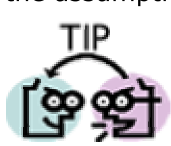

✓ Mauchly's test should be nonsignificant if we are to assume that the condition of sphericity has been met.

SPSS Output 2 shows Mauchly's test for the Santa data, and the important column is the one containing the significance vale. The significance value is .009, which is less than .05, so we must accept the hypothesis that the variances of the differences between levels were significantly different. In other words the assumption of sphericity has been violated.

Mauchly's Test of Sphericity[b]

Measure: MEASURE_1

Within Subjects Effect	Mauchly's W	Approx. Chi-Square	df	Sig.	Epsilon[a] Greenhouse-Geisser	Huynh-Feldt	Lower-bound
SANTA	.094	9.437	2	.009	.525	.544	.500

Tests the null hypothesis that the error covariance matrix of the orthonormalized transformed dependent variables is proportional to an identity matrix.

a. May be used to adjust the degrees of freedom for the averaged tests of significance. Corrected tests are displayed in the Tests of Within-Subjects Effects table.

b.
Design: Intercept
Within Subjects Design: SANTA

SPSS Output 2

The Main ANOVA

SPSS Output 3 shows the results of the ANOVA for the within-subjects variable. The table you see will look slightly different (it will look like SPSS Output 4 in fact), but for the time being I've simplified it a bit. Bear with me for now. This table can be read much the same as for One-way between-group ANOVA (see your handout). There is a sum of squares for the within-subject effect of Santa, which tells us how much of the total variability is explained by the experimental effect (i.e. differences in ratings of the three Santas). There is also an error term, which is the amount of unexplained variation across the conditions of the repeated measures variable. These sums of squares are converted into mean squares by dividing by the degrees of freedom[2].

The *F*-ratio is obtained by dividing the mean squares for the experimental effect (22.167) by the error mean squares (2.7). As with between-group ANOVA, this test statistic represents the ratio of systematic variance to unsystematic variance. The value of the *F*-ratio (22.167/2.7 = 8.21) is then compared against a critical value for 2 and 10 degrees of freedom. SPSS displays the exact significance level for the *F*-ratio. The significance of *F* is .008 which is significant because it is less than the criterion value of .05. We can, therefore, conclude that there was a significant difference between the elves' ratings of the three different Santas. However, this main test does not tell us which Santa's differed from each other in the ratings awarded to them by the elves.

Tests of Within-Subjects Effects

Measure: MEASURE_1
Sphericity Assumed

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
SANTA	44.333	2	22.167	8.210	.008
Error(SANTA)	27.000	10	2.700		

SPSS Output 3

[2] If you're interested in how the degrees of freedom are calculated read my book (Chapter 11).

Although this result seems very plausible, we saw earlier that the assumption of sphericity had been violated. I also mentioned that a violation of the sphericity assumption makes the *F*-test inaccurate. So, what do we do? Well, I mentioned earlier on that we can correct the degrees of freedom in such a way that it is accurate when sphericity is violated. This is what SPSS does. SPSS Output 4 (which is the output you will see in your own SPSS analysis) shows the main ANOVA. As you can see in this output, the value of *F* does not change, only the degrees of freedom[3]. But the effect of changing the degrees of freedom is that the *significance* of the value of *F* changes: the effect of Santa is less significant after correcting for sphericity.

Tests of Within-Subjects Effects

Measure: MEASURE_1

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
SANTA	Sphericity Assumed	44.333	2	22.167	8.210	.008
	Greenhouse-Geisser	44.333	1.050	42.239	8.210	.033
	Huynh-Feldt	44.333	1.088	40.752	8.210	.031
	Lower-bound	44.333	1.000	44.333	8.210	.035
Error(SANTA)	Sphericity Assumed	27.000	10	2.700		
	Greenhouse-Geisser	27.000	5.248	5.145		
	Huynh-Feldt	27.000	5.439	4.964		
	Lower-bound	27.000	5.000	5.400		

SPSS Output 4

The next issue is which of the three corrections to use. Earlier I gave you some tips and they were that when ε > .75 then use the Huynh-Feldt correction, and when ε < 0.75, or nothing is known about sphericity at all, then use the Greenhouse-Geisser correction. *E* is the estimate of sphericity from SPSS output 2 and these values are .525 and .544; because these values are less than .75 we should use the Greenhouse-Geisser corrected values. Using this correction, *F* is still significant because its *p* value is .033, which is less than the normal criterion of .05.

Post Hoc Tests

Pairwise Comparisons

Measure: MEASURE_1

(I) SANTA	(J) SANTA	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a] Lower Bound	Upper Bound
1	2	-2.833*	.401	.003	-4.252	-1.415
	3	.833	.946	1.000	-2.509	4.176
2	1	2.833*	.401	.003	1.415	4.252
	3	3.667	1.282	.106	-.865	8.199
3	1	-.833	.946	1.000	-4.176	2.509
	2	-3.667	1.282	.106	-8.199	.865

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Bonferroni.

SPSS Output 5

The arrangement of the table in SPSS Output 5 is similar to the table produced for between-group *post hoc* tests (see your handout on this of Field, 2005): the difference between group means is displayed, the standard error, the significance value and a confidence interval for the difference between means. By looking at the significance values we can see that the only

[3] SPSS corrects the degrees of freedom by multiplying them by the estimates of sphericity in SPSS Output 2. If you want a more detailed explanation and an example see: http://www.sussex.ac.uk/Users/andyf/research/articles/sphericity.pdf

difference between group means is between Santa 1 and Santa 2. Looking at the means of these groups (SPSS Output 1) we can see that Santa 2 was rated significantly more highly than Santa 1. So, although the Elves might confidently conclude that Santa 2 is more likely to be the real Santa than Santa 1, there is nothing in the data to suggest that they were more confident that Santa 2 was the real Santa compared to Santa 3. This provides them with something of a dilemma.

Reporting One-Way Repeated Measures ANOVA

We can report repeated measures ANOVA in the same way as an independent ANOVA (see your handout). The only additional thing we should concern ourselves with is reporting the corrected degrees of freedom if sphericity was violated. Personally, I'm also keen on reporting the results of sphericity tests as well. Therefore, we could report the main finding as:

- ✓ Mauchly's test indicated that the assumption of sphericity had been violated ($\chi^2(2) = 9.44$, $p < .05$), therefore degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity ($\varepsilon = .53$). The results show that the elves' ratings of the three Santas differed significantly, $F(1.05, 5.25) = 8.21$, $p < .05$. Post hoc tests revealed that although Elves rated Santa 2 as significantly more likely to be the real Santa than Santa 1 ($P < .01$), Santa 3 was not rated significantly differently from either of the other Santas (both *ps* > .05).

A Happy Ending?

Well, you're probably wondering what happened? The Elves had no firm evidence about which Santa they thought was the real one. Fortunately, just as the elves were about to admit defeat Rudolph, who'd been enjoying a nice bucket water outside, walked in and licked the three Santas (which the Elves couldn't do because of Elf & Safety regulations …) and identified Santa 2 as the real one (the Real Santa tastes of Christmas pudding apparently … you can put it to the test in a couple of weeks should you want to …). They got him back to Lapland in time, and all the Children got their presents and the world was a happy place. Hooray!

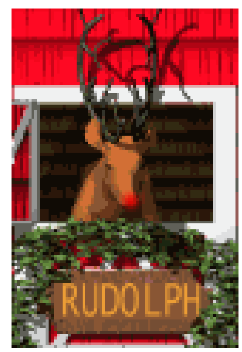

Two-Way Repeated Measures ANOVA Using SPSS

As we have seen before, the name of any ANOVA can be broken down to tell us the type of design that was used. The 'two-way' part of the name simply means that two independent variables have been manipulated in the experiment. The 'repeated measures' part of the name tells us that the same participants have been used in all conditions. Therefore, this analysis is appropriate when you have two repeated-measures independent variables: each participant does all of the conditions in the experiment, and provides a score for each permutation of the two variables.

An Example

At Christmas we normally leave treats for Santa Claus and his helpers (mince pies, a glass of sherry and a bucket of water for Rudolph). Santa Claus noticed that he was struggling to deliver all the presents on Christmas Eve and wondered whether these treats might be slowing down his Elves[4]. So, one Christmas Santa did a little experiment. He randomly selected 10 Elves from his workforce and timed how long it took each of them to deliver the presents

[4] He was starting to wonder if the alcohol and cakes were bad for their 'elf … (boy, am I going to get some mileage out of that pitiful attempt at a joke).

to 5 houses. Half of the elves were told that they could eat any mince pies or Christmas pudding but that they must not have any sherry, while the other half were told to drink sherry but not to eat any food that was left for them. The following year Santa took the same 10 elves and again timed how long it took them to deliver presents to the same 5 houses as the previous year. This time, however, the ones who had drunk sherry the previous year were banned from drinking it and told instead to eat any mince pies or Christmas pudding. Conversely, the ones who had eaten treats the year before were told this year only to drink sherry and not to eat any treats. As such, over the two years each of the 10 elves was timed for their speed of present delivery after 1, 2, 3, 4 and 5 doses of sherry, and also after 1, 2, 3, 4, & 5 doses of mince pies.

✓ Why do you think Santa got half of the elves to drink sherry the first year and ate treats the second year, while the other half ate treats the first year and drank sherry the second year?

Think about the design of this study for a moment. We have the following variables:

- **Treat**: Independent Variable 1 is the treat that was consumed by the elves and it has 2 levels: Sherry or Mince Pies.
- **Dose**: Independent Variable 2 is the dose of the treat (remember each elf had a treat at the five houses to which they delivered and so the total quantity consumed increased across the houses). This variable has 5 levels: house 1, house 2, house 3, house 4 & house 5.
- The dependent variable was the time taken to deliver the presents to a given house (in nanoseconds: elves deliver very quickly!)

These data could, therefore, be analysed with a 2 × 5 two-way repeated-measures ANOVA. As with other ANOVA designs, there is in principle no limit to the number of conditions for each of the independent variables in the experiment. In practice, however, you'll find that your participants get very bored and inattentive if there were too many conditions (although if you're feeding them cake and sherry then perhaps not …)!

Treat:	**Sherry**					**Mince Pies**				
Dose:	**1**	**2**	**3**	**4**	**5**	**1**	**2**	**3**	**4**	**5**
Bingo	10	15	18	22	37	9	13	13	18	22
Flingo	10	18	10	42	60	7	14	20	21	32
Lardy	7	11	28	31	56	9	13	24	30	35
Alchy	9	19	36	45	60	7	14	9	20	25
Goody	15	14	29	33	37	14	13	20	22	29
Pongo	14	13	26	26	49	5	12	17	16	33
Beadle	9	12	19	37	48	5	15	12	17	24
Groucho	9	18	22	31	39	13	13	14	17	17
Chunder	12	14	24	28	53	12	13	21	19	22
Flouncy	7	11	21	23	45	12	14	20	21	29
Mean	**10.2**	**14.5**	**23.3**	**31.8**	**48.4**	**9.3**	**13.4**	**17.0**	**20.1**	**26.8**

Table 2: Data for example two

Entering the Data

As with the previous example, To enter these data into SPSS we need to recap the golden rule of the data editor:

- ✓ Each row represents a single participant's data.

If a person participates in all experimental conditions (in this case all elves experience both Sherry and Mince Pies in the different doses) then each experimental condition must be represented by a column in the data editor. In this experiment there are ten experimental conditions and so the data need to be entered in ten columns (so, the format is identical to the original table in which I put the data). You should create the following ten variables in the data editor with the names as given. For each one, you should also enter a full variable name for clarity in the output.

Sherry1	1 Dose of Sherry
Sherry2	2 Doses of Sherry
Sherry3	3 Doses of Sherry
Sherry4	4 Doses of Sherry
Sherry5	5 Doses of Sherry
Pie1	1 Mince Pie
Pie2	2 Mince Pies
Pie3	3 Mince Pies
Pie4	4 Mince Pies
Pie5	5 Mince Pies

Once these variables have been created, enter the data as in Table 2 (above) and save the file onto a disk with the name **santa.sav**.

Running the analysis

The analysis is run in the same way as for one-way repeated measures ANOVA: access the *define factors* dialog box use the menu path **Analyze⇒General Linear Model⇒Repeated Measures …**. In the *define factors* dialog box you are asked to supply a name for the within-subject (repeated measures) variable. In this case there are two within-subject factors: **treat** (Sherry or Mince Pie) and **dose** (1, 2, 3 4 or 5 doses). Replace the word *factor1* with the word *treat*. When you have given this repeated measures factor a name, you have to tell the computer how many levels there were to that variable. In this case, there were two types of treat, so we have to enter the number 2 into the box labelled *Number of Levels*. Click on Add to add this variable to the list of repeated measures variables. This variable will now appear in the white box at the bottom of the dialog box and appears as *treat(2)*. We now have to repeat this process for the second independent variable. Enter the word *dose* into the space labelled *Within-Subject Factor Name* and then, because there were five levels of this variable, enter the number 5 into the space labelled *Number of Levels*. Click on Add to include this variable in the list of factors; it will appear as *dose(5)*. The finished dialog box is shown in Figure 3. When you have entered both of the within-subject factors click on Define to go to the main dialog box.

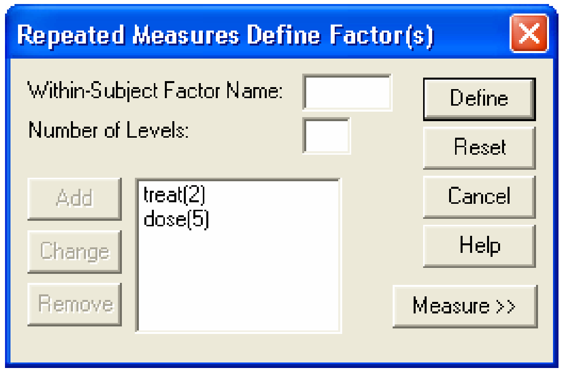

Figure 3: *Define factors* dialog box for factorial repeated measures ANOVA

The main dialog box is essentially the same as when there is only one independent variable (see previous handout) except that there are now ten question marks (Figure 4). At the top of the *Within-Subjects Variables* box, SPSS states that there are two factors: **treat** and **dose**. In the box below there is a series of question marks followed by bracketed numbers. The numbers in brackets represent the levels of the factors (independent variables).

_?_(1,1)	variable representing 1st level of treat and 1st level of dose
_?_(1,2)	variable representing 1st level of treat and 2nd level of dose
_?_(1,3)	variable representing 1st level of treat and 3rd level of dose
_?_(1,4)	variable representing 1st level of treat and 4th level of dose
_?_(1,5)	variable representing 1st level of treat and 5th level of dose
_?_(2,1)	variable representing 2nd level of treat and 1st level of dose
_?_(2,2)	variable representing 2nd level of treat and 2nd level of dose
_?_(2,3)	variable representing 2nd level of treat and 3rd level of dose
_?_(2,4)	variable representing 2nd level of treat and 4th level of dose
_?_(2,5)	variable representing 2nd level of treat and 5th level of dose

In this example, there are two independent variables and so there are two numbers in the brackets. The first number refers to levels of the first factor listed above the box (in this case **treat**). The second number in the bracket refers to levels of the second factor listed above the box (in this case **dose**). As with one-way repeated measures ANOVA, you are required to replace these question marks with variables from the list on the left-hand side of the dialog box. With between-group designs, in which coding variables are used, the levels of a particular factor are specified by the codes assigned to them in the data editor. However, in repeated measures designs, no such coding scheme is used and so we determine which condition to assign to a level at this stage. For example, if we entered **sherry1** into the list first, then SPSS will treat sherry as the first level of **treat**, and dose 1 as the first level of the **dose** variable. However, if we entered **pie5** into the list first, SPSS would consider mince pies as the first level of **treat**, and dose 5 as the first level of **dose**

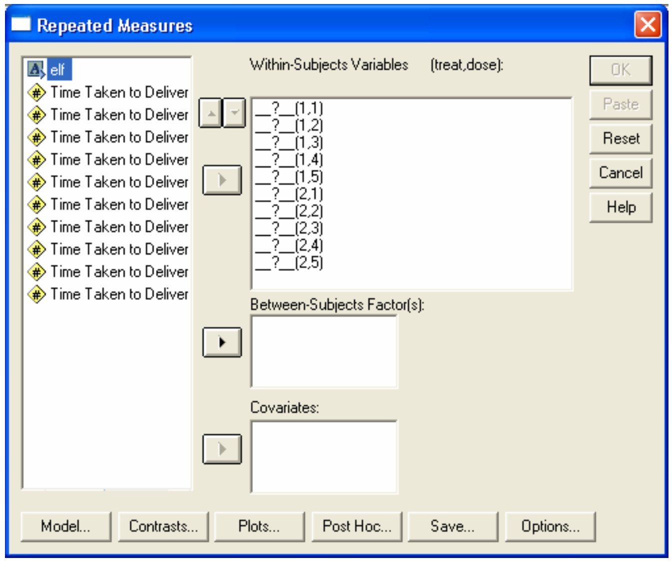

Figure 4

It should be reasonably obvious that it doesn't really matter which way round we specify the treats, but is very important that we specify the doses in the correct order. Therefore, the variables could be entered as follows:

Sherry1 _?_(1,1)
Sherry2 _?_(1,2)
Sherry3 _?_(1,3)
Sherry4 _?_(1,4)
Sherry5 _?_(1,5)
Pie1 _?_(2,1)
Pie2 _?_(2,2)
Pie3 _?_(2,3)
Pie4 _?_(2,4)
Pie5 _?_(2,5)

When these variables have been transferred, the dialog box should look exactly like Figure 5. The buttons at the bottom of the screen have already been described for the one independent variable case and so I will describe only the most relevant.

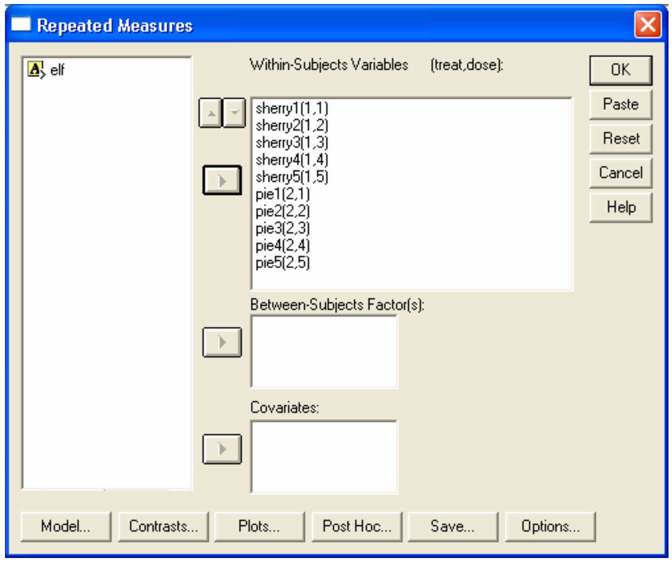

Figure 5

Graphing Interactions

When we had only one independent variable, we ignored the *plots* dialog box; however, if there are two or more factors, the *plots* dialog box is a convenient way to plot the means for each level of the factors. This plot will be useful for interpreting the meaning of the treat × dose interaction. To access this dialog box click on Plots... . Select **dose** from the variables list on the left-hand side of the dialog box and transfer it to the space labelled *Horizontal Axis* by clicking on ▶. In the space labelled *Separate Lines* we need to place the remaining independent variable: **treat**. As before, it is down to your discretion which way round the graph is plotted, but it actually makes sense this time to have dose on the horizontal axis. When you have moved the two independent variables to the appropriate box, click on Add and this interaction graph will be added to the list at the bottom of the box (see Figure 6). When you have finished specifying graphs, click on Continue to return to the main dialog box.

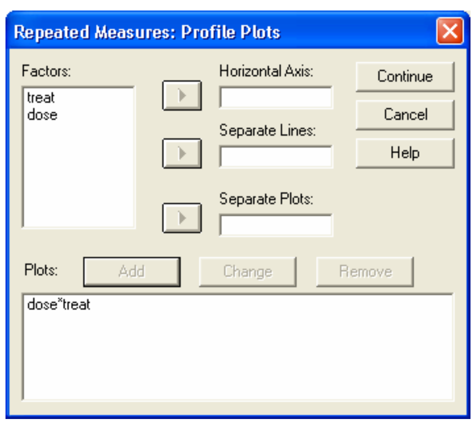

Figure 6

Other Options

As before, *post hoc* tests are disabled for solely repeated measures designs. Therefore, the only remaining options are in the *options* dialog box, which is accessed by clicking on Options... . The options here are the same as for the one-way ANOVA. I recommend selecting some descriptive statistics and you might also want to select some multiple comparisons by selecting all factors in the box labelled *Factor(s) and Factor Interactions* and transferring them to the box labelled *Display Means for* by clicking on ▶ (see Figure 7). Having selected these variables, you should tick the box labelled *Compare main effects* (☑ Compare main effects) and select an appropriate correction (I chose Bonferroni).

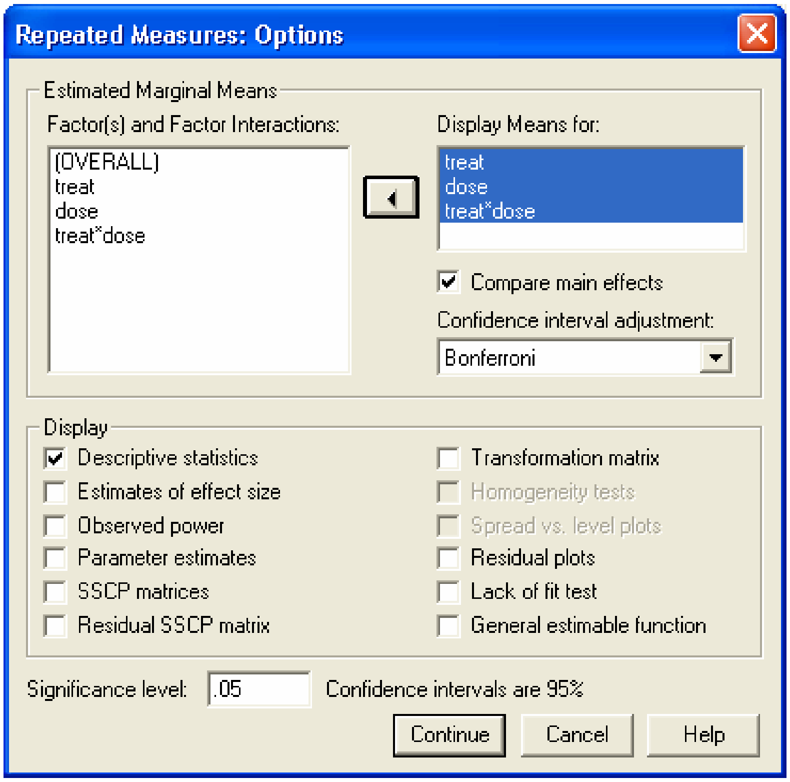

Figure 7

Interpreting the Output from Factorial Repeated Measures ANOVA

Descriptives and Main Analysis

SPSS Output 6 shows the initial output from this ANOVA. The first table merely lists the variables that have been included from the data editor and the level of each independent variable that they represent. This table is more important than it might seem, because it enables you to verify that the variables in the SPSS data editor represent the correct levels of the independent variables. The second table is a table of descriptives and provides the mean and standard deviation for each of the ten conditions. The names in this table are the names I gave the variables in the data editor (therefore, if you didn't give these variables full names, this table will look slightly different).

The descriptives are interesting in that they tell us that the variability among scores was greatest after 5 Sherries and was generally higher when sherry was consumed (compare the standard deviations of the levels of the sherry variable compared to those of the mince pie variable). The values in this table will help us later to interpret the main effects of the analysis.

Within-Subjects Factors

Measure: MEASURE_1

TREAT	DOSE	Dependent Variable
1	1	SHERRY1
	2	SHERRY2
	3	SHERRY3
	4	SHERRY4
	5	SHERRY5
2	1	PIE1
	2	PIE2
	3	PIE3
	4	PIE4
	5	PIE5

Descriptive Statistics

	Mean	Std. Deviation	N
Time Taken to Deliver Presents After 1 Sherry	10.2000	2.6998	10
Time Taken to Deliver Presents After 2 Sherries	14.5000	2.9533	10
Time Taken to Deliver Presents After 3 Sherries	23.3000	7.1032	10
Time Taken to Deliver Presents After 4 Sherries	31.8000	7.6710	10
Time Taken to Deliver Presents After 5 Sherries	48.4000	8.8719	10
Time Taken to Deliver Presents After 1 Mince Pie	9.3000	3.3015	10
Time Taken to Deliver Presents After 2 Mince Pies	13.4000	.8433	10
Time Taken to Deliver Presents After 3 Mince Pies	17.0000	4.7842	10
Time Taken to Deliver Presents After 4 Mince Pies	20.1000	4.0125	10
Time Taken to Deliver Presents After 5 Mince Pies	26.8000	5.7310	10

SPSS Output 6

SPSS Output 7 shows the results of Mauchly's sphericity test (see earlier and chapter 11 of Field, 2005) for each of the three effects in the model (two main effects and one interaction). The significance values of these tests indicate that the main effect of **dose** has violated this assumption and so the *F*-value should be corrected (see earlier and chapter 11 of Field, 2005). For the main effect of **treat** and the interaction the assumption of sphericity is met (because $p > .05$) so we need not correct the *F*-ratios for these effect.

Mauchly's Test of Sphericity[b]

Measure: MEASURE_1

Within Subjects Effect	Mauchly's W	Approx. Chi-Square	df	Sig.	Epsilon[a] Greenhouse-Geisser	Huynh-Feldt	Lower-bound
TREAT	1.000	.000	0	.	1.000	1.000	1.000
DOSE	.092	17.685	9	.043	.552	.740	.250
TREAT * DOSE	.425	6.350	9	.712	.747	1.000	.250

Tests the null hypothesis that the error covariance matrix of the orthonormalized transformed dependent variables is proportional to an identity matrix.

a. May be used to adjust the degrees of freedom for the averaged tests of significance. Corrected tests are displayed in the Tests of Within-Subjects Effects table.

b. Design: Intercept
Within Subjects Design: TREAT+DOSE+TREAT*DOSE

SPSS Output 7

SPSS Output 8 shows the results of the ANOVA (with corrected *F* values). The output is split into sections that refer to each of the effects in the model and the error terms associated with these effects (a bit like the general table earlier on in this handout). The interesting part is the significance values of the *F*-ratios. If these values are less than .05 then we can say that an effect is significant. Looking at the significance values in the table it is clear that there is a significant effect of the type of treat consumed by the elves, a significant main effect of the number of treats consumed (dose), and a significant interaction between these two variables. I will examine each of these effects in turn.

Tests of Within-Subjects Effects

Measure: MEASURE_1

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
TREAT	Sphericity Assumed	1730.560	1	1730.560	34.078	.000
	Greenhouse-Geisser	1730.560	1.000	1730.560	34.078	.000
	Huynh-Feldt	1730.560	1.000	1730.560	34.078	.000
	Lower-bound	1730.560	1.000	1730.560	34.078	.000
Error(TREAT)	Sphericity Assumed	457.040	9	50.782		
	Greenhouse-Geisser	457.040	9.000	50.782		
	Huynh-Feldt	457.040	9.000	50.782		
	Lower-bound	457.040	9.000	50.782		
DOSE	Sphericity Assumed	9517.960	4	2379.490	83.488	.000
	Greenhouse-Geisser	9517.960	2.209	4309.021	83.488	.000
	Huynh-Feldt	9517.960	2.958	3217.666	83.488	.000
	Lower-bound	9517.960	1.000	9517.960	83.488	.000
Error(DOSE)	Sphericity Assumed	1026.040	36	28.501		
	Greenhouse-Geisser	1026.040	19.880	51.613		
	Huynh-Feldt	1026.040	26.622	38.541		
	Lower-bound	1026.040	9.000	114.004		
TREAT * DOSE	Sphericity Assumed	1495.240	4	373.810	20.730	.000
	Greenhouse-Geisser	1495.240	2.989	500.205	20.730	.000
	Huynh-Feldt	1495.240	4.000	373.810	20.730	.000
	Lower-bound	1495.240	1.000	1495.240	20.730	.001
Error(TREAT*DOSE)	Sphericity Assumed	649.160	36	18.032		
	Greenhouse-Geisser	649.160	26.903	24.129		
	Huynh-Feldt	649.160	36.000	18.032		
	Lower-bound	649.160	9.000	72.129		

SPSS Output 8

The Effect of Treat

The first part of SPSS Output 8 tells us the effect of the type of treat consumed by the elves. For this effect there was no violation of sphericity and so we can look at the uncorrected *F*-ratios.

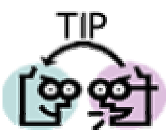

This effect should be reported as:

- ✓ There was a significant main effect of the type of treat, $F(1, 9) = 34.08$, $p < .001$

This effect tells us that if we ignore the number of treats consumed, the elves were slower at delivering presents after one type of treat than after the other type.

You can request that SPSS produce means of the main effects (see Field, 2005) and if you do this, you'll find the table in SPSS Output 9 in a section headed *Estimated Marginal Means*.[5] SPSS Output 9 is a table of means for the main effect of treat with the associated standard errors. The levels of this variable are labelled 1 and 2 and so we must think back to how we entered the variable to see which row of the table relates to which condition. We entered this variable with the sherry condition first and the mince pie condition last. Figure 8 uses this information to display the means for each condition. It is clear from this graph that mean delivery times were higher after sherry ($M = 25.64$) than after mince pies ($M = 17.32$). Therefore, sherry slowed down present delivery significantly compared to mince pies.

[5] These means are obtained by taking the average of the means in Table 2 for a given condition. For example, the mean for the mince pie condition (ignoring the dose) is

$$\bar{X}_{\text{mincepie}} = \left(\bar{X}_{1\text{mincepie}} + \bar{X}_{2\text{mincepies}} + \bar{X}_{3\text{mincepies}} + \bar{X}_{4\text{mincepies}} + \bar{X}_{5\text{mincepies}}\right)/5$$
$$= (9.3 + 13.4 + 17.0 + 20.1 + 26.8)/5 = 17.32.$$

1. TREAT

Measure: MEASURE_1

TREAT	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	25.640	1.167	22.999	28.281
2	17.320	.815	15.476	19.164

SPSS Output 9

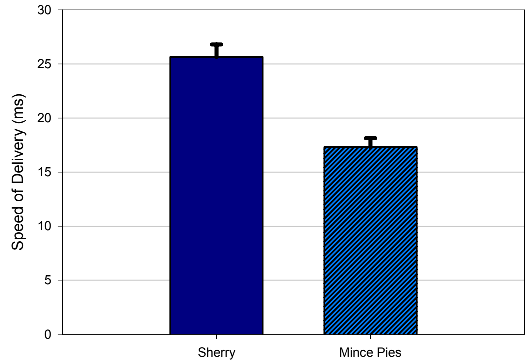

Figure 8

The Effect of Dose

SPSS Output 8 also tells us the effect of the number of treats consumed (dose) by the elves. For this effect there was a violation of sphericity and so we must look at the corrected *F*-ratios. All of the corrected values are highly significant and so we can report the Greenhouse-Geisser corrected values as these are the most conservative. You should report the sphericity data (mauchley's test etc.) as explained in the first example. The effect itself could be reported as:

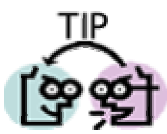

✓ There was a significant main effect of the number of treats consumed, $F(2.21, 19.88) = 83.49$, $p < .001$.

Note the degrees of freedom represent the Greenhouse-Geisser corrected values. This effect tells us that if we ignore the type of treats that was consumed, the elves were slower at delivering presents after consuming certain amounts of treats. We don't know from this effect, which amounts of treats (doses) in particular slowed the elves down, but we could look at this with post hoc tests – see example 1.

2. DOSE

Measure: MEASURE_1

DOSE	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	9.750	.704	8.157	11.343
2	13.950	.491	12.839	15.061
3	20.150	1.250	17.323	22.977
4	25.950	1.427	22.722	29.178
5	37.600	1.973	33.136	42.064

SPSS Output 10

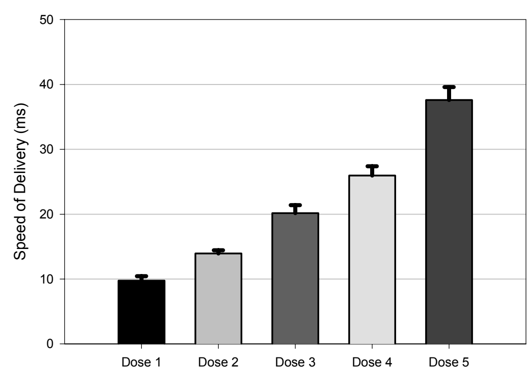

Figure 9

If we requested means of the main effects (see Field, 2005 section 11.8.4) then you'll see the table in SPSS Output 10, which is a table of means for the main effect of dose with the associated standard errors. The levels of this variable are labelled 1, 2, 3, 4, & 5 and so we must think back to how we entered the variables to see which row of the table relates to which condition. Figure 9 uses this information to display the means for each condition. It is clear from this graph that mean delivery times got progressively higher as more treats were consumed (in fact the trend looks linear).

The Interaction Effect (Treat × Dose)

SPSS Output 8 indicated that the number of treats consumed interacted in some way with the type of treat. In other words, the effect that the number of treats (dose) had on the speed of delivery was different for mince pies and sherry. The means for all conditions can be seen in SPSS Output 11 (and these values are the same as in the table of descriptives).

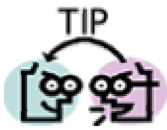

The interaction did not violate sphericity and so we can report from the ANOVA table that:

- ✓ There was a significant interaction between the type of treat consumed and the number of treats consumed, $F(4, 36) = 20.73$, $p < .001$.

This effect tells us that the effect of consuming more treats was stronger for one of the treats than for the other.

3. TREAT * DOSE

Measure: MEASURE_1

TREAT	DOSE	Mean	Std. Error	95% Confidence Interval Lower Bound	95% Confidence Interval Upper Bound
1	1	10.200	.854	8.269	12.131
	2	14.500	.934	12.387	16.613
	3	23.300	2.246	18.219	28.381
	4	31.800	2.426	26.312	37.288
	5	48.400	2.806	42.053	54.747
2	1	9.300	1.044	6.938	11.662
	2	13.400	.267	12.797	14.003
	3	17.000	1.513	13.578	20.422
	4	20.100	1.269	17.230	22.970
	5	26.800	1.812	22.700	30.900

SPSS Output 11

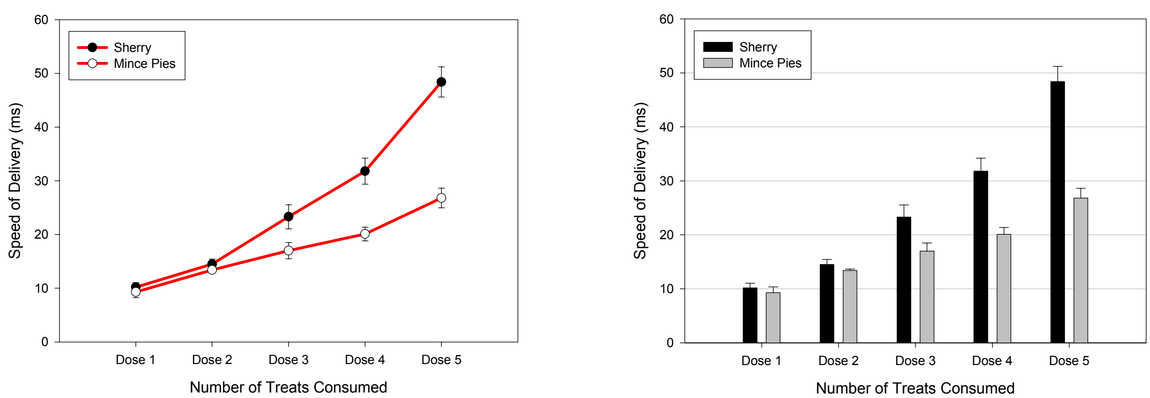

Figure 10

We can use the means in SPSS Output 11 to plot an interaction graph, which is essential for interpreting the interaction. **Figure 10** shows two interaction graphs of these data (just to illustrate that you can present them as either bars or lines). The graph shows that the pattern of responding for the two treats is very similar for small doses (the lines are almost identical for 1 and two doses and the bars are the same height). However, as more treats are consumed, the effect of drinking sherry becomes more pronounced (delivery times are higher) than when mince pies are eaten. This effect is shown by the

fact that the line representing sherry starts to deviate away from the line for mince pies (the lines become non-parallel). In the bar chart this is shown by the increasingly large differences between the pairs of bars for large numbers of treats. To verify the interpretation of the interaction effect, we would need to look at some contrasts (see Field, 2005, chapter 11). However, in general terms, Santa Claus should conclude that the number of treats consumed had a much greater effect in slowing down elves when the treat was sherry (presumably because they all get shit-faced and start staggering around being stupid), but much less of an effect when the treats were mince pies although even the pies did slow them down to some extent).

Guided Example:

A clinical psychologist was interested in the effects of antidepressants and cognitive behaviour therapy on suicidal thoughts. Four depressives took part in four conditions: placebo tablet with no therapy for one month, placebo tablet with cognitive behaviour therapy (CBT) for one month, antidepressant with no therapy for one month, and antidepressant with cognitive behaviour therapy (CBT) for one month. The order of conditions was fully counterbalanced across the 4 participants. Participants recorded the number of suicidal thoughts they had during the final week of each month. The data are below:

Drug:	**Placebo**		**Antidepressant**	
Therapy:	**None**	**CBT**	**None**	**CBT**
Andy	70	60	81	52
Leslie	66	52	70	40
Verena	56	41	60	31
Shane	68	59	77	49
Mean	**65**	**53**	**72**	**43**

The SPSS output you get for these data should look like the following:

Within-Subjects Factors

Measure: MEASURE_1

DRUG	THERAPY	Dependent Variable
1	1	PLNONE
	2	PLCBT
2	1	ANTNONE
	2	ANTCBT

Descriptive Statistics

	Mean	Std. Deviation	N
Placebo - No Therapy	65.0000	6.2183	4
Placebo - CBT	53.0000	8.7560	4
Antidepressant - No Therapy	72.0000	9.2014	4
Antidepressant - CBT	43.0000	9.4868	4

Mauchly's Test of Sphericity[b]

Measure: MEASURE_1

					Epsilon[a]		
Within Subjects Effect	Mauchly's W	Approx. Chi-Square	df	Sig.	Greenhouse-Geisser	Huynh-Feldt	Lower-bound
DRUG	1.000	.000	0	.	1.000	1.000	1.000
THERAPY	1.000	.000	0	.	1.000	1.000	1.000
DRUG * THERAPY	1.000	.000	0	.	1.000	1.000	1.000

Tests the null hypothesis that the error covariance matrix of the orthonormalized transformed dependent variables is proportional to an identity matrix.

a. May be used to adjust the degrees of freedom for the averaged tests of significance. Corrected tests are displayed in the Tests of Within-Subjects Effects table.

b.

Design: Intercept
Within Subjects Design: DRUG+THERAPY+DRUG*THERAPY

Tests of Within-Subjects Effects

Measure: MEASURE_1

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
DRUG	Sphericity Assumed	9.000	1	9.000	1.459	.314
	Greenhouse-Geisser	9.000	1.000	9.000	1.459	.314
	Huynh-Feldt	9.000	1.000	9.000	1.459	.314
	Lower-bound	9.000	1.000	9.000	1.459	.314
Error(DRUG)	Sphericity Assumed	18.500	3	6.167		
	Greenhouse-Geisser	18.500	3.000	6.167		
	Huynh-Feldt	18.500	3.000	6.167		
	Lower-bound	18.500	3.000	6.167		
THERAPY	Sphericity Assumed	1681.000	1	1681.000	530.842	.000
	Greenhouse-Geisser	1681.000	1.000	1681.000	530.842	.000
	Huynh-Feldt	1681.000	1.000	1681.000	530.842	.000
	Lower-bound	1681.000	1.000	1681.000	530.842	.000
Error(THERAPY)	Sphericity Assumed	9.500	3	3.167		
	Greenhouse-Geisser	9.500	3.000	3.167		
	Huynh-Feldt	9.500	3.000	3.167		
	Lower-bound	9.500	3.000	3.167		
DRUG * THERAPY	Sphericity Assumed	289.000	1	289.000	192.667	.001
	Greenhouse-Geisser	289.000	1.000	289.000	192.667	.001
	Huynh-Feldt	289.000	1.000	289.000	192.667	.001
	Lower-bound	289.000	1.000	289.000	192.667	.001
Error(DRUG*THERAPY)	Sphericity Assumed	4.500	3	1.500		
	Greenhouse-Geisser	4.500	3.000	1.500		
	Huynh-Feldt	4.500	3.000	1.500		
	Lower-bound	4.500	3.000	1.500		

1. DRUG

Measure: MEASURE_1

DRUG	Mean	Std. Error	95% Confidence Interval Lower Bound	Upper Bound
1	59.000	3.725	47.146	70.854
2	57.500	4.668	42.644	72.356

2. THERAPY

Measure: MEASURE_1

THERAPY	Mean	Std. Error	95% Confidence Interval Lower Bound	Upper Bound
1	68.500	3.824	56.329	80.671
2	48.000	4.546	33.532	62.468

3. DRUG * THERAPY

Measure: MEASURE_1

DRUG	THERAPY	Mean	Std. Error	95% Confidence Interval Lower Bound	Upper Bound
1	1	65.000	3.109	55.105	74.895
	2	53.000	4.378	39.067	66.933
2	1	72.000	4.601	57.358	86.642
	2	43.000	4.743	27.904	58.096

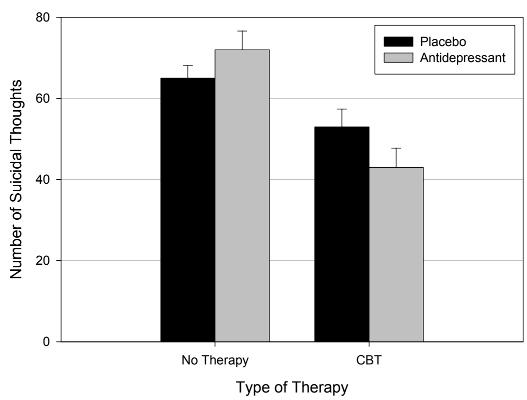

- ✓ Enter the data into SPSS.
- ✓ Save the data onto a disk in a file called **suicidaltutors.sav**.
- ✓ Conduct the appropriate analysis to see whether the number of suicidal thoughts patients had was significantly affected by the type of drug they had, the therapy they received or the interaction of the two..

What are the independent variables and how many levels do they have?

Your Answer:

What is the dependent variable?

Your Answer:

What analysis have you performed?

Your Answer:

Describe the assumption of sphericity. Has this assumption been met? (Quote relevant statistics in APA format).

Your Answer:

Report the main effect of therapy in APA format. Is this effect significant and how would you interpret it?

Your Answer:

Report the main effect of 'drug' in APA format. Is this effect significant and how would you interpret it?

Your Answer:

Report the interaction effect between drug and therapy in APA format. Is this effect significant and how would you interpret it?

Your Answer:

Unguided Example 1:

There is a lot of concern among students as to the consistency of marking between lecturers. It is pretty common that lecturers obtain reputations for being 'hard markers' or 'light markers' but there is often little to substantiate these reputations. So, a group of students investigated the consistency of marking by submitting the same essay to four different lecturers. The mark given by each lecturer was recorded for each of the 8 essays. It was important that the same essays were used for all lecturers because this eliminated any individual differences in the standard of work that each lecturer was marking. The data are below.

Essay	Tutor 1 (Dr. Field)	Tutor 2 (Dr. Smith)	Tutor 3 (Dr. Scrote)	Tutor 4 (Dr. Death)
1	62	58	63	64
2	63	60	68	65
3	65	61	72	65
4	68	64	58	61
5	69	65	54	59
6	71	67	65	50
7	78	66	67	50
8	75	73	75	45

Self-Test

- ✓ Enter the data into SPSS.
- ✓ Save the data onto a disk in a file called **tutor.sav**.
- ✓ Conduct the appropriate analysis to see whether the tutor who marked the essay had a significant effect on the mark given.
- ✓ What analysis have you performed?
- ✓ Report the results in APA format?
- ✓ Do the findings support the idea that some tutors give more generous marks than others?

The answers to this task are in Field (2005) Chapter 11.

Unguided Example 2:

In a previous handout we came across the beer-goggles effect: a severe perceptual distortion after imbibing vast quantities of alcohol. Imagine we wanted to follow this finding up to look at what factors mediate the beer goggles effect. Specifically, we thought that the beer goggles effect might be made worse by the fact that it usually occurs in clubs, which have dim lighting. We took a sample of 26 men (because the effect is stronger in men) and gave them various doses of alcohol over four different weeks (0 pints, 2 pints, 4 pints and 6 pints of lager). This is our first independent variable, which we'll call *alcohol consumption*, and it has four levels. Each week (and, therefore, in each state of drunkenness) participants were asked to select a mate in a normal club (that had dim lighting) and then select a second mate in a specially designed club that had bright lighting. As such, the second independent variable was whether the club had dim or bright lighting. The outcome measure was the attractiveness of each mate as assessed by a panel of independent judges. To recap, all participants took part in all levels of the alcohol consumption variable, and selected mates in both brightly- and dimly-lit clubs. This is the example I presented in my handout and lecture in writing up laboratory reports.

Self-Test

- ✓ Enter the data into SPSS.
- ✓ Save the data onto a disk in a file called **BeerGogglesLighting.sav**.
- ✓ Conduct the appropriate analysis to see whether the amount drunk and lighting in the club have a significant effect on mate selection.
- ✓ What analysis have you performed?
- ✓ Report the results in APA format?
- ✓ Do the findings support the idea that mate selection gets worse as lighting dims and alcohol is consumed?

Dim Lighting				Bright Lighting			
0 Pints	2 Pints	4 Pints	6 Pints	0 Pints	2 Pints	4 Pints	6 Pints
58	65	44	5	65	65	50	33
67	64	46	33	53	64	34	33
64	74	40	21	74	72	35	63
63	57	26	17	61	47	56	31
48	67	31	17	57	61	52	30
49	78	59	5	78	66	61	30
64	53	29	21	70	67	46	46
83	64	31	6	63	77	36	45
65	59	46	8	71	51	54	38
64	64	45	29	78	69	58	65
64	56	24	32	61	65	46	57
55	78	53	20	47	63	57	47
81	81	40	29	57	78	45	42
58	55	29	42	71	62	48	31
63	67	35	26	58	58	42	32
49	71	47	33	48	48	67	48
52	67	46	12	58	66	74	43
77	71	14	15	65	32	47	27
74	68	53	15	50	67	47	45
73	64	31	23	58	68	47	46
67	75	40	28	67	69	44	44
58	68	35	13	61	55	66	50
82	68	22	43	66	61	44	44
64	70	44	18	68	51	46	33
67	55	31	13	37	50	49	22
81	43	27	30	59	45	69	35

The answer to this question is in the file **Answers(Chapter11).pdf** on the CR-ROM of Field (2005).

Unguided Example 3

Imagine I wanted to look at the effect alcohol has on the 'roving eye' (because, I seem to be rather obsessed with experiments involving alcohol and dating for some bizarre reason). The 'roving eye' effect is the propensity of people in relationships to 'eye-up' members of the opposite sex. I took 20 men and fitted them with incredibly sophisticated glasses that could track their eye movements and record both the movement and the object being observed (this is the point at which it should be apparent that I'm making it up as I go along). Over 4 different nights I plied these poor souls with either 1, 2, 3 or 4 pints of strong lager in a pub. Each night I measured how many different women they eyed-up (a women was categorized as having been eyed up if the man's eye moved from her head to toe and back up again). To validate this measure we also collected the amount of dribble on the man's chin while looking at a woman.

1 Pint	2 Pints	3 Pints	4 Pints
15	13	18	13
3	5	15	18
3	6	15	13
17	16	15	14
13	10	8	7
12	10	14	16
21	16	24	15
10	8	14	19
16	20	18	18
12	15	16	13
11	4	6	13
12	10	8	23
9	12	7	6
13	14	13	13
12	11	9	12
11	10	15	17
12	19	26	19
15	18	25	21
6	6	20	21
12	11	18	8

- ✓ Enter the data into SPSS.
- ✓ Save the data onto a disk in a file called **RovingEye.sav**.
- ✓ Conduct the appropriate analysis to see whether the amount drunk has a significant effect on the roving eye.
- ✓ What analysis have you performed?
- ✓ Report the results in APA format?
- ✓ Do the findings support the idea that males tend to eye up females more after they drink alcohol?

The answer to this question is in the file **Answers(Chapter11).pdf** on the CR-ROM of Field (2005).

Unguided Example 4:

Western people can become obsessed with body weight and diets, and because the media are insistent on ramming ridiculous images of stick-thin celebrities down into our eyes and brainwashing us into believing that these emaciated corpses are actually attractive, we all end up terribly depressed that we're not perfect (because we don't have a couple of red slugs stuck to our faces instead of lips). This gives evil corporate types the opportunity to jump on our vulnerability by making loads of money on diets that will apparently help us attain the body beautiful! Well, not wishing to miss out on this great opportunity to exploit people's insecurities I came up with my own diet called the 'Andikins diet'[6]. The basic principle is that you eat like me: you eat no meat, drink lots of Darjeeling tea, eat shed-loads of smelly European cheese with lots of fresh crusty bread, pasta, and eat chocolate at every available opportunity, and enjoy a few beers at the weekend. To test the efficacy of my wonderful new diet, I took 10 people who considered themselves to be in need of losing weight (this was for ethical reasons – you can't force people to diet!) and put them on this diet for two months. Their weight was measured in Kilograms at the start of the diet and then after 1 month and 2 months.

Before Diet	After 1 Month	After 2 Months
63.75	65.38	81.34
62.98	66.24	69.31
65.98	67.70	77.89
107.27	102.72	91.33
66.58	69.45	72.87
120.46	119.96	114.26
62.01	66.09	68.01
71.87	73.62	55.43
83.01	75.81	71.63
76.62	67.66	68.60

- ✓ Enter the data into SPSS.
- ✓ Save the data onto a disk in a file called **AndikinsDiet.sav**.
- ✓ Conduct the appropriate analysis to see whether the diet is effective.
- ✓ What analysis have you performed?
- ✓ Report the results in APA format?
- ✓ Does the diet work?

The results should be similar to those I presented in the lecture☺ (Look at my lecture slides on the course website).

[6] Not to be confused with the Atkins diet obviously☺

… And Finally The Multiple Choice Test!

Go to http://www.sagepub.co.uk/field/multiplechoice.html and test yourself on the multiple choice questions for Chapter 11. If you get any wrong, re-read this handout (or Field, 2005, Chapter 11) and do them again until you get them all correct.

This handout doesn't particularly contain material from:

Field, A. P. (2005). *Discovering statistics using SPSS (2nd edition)*. London: Sage.

but some of the other examples are from it and please consult this book (chapter 11) for more detail on how to use SPSS to analyse factorial ANOVAs with repeated measures.