



Mixed Factorial ANOVA

Introduction

The final ANOVA design that we need to look at is one in which you have a mixture of between-group and repeated measures variables. It should be obvious that you need at least two independent variables for this type of design to be possible, but you can have more complex scenarios too (e.g. two between-group and one repeated measures, one between-group and two repeated measures, or even two of each). SPSS allows you to test almost any design you might want to of virtually any degree of complexity. However, interaction terms are difficult enough to interpret with only two variables so imagine how difficult they are if you include, for example, four!

Two-Way Mixed 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 'mixed' part of the name tells us that the same participants have been used to manipulate one independent variable, but different participants have been used when manipulating the other. Therefore, this analysis is appropriate when you have one repeated-measures independent variables, and one between-group independent variables.

An Example: The Real Santa




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 immovable 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 Hypothesis¹ 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 the tip of a red hat, they could hear a hearty chuckle.



Imagine their surprise as the last book was removed revealing three identical Santas ... All three were static as if hypnotised. The Statisticians cackled with joy, and Professor Life shouted "ha ha, Santa is transfixed, held in a catatonic state by the power of statistical equations. It is statistically improbable that you will identify the real Santa because you have but a 1 in 3 chance of guessing correctly. The odds are stacked against you, you pointy eared, fun loving, non-significant elves; so confident am I that you won't identify the real Santa, that if you can I will release him and renounce my life of numbers to join the elfin clan"

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



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.

DISCOVERING STATISTICS



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). Just to make doubly sure, they decided to enlist the help of Rudolph and some of his fellow reindeers who'd been enjoying a nice bucket of water outside. These Reindeers have particularly sensitive taste buds and each one licked the three Santas (which the Elves couldn't do because of Elf & Safety regulations ...) and like the elves gave each Santa a rating from 0 to 10 based on his taste.

Table 1: Data for the Santa example

Rater	Ratings of Santa 1	Ratings of Santa 2	Ratings of Santa 3
Elves	1	3	1
	2	5	3
	4	6	6
	5	7	4
	5	9	1
	6	9	3
Reindeer	1	10	2
	4	8	1
	5	7	3
	4	9	2
	2	10	4
	5	10	2

Entering the Data

The independent variables were the Santa that was being assessed (Santa 1 2 or 3) and whether the rating was made by an elf or a reindeer. The dependent variable was the elf's rating out of 10.

To enter these data into SPSS we use the same procedure as the repeated measures ANOVA that we came across last week, except that we also need a variable (column) that codes whether the helper was an elf or a reindeer.



- Levels of repeated measures variables go in different columns of the SPSS data editor.
- Data from different people go in different rows of the data editor, therefore, levels of between-group variables go in a single column (a coding variable).

Therefore, separate columns should represent each level of a repeated measures variable and a fourth column should be made with numbers representing whether the rater was an elf or a reindeer. So, create a column and call it **rater**. Use the value 1 to represent elves, and 2 to represent reindeer (and remember to change the 'values' property so that we know what these numbers represent, and to change the 'measure' property to 'nominal' so that SPSS knows that **rater** is a categorical variable.

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

The Main Analysis

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

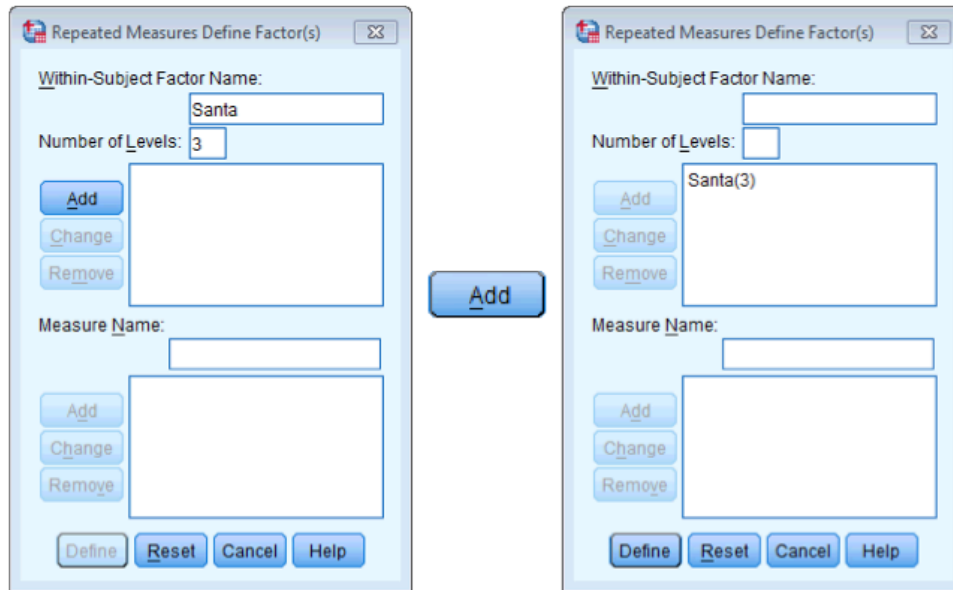




Figure 1: *Define Factors* dialog box for repeated measures ANOVA

In the *Define Factors* 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/Reindeer tested, so replace the word *factor1* with the word *Santa*. The name you give to the repeated measures variable cannot have spaces. 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/Reindeers 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 preceded 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 dragging them or 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.

So far the procedure has been similar to a one-way repeated measures design (last week). However, we have a mixed design here, and so we also need to specify our between-group factor as well. We do this by selecting **rater** in the variables list and dragging it to the box labelled *Between-Subjects Factors* (or click on ). The completed dialog box should look exactly like Figure 3. I've already discussed the options for the buttons at the bottom of this dialog box, so I'll talk only about the ones of particular interest for this example.

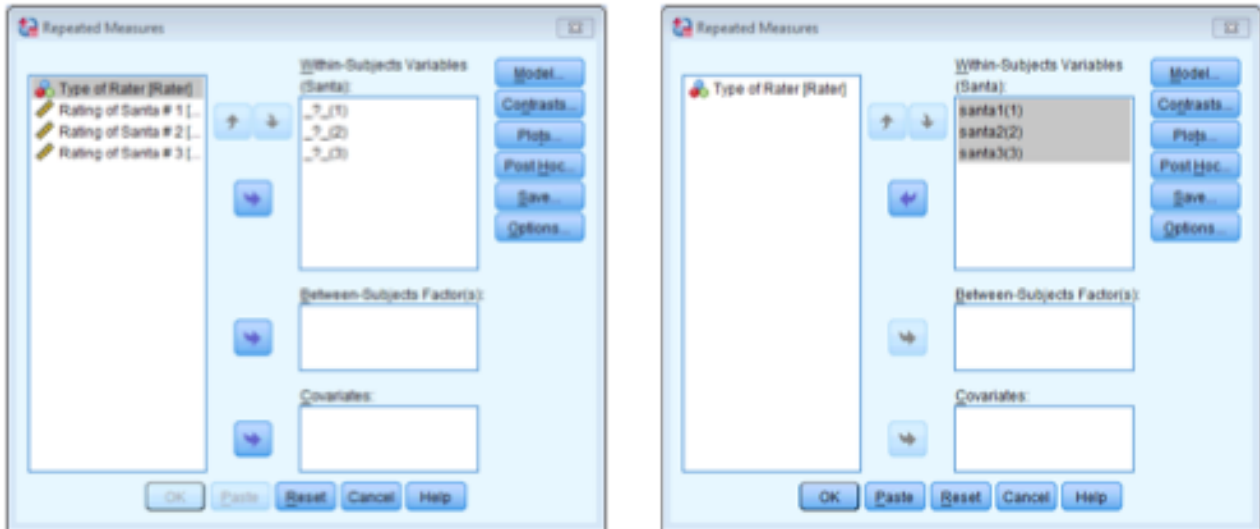


Figure 2: Main dialog box for repeated measures ANOVA

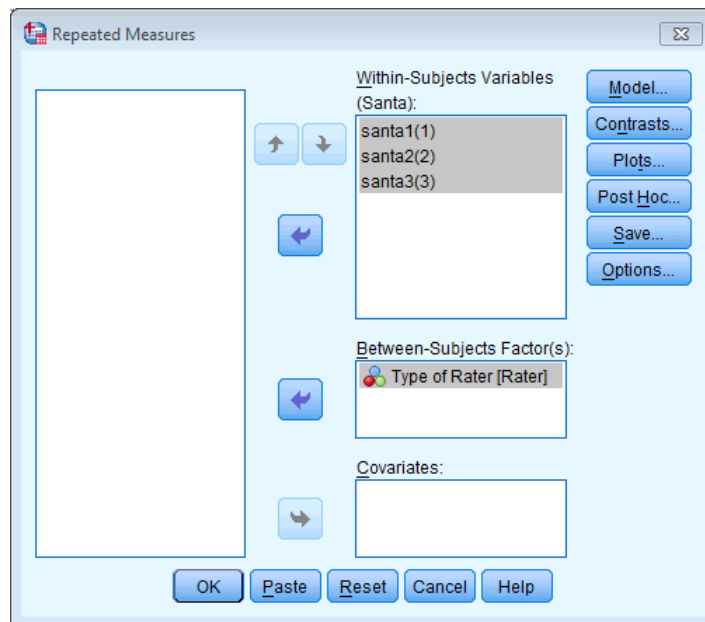


Figure 3

Post Hoc Tests

There is no proper facility for producing post hoc tests for repeated measures variables in SPSS! However, consult your handout from last week to tell you about using Bonferroni corrected t -tests.

Graphing Interactions

When 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 interaction effects. To access this dialog box click on **Plots...**. Select **santa** from the variables list on the left-hand side of the dialog box and drag it to the space labelled Horizontal Axis (or click on **Add**). In the space labelled Separate Lines we need to place the remaining independent variable: **Rater**. It is down to your discretion which way round the graph is plotted. 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 4). When you have finished specifying graphs, click on **Continue** to return to the main dialog box.

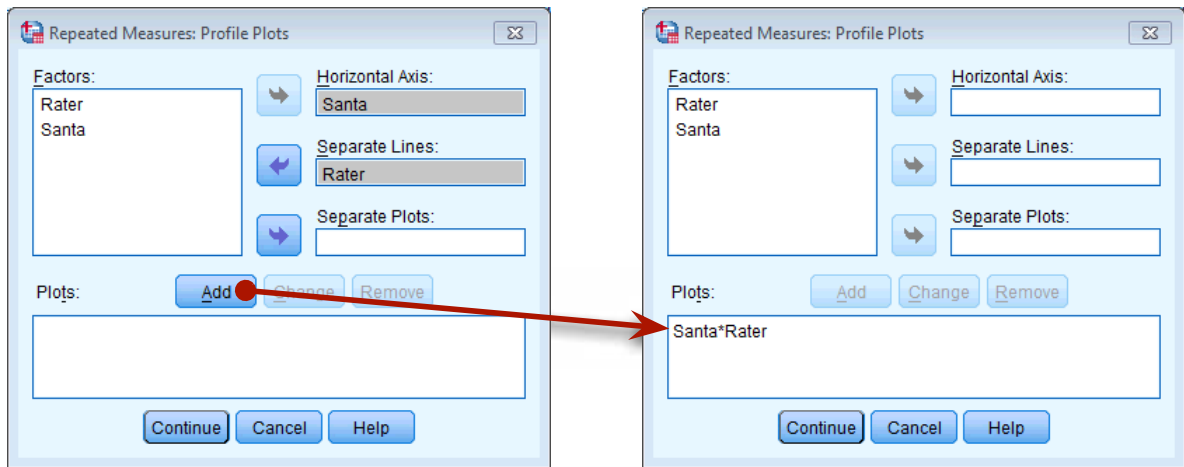


Figure 4

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 5) 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 is active because there is a between group factor and we should select this to get Levene's test (see your handout on Bias from week 1).

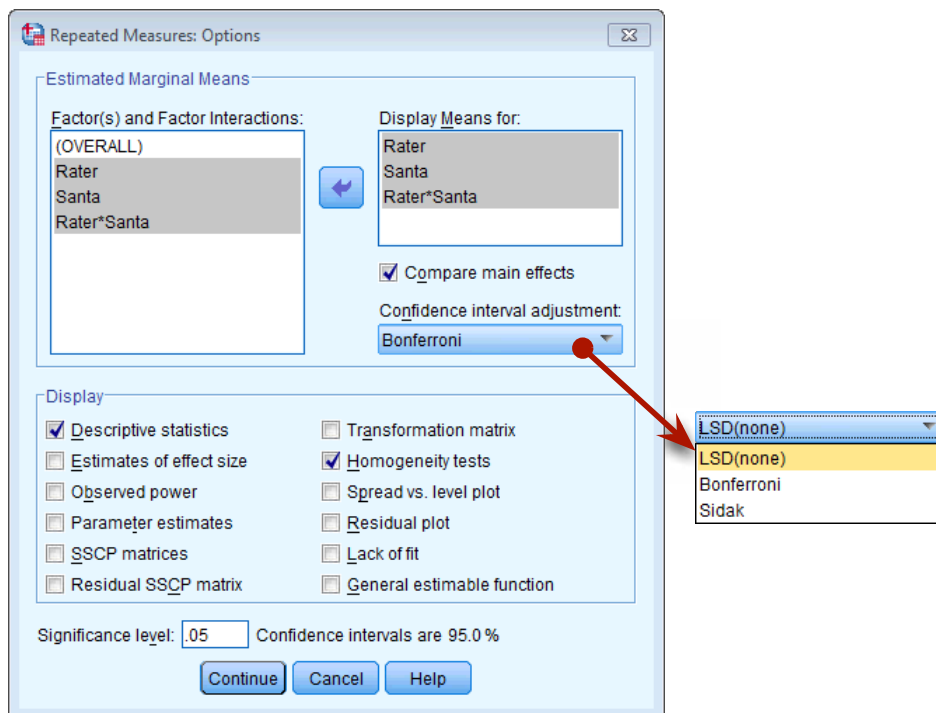

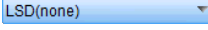


Figure 5: Options dialog box

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 drag it to the box labelled *Display Means for* (or click on ). Once a variable has been transferred, the box labelled *Compare main effects* (*Compare main effects*) becomes active and you should select this option. 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  to return to the main dialog box, and then click on  to run the analysis.

Descriptive statistics and other Diagnostics

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, Santa 2 was rated highest (i.e. most likely to be the real Santa) by both elves and reindeer, note that the reindeers gave particularly high ratings of Santa 2.

Within-Subjects Factors

Measure: MEASURE_1

Santa	Dependent Variable
1	santa1
2	santa2
3	santa3

Descriptive Statistics


	Type of Rater	Mean	Std. Deviation	N
Rating of Santa # 1	Elf	3.83	1.941	6
	Reindeer	3.50	1.643	6
	Total	3.67	1.723	12
Rating of Santa # 2	Elf	6.50	2.345	6
	Reindeer	9.00	1.265	6
	Total	7.75	2.221	12
Rating of Santa # 3	Elf	3.00	1.897	6
	Reindeer	2.33	1.033	6
	Total	2.67	1.497	12

SPSS Output 1

Assessing Sphericity

Last week 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.

TIP



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

→ If it is significant we must use Greenhouse-Geisser or Huyn-Feldt corrected degrees of freedom to assess the significance of the corresponding *F*.

SPSS Output 2 shows Mauchly's test for the Santa data, and the important column is the one containing the significance value. The significance value is .788, which is more than .05, so we can reject the hypothesis that the variances of the differences between levels were significantly different. In other words the assumption of sphericity has been met.

SPSS Output 3 shows the results of the 'repeated measures' part 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 both of the effects are significant. The effect of the between group variable (**Rater**) is found in a different table, which we'll look at next because I will examine each of the effects that we need to analyse in turn.



Mauchly's Test of Sphericity

Measure: MEASURE_1

Within Subjects Effect	Mauchly's W	Approx. Chi-Square	df	Sig.	Epsilon ^a		
					Greenhouse-Geisser	Huynh-Feldt	Lower-bound
Santa	.948	.477	2	.788	.951	1.000	.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+Rater
Within Subjects Design: Santa

SPSS Output 2

Tests of Within-Subjects Effects

Measure: MEASURE_1

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
Santa	Sphericity Assumed	174.056	2	87.028	36.946	.000
	Greenhouse-Geisser	174.056	1.902	91.519	36.946	.000
	Huynh-Feldt	174.056	2.000	87.028	36.946	.000
	Lower-bound	174.056	1.000	174.056	36.946	.000
Santa * Rater	Sphericity Assumed	18.167	2	9.083	3.856	.038
	Greenhouse-Geisser	18.167	1.902	9.552	3.856	.041
	Huynh-Feldt	18.167	2.000	9.083	3.856	.038
	Lower-bound	18.167	1.000	18.167	3.856	.078
Error(Santa)	Sphericity Assumed	47.111	20	2.356		
	Greenhouse-Geisser	47.111	19.019	2.477		
	Huynh-Feldt	47.111	20.000	2.356		
	Lower-bound	47.111	10.000	4.711		

SPSS Output 3

The Effect of Rater



The main effect of **Rater** is listed separately from the repeated measure effects in a table labelled *Tests of Between-Subjects Effects*. Before looking at this table it is important to check the assumption of homogeneity of variance using Levene's test (see handout on exploring data). SPSS produces a table listing Levene's test for each of the repeated measures variables in the data editor, and we need to look for any variable that has a significant value. SPSS Output 4 shows both tables. The table showing Levene's test indicates that variances are homogeneous (i.e. more or less the same for elves and Reindeers) for all levels of the repeated measures variables (because all significance values are greater than .05). Had any of the values been significant then it would have compromised the accuracy of the *F*-test for Rater and you could consider transforming all of the data to stabilize the variances between groups (see Field, 2009, Chapter 5 or your handout). Luckily for us, this wasn't the case for these data. The second table shows the ANOVA summary table for the main effect of Rater, and this reveals a non-significant effect (because the significance of .491 is more than the standard cut-off point of .05).

If you requested that SPSS display means for the Rater effect you should scan through your output and find the table in a section headed *Estimated Marginal Means*. SPSS Output 5 is a table of means for the main effect of Rater with the associated standard errors. This information is plotted in Figure 6. It is clear from this graph that elves and reindeer's ratings were generally the same (although remember that this does not take into account which Santa was being rated, it just takes the average rating across all Santas).



- The F ratio represents the ratio of the experimental effect compared to 'error'. Therefore, when it is less than 1 it means there was more error than variance created by the experiment. As such, an $F < 1$ is always non-significant and sometimes people report these statistics as simply $F < 1$. However, it is more informative to report the exact value.
- We can report that 'there was a non-significant main effect of rater, $F(1, 10) = 0.51, p = .491$.
- This effect tells us that if we ignore the santa being rated, elves and reindeers gave similar ratings'.

Levene's Test of Equality of Error Variances

	F	df1	df2	Sig.
Rating of Santa # 1	.207	1	10	.659
Rating of Santa # 2	2.232	1	10	.166
Rating of Santa # 3	1.025	1	10	.335

Tests the null hypothesis that the error variance of the dependent variable is equal across groups.

a.

Design: Intercept+Rater
Within Subjects Design: Santa

Tests of Between-Subjects Effects

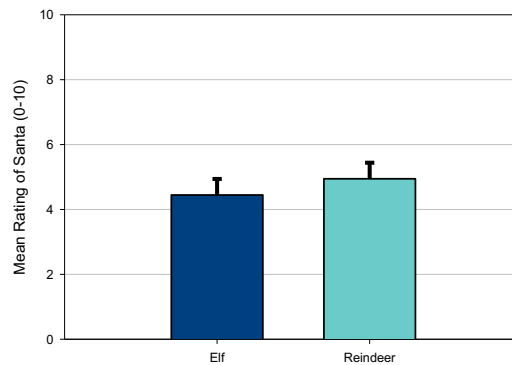
Measure: MEASURE_1
Transformed Variable: Average

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Intercept	793.361	1	793.361	180.082	.000
Rater	2.250	1	2.250	.511	.491
Error	44.056	10	4.406		

SPSS Output 4

Estimates

Type of Rater	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
Elf	4.444	.495	3.342	5.547
Reindeer	4.944	.495	3.842	6.047



SPSS Output 5

Figure 6

The Effect of Santa



The first part of SPSS Output 3 tells us the effect of the Santa that was evaluated by the elves/reindeers (i.e. did the three Santa's receive different average ratings?). For this effect sphericity wasn't an issue and so we can look at the uncorrected F -ratio, which was significant.

You can request that SPSS produce means of the main effects (see Field, 2009) and if you do this, you'll find the table in SPSS Output 6 in a section headed *Estimated Marginal Means*. SPSS Output 6 is a table of means for the main effect of Santa with the associated standard errors. The levels of this variable are labelled 1, 2 and 3 and 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 in the order of Santas rated (Santa #1, Santa #2 and Santa #3). Figure 7 uses this information to display the means for each Santa. It is clear from this graph that the mean Rating was highest for Santa #2, $M = 7.75$ (i.e. the elves and Reindeers were most confident that this Santa was the real Santa). The other Santa's were rated fairly equally, Santa #1 $M = 3.67$ and Santa #3 $M = 2.67$. Therefore, elves and Reindeers in combination were most confident that Santa #2 was the real Santa. If




you asked for Post Hoc tests (SPSS Output 7) you could go on to say that ratings of Santa #1 and Santa #3 did not significantly differ, but ratings of Santa #2 were significantly higher than both Santa #1 and Santa #3.

Estimates

Measure: MEASURE_1

Santa	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	3.667	.519	2.510	4.823
2	7.750	.544	6.538	8.962
3	2.667	.441	1.684	3.649



SPSS Output 6

Figure 7

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	-4.083*	.554	.000	-5.673	-2.493
	3	1.000	.643	.453	-.846	2.846
2	1	4.083*	.554	.000	2.493	5.673
	3	5.083*	.676	.000	3.143	7.023
3	1	-1.000	.643	.453	-2.846	.846
	2	-5.083*	.676	.000	-7.023	-3.143

Based on estimated marginal means

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

a. Adjustment for multiple comparisons: Bonferroni.

SPSS Output 7




This effect should be reported as:

- There was a significant main effect of the Santa being rated, $F(2, 20) = 36.95, p < .001$
- This effect tells us that if we ignore whether the rating came from an elf or reindeer, the ratings of the three Santas significantly differed.
- Bonferroni corrected post hoc tests showed that ratings of Santa #1 and Santa #3 did not significantly differ ($p = .45$), but ratings of Santa #2 were significantly higher than both Santa #1 and Santa #3 (both $ps < .001$).

The Santa × Rater Interaction Effect

SPSS Output 3 indicated that the organism doing the rating (elf or reindeer) interacted with the Santa that was being rated. In other words, the ratings for the three Santas differed in some way in elves and reindeers. The means for all conditions can be seen in SPSS Output 8.



This effect should be reported as:

- There was a significant Santer × Rater interaction, $F(2, 20) = 3.86, p = .038$. This effect tells us that the ratings of the three Santas significantly differed in elves and reindeers.




We can use the means in SPSS Output 8 to plot an interaction graph, which is essential for interpreting the interaction. Figure 8 shows that the ratings of elves and reindeers were fairly similar for Santa #2 and Santa #3 (the bars are similar heights). However, for Santa #2 reindeer ratings seem to be quite a lot higher than elves. So, although both elves and reindeers seemed most confident that Santa #2 was the real Santa, reindeers appeared to be more confident about this than elves (their ratings were higher). To verify the interpretation of the interaction effect, we would need to look at some contrasts (see Field, 2009, chapter 13).

3. Type of Rater * Santa

Measure: MEASURE_1

Type of Rater	Santa	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Elf	1	3.833	.734	2.198	5.469
	2	6.500	.769	4.786	8.214
	3	3.000	.624	1.611	4.389
Reindeer	1	3.500	.734	1.864	5.136
	2	9.000	.769	7.286	10.714
	3	2.333	.624	.944	3.723



SPSS Output 8

Figure 8

A Happy Ending?

Well, you're probably wondering what happened? Both the elves and the reindeers were most confident that Santa #2 was the real Santa and that's the Santa they chose. On hearing their choice, Professor N. O. Life let out a blood-curdling cry, "Damn you, you pesky elves and reindeers" he wailed, "you have correctly identified the real Santa". True to his word he released the real Santa, and Santas #1 and 3 removed their very realistic disguises to reveal two more statisticians. Santa emerged from his spell, "Thank you, my friends" he said "my head was filled with numbers, I couldn't move, I feared for my mental 'elf". "That's Ok" said Twallybliddle, "The kids need you to deliver the presents, and we can't let them down". "But how on earth did you identify me" said Santa? "Well", said Twallybliddle "my elves tickled you and we have big pointy super-sensitive ears that allowed us to distinguish your laugh from the impostors"; at which point Rudolph butted in and said "and my reindeer friends and I have big tongues, and we tasted you, and as everyone knows the Real Santa tastes of Christmas pudding²", he added "the other two Tasted of B.O., so we knew they must be mathematicians in disguise". Santa roared a big belly laugh "Ho Ho Ho" he said, "come on, there's work to be done". "A bit less work" added Twallybliddle, "because we have an extra elf helper", and as he passed a pair of little green shorts and a pointy green hat to Professor N. O. Life everyone began to laugh, even Professor N. O. Life who had already began calculating an equation to minimise Santa's flight path on Christmas night. They all got back to Lapland in time, and all the Children got their presents and the world was a happy place. Hooray!



Three-Way Mixed ANOVA

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 'three-way' part of the name simply means that three independent variables have been manipulated in the experiment. The 'mixed' part of the name tells us that the same participants have been used to manipulate one or more independent variables, but different participants have been used when manipulating one or more independent variables. Therefore,

² You can put it to the test in a couple of weeks should you want to ...



this analysis is appropriate when you have one or more repeated-measures independent variables, and one or more between-group independent variables.

An Example: Are Fairies or Elves More Susceptible to Christmas Treats?



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 Elves³. He also wanted to see whether a different type of helper might be less susceptible to these treats. So, Santa did a little experiment. He randomly selected 9 Elves from his workforce and also took on 9 new helpers from the Fairy kingdom and timed how long it took each of them to deliver the presents to 5 houses. About half of the elves/fairies 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 9 elves and the same 9 fairies 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 9 elves and 9 fairies 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. (Fairy image is from www.drawingbusiness.com/)



- Why do you think Santa got half of the elves/fairies 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/fairies and it has 2 levels: Sherry or Mince Pies.
- **Dose:** Independent Variable 2 is the dose of the treat (remember each elf/fairy 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.
- **Helper:** Independent variable 3 was whether the helper was an elf or a fairy.
- 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 $2 \times 5 \times 2$ three-way mixed ANOVA. As with other ANOVA designs, there is no limit to the number of conditions for each of the independent variables in the experiment; however, in practice, 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/Christmas Pudding				
Elf	Dose:	1	2	3	4	5	1	2	3	4	5
	Trampy	11	11	30	15	42	8	15	17	28	18
	Hughy	15	16	19	28	43	13	8	27	21	31
	Lardy	15	16	15	22	45	11	15	31	12	26
	Alchi	12	9	31	33	25	11	16	20	38	41
	Goody	7	14	23	29	39	14	7	9	13	42
	Pongo	8	14	21	34	36	10	12	17	24	29


³ 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).

Fairy	Bringitup	16	14	15	21	50	16	8	15	14	33
	Dio	12	15	24	22	35	12	16	12	30	43
	Chunder	6	15	32	34	39	13	11	15	14	3
	Duncan	12	13	35	35	53	7	12	20	7	32
	Ozzy	6	16	35	40	59	13	16	20	28	28
	Skinny	6	18	29	52	72	7	15	25	32	27
	Speedy	8	10	23	50	65	6	18	18	19	37
	Floaty	10	12	25	30	82	7	13	8	20	31
	Smiley	17	16	22	19	36	10	13	20	16	25
	Retch	6	8	12	50	46	12	9	24	18	40
	Wibble	9	10	35	38	36	12	14	23	23	33
	Condrick	9	15	22	37	72	12	14	22	14	19

Table 2: Data for example two

Entering the Data

To enter these data into SPSS we use the same procedure as the previous example, remembering the golden rules of the data editor:



TIP

- Levels of repeated measures variables go in different columns of the SPSS data editor.
- Data from different people go in different rows of the data editor, therefore, levels of between-group variables go in a single column (a coding variable).

If a person participates in all experimental conditions (in this case all elves and fairies 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 (variable view) 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

In addition, we need a variable (column) that codes whether the helper was an elf or a fairy. So, create a column and call it **Helper**. Use the value 1 to represent elves, and 2 to represent fairies (and remember to change the 'values' property so that we know what these numbers represent, and to change the 'measure' property to 'nominal' so that SPSS knows that Helper is a categorical variable).

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 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 9**. When you have entered both of the within-subject factors click on **Define** to go to the main dialog box.

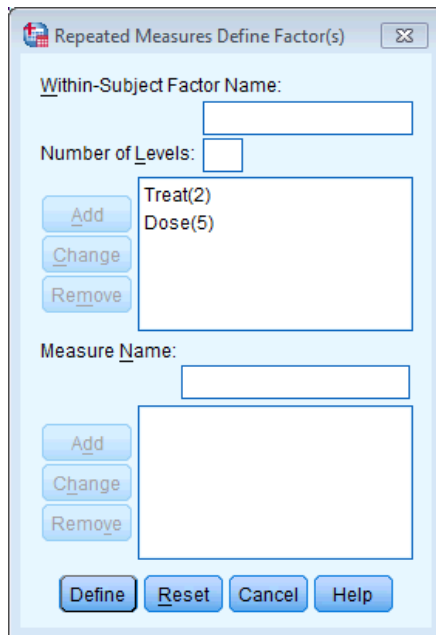


Figure 9: Define factors dialog box for factorial repeated measures ANOVA

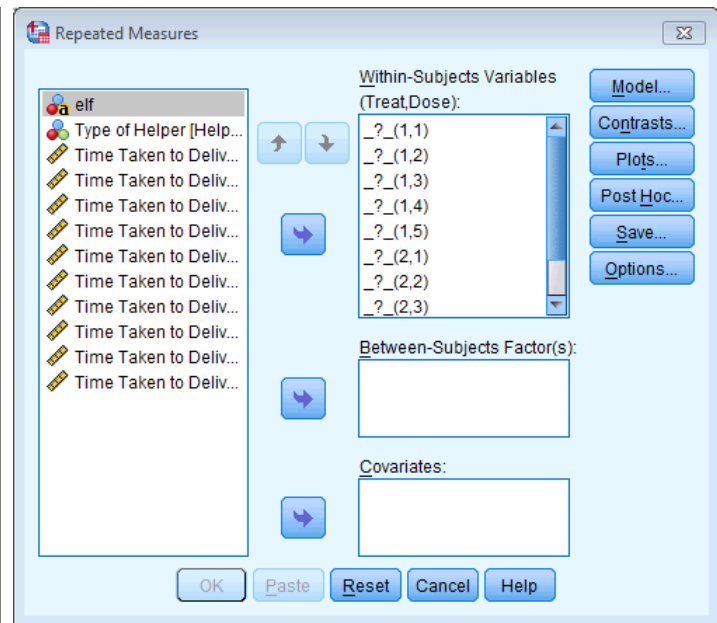


Figure 10

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 10). 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**

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 11. 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.

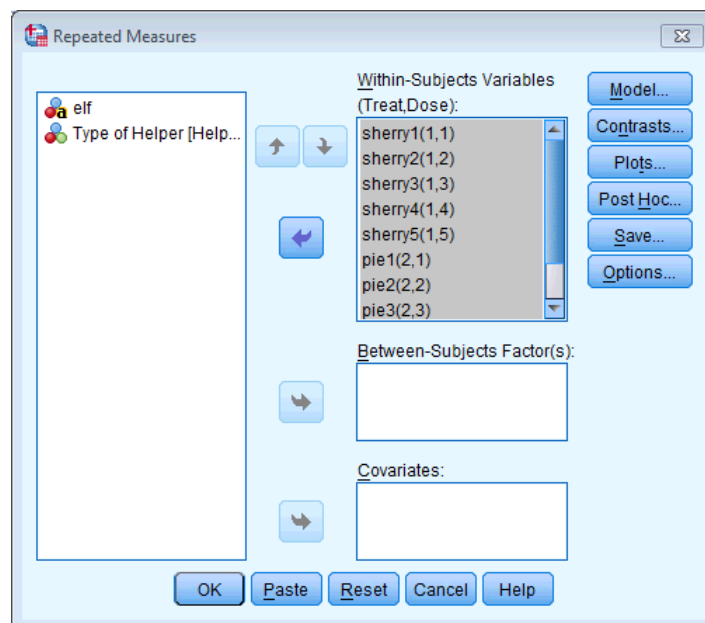



Figure 11

So far the procedure has been similar to other factorial repeated measures designs. However, we have a mixed design here, and so we also need to specify our between-group factor as well. We do this by selecting **helper** in the variables list and dragging it (or clicking on ) to transfer it to the box labelled *Between-Subjects Factors*. The completed dialog box should look exactly like Figure 12. I've already discussed the options for the buttons at the bottom of this dialog box, so I'll talk only about the ones of particular interest for this example.

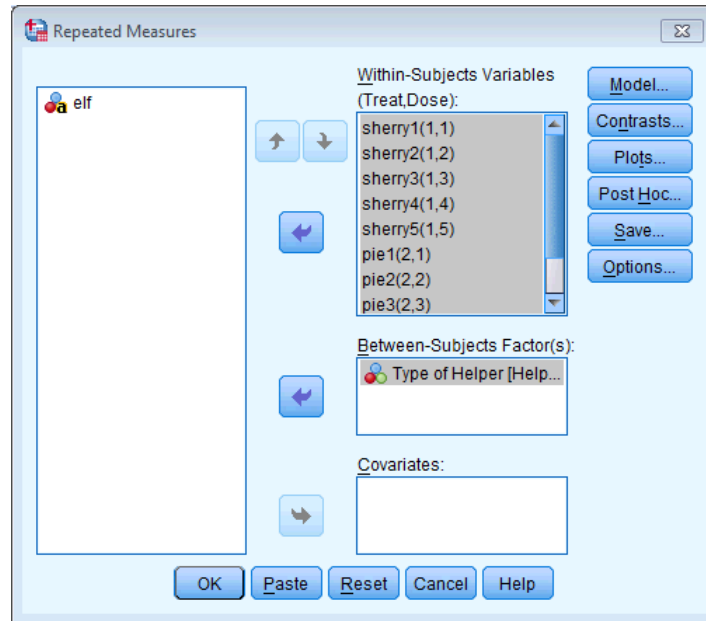


Figure 12

Graphing Interactions



The addition of an extra variable makes it necessary to choose a different graph to the one in the previous example. Click on **Plots...** to access the dialog box in Figure 13. Place **dose** in the slot labelled *Horizontal Axis:* and **treat** in the slot labelled *Separate Line:*, finally, place **helper** in the slot labelled *Separate Plots*. When all three variables have been specified, don't forget to click on **Add** to add this combination to the list of plots. By asking SPSS to plot the dose \times treat \times helper interaction, we should get the interaction graph for dose and treat, but a separate version of this graph will be produced for elves and fairies. You could also think about plotting graphs for the two way interactions (e.g. dose \times treat, dose \times helper, and treat \times helper). 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 because this variable is ordered (2 doses are bigger than 1 and so on), and by placing this variable on the horizontal axis it will enable us to easily track the effect on delivery times as the dose increases. Also, because dose has lots of levels (5), if we asked for this variable to be plotted as separate lines then we'd have a lot of lines (very confusing), or worse still, if we asked for this variable to be plotted on different graphs we'd end up with 5 different graphs (the road to madness). However, ultimately it's up to you, what I have suggested is simply what I find easiest. When you have finished specifying graphs, click on **Continue** to return to the main dialog box.

Other Options

As far as other options are concerned, you should select the same ones that were chosen for the previous example. It is worth selecting estimated marginal means for all effects (because these values will help you to understand any significant effects). When all of the appropriate options have been selected, run the analysis.

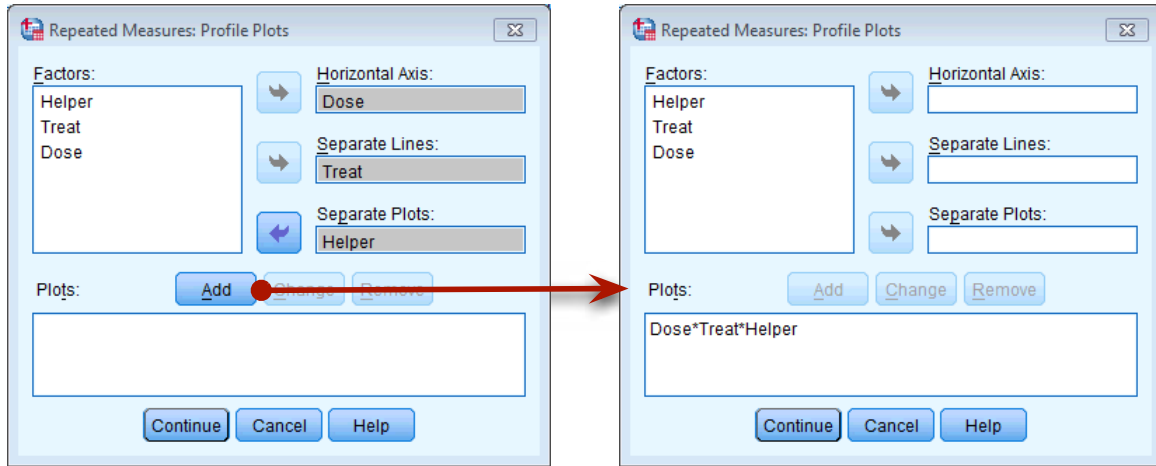


Figure 13

Interpreting the Output from Three-Way Mixed ANOVA



Descriptives and Main Analysis SPSS Output 9 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).

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

	Type of Helper	Mean	Std. Deviation	N
Time Taken to Deliver Presents After 1 Sherry	Elf	11.33	3.674	9
	Fairy	9.22	3.563	9
	Total	10.28	3.675	18
Time Taken to Deliver Presents After 2 Sherries	Elf	13.78	2.333	9
	Fairy	13.11	3.371	9
	Total	13.44	2.833	18
Time Taken to Deliver Presents After 3 Sherries	Elf	23.33	6.538	9
	Fairy	26.44	7.812	9
	Total	24.89	7.169	18
Time Taken to Deliver Presents After 4 Sherries	Elf	26.44	6.766	9
	Fairy	39.00	10.689	9
	Total	32.72	10.818	18
Time Taken to Deliver Presents After 5 Sherries	Elf	39.33	7.089	9
	Fairy	57.89	16.412	9
	Total	48.61	15.542	18
Time Taken to Deliver Presents After 1 Mince Pie	Elf	12.00	2.345	9
	Fairy	9.56	2.789	9
	Total	10.78	2.798	18
Time Taken to Deliver Presents After 2 Mince Pies	Elf	12.00	3.674	9
	Fairy	13.78	2.539	9
	Total	12.89	3.197	18
Time Taken to Deliver Presents After 3 Mince Pies	Elf	18.11	6.990	9
	Fairy	20.00	5.025	9
	Total	19.06	5.985	18
Time Taken to Deliver Presents After 4 Mince Pies	Elf	21.56	9.139	9
	Fairy	19.67	7.433	9
	Total	20.61	8.140	18
Time Taken to Deliver Presents After 5 Mince Pies	Elf	29.56	12.905	9
	Fairy	30.22	6.340	9
	Total	29.89	9.869	18

SPSS Output 9

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 standard deviations also look bigger for fairies compared to elves. The values in this table will help us later to interpret the main effects of the analysis.

SPSS Output 10 shows the results of Mauchly's sphericity test for each of the three repeated measures effects in the model (two main effects and one interaction). The significance values of these tests indicate that the main effect of **dose** and the **dose × treat** interaction has violated this assumption and so the *F*-values for any effect involving dose or the dose × treat interaction term should be corrected (see earlier and chapter 13 of Field, 2009).



- Why is the Sig. column of Mauchley's test sometimes empty?**
- The assumption of sphericity is all about the variance of the differences of different conditions being equal. When you have only two levels of a variable (as we have with treat), then there is only one set of differences (sherry compared to mince pies), so there isn't anything to compare the variance of these differences against.
 - Therefore, when you have only two levels of a repeated measures variable, Sphericity simply isn't an issue.
 - That's why SPSS leaves the Sig column blank, because sphericity can't be tested.
 - The main effect of **treat** has two levels so the assumption of sphericity is not an issue and we need not correct its *F*-ratio.

Mauchly's Test of Sphericity^b

Measure: MEASURE_1

Within Subjects Effect	Mauchly's W	Approx. Chi-Square	df	Sig.	Epsilon ^a		
					Greenhouse e-Geisser	Huynh-Feldt	Lower-bound
Treat	1.000	.000	0	.	1.000	1.000	1.000
Dose	.099	33.321	9	.000	.583	.732	.250
Treat * Dose	.198	23.332	9	.006	.648	.833	.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+Helper
Within Subjects Design: Treat+Dose+Treat*Dose

SPSS Output 10

SPSS Output 11 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 all of the effects are significant. I will examine each of these effects in turn.

The Effect of Helper

The main effect of **Helper** is listed separately from the repeated measure effects in a table labelled *Tests of Between-Subjects Effects*. Before looking at this table it is important to check the assumption of homogeneity of variance using Levene's test (see handout on exploring data). SPSS produces a table listing Levene's test for each of the repeated measures variables in the data editor, and we need to look for any variable that has a significant value. SPSS Output 12 shows both tables. The table showing Levene's test indicates that variances are homogeneous (i.e. more or less the same for elves and fairies) for all levels of the repeated measures variables (because all significance values are greater than .05) except for 5 doses of sherry. This significant value compromises the accuracy of the *F*-test for Helper; however, because it's only one of the 10 combinations of levels that is significant we can probably not lose sleep over this result. You could consider transforming all of the data though to stabilize the variances between groups (see Field, 2009,

Chapter 5 or your handout). The second table shows the ANOVA summary table for the main effect of Helper, and this reveals a significant effect (because the significance of .006 is less than the standard cut-off point of .05).



- We can report that ‘there was a significant main effect of Helper, $F(1, 16) = 9.87, p = .006$ ’.
- This effect tells us that if we ignore the type of treat that was given and how many of those treats were consumed, elves and fairies differed in their delivery speeds’.

Tests of Within-Subjects Effects

Measure: MEASURE_1

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
Treat	Sphericity Assumed	2427.339	1	2427.339	42.352	.000
	Greenhouse-Geisser	2427.339	1.000	2427.339	42.352	.000
	Huynh-Feldt	2427.339	1.000	2427.339	42.352	.000
	Lower-bound	2427.339	1.000	2427.339	42.352	.000
Treat * Helper	Sphericity Assumed	444.939	1	444.939	7.763	.013
	Greenhouse-Geisser	444.939	1.000	444.939	7.763	.013
	Huynh-Feldt	444.939	1.000	444.939	7.763	.013
	Lower-bound	444.939	1.000	444.939	7.763	.013
Error(Treat)	Sphericity Assumed	917.022	16	57.314		
	Greenhouse-Geisser	917.022	16.000	57.314		
	Huynh-Feldt	917.022	16.000	57.314		
	Lower-bound	917.022	16.000	57.314		
Dose	Sphericity Assumed	19025.256	4	4756.314	92.188	.000
	Greenhouse-Geisser	19025.256	2.334	8152.053	92.188	.000
	Huynh-Feldt	19025.256	2.928	6498.711	92.188	.000
	Lower-bound	19025.256	1.000	19025.256	92.188	.000
Dose * Helper	Sphericity Assumed	748.144	4	187.036	3.625	.010
	Greenhouse-Geisser	748.144	2.334	320.569	3.625	.030
	Huynh-Feldt	748.144	2.928	255.554	3.625	.020
	Lower-bound	748.144	1.000	748.144	3.625	.075
Error(Dose)	Sphericity Assumed	3302.000	64	51.594		
	Greenhouse-Geisser	3302.000	37.341	88.429		
	Huynh-Feldt	3302.000	46.841	70.494		
	Lower-bound	3302.000	16.000	206.375		
Treat * Dose	Sphericity Assumed	2358.744	4	589.686	10.217	.000
	Greenhouse-Geisser	2358.744	2.594	909.419	10.217	.000
	Huynh-Feldt	2358.744	3.333	707.646	10.217	.000
	Lower-bound	2358.744	1.000	2358.744	10.217	.006
Treat * Dose * Helper	Sphericity Assumed	761.589	4	190.397	3.299	.016
	Greenhouse-Geisser	761.589	2.594	293.632	3.299	.035
	Huynh-Feldt	761.589	3.333	228.484	3.299	.023
	Lower-bound	761.589	1.000	761.589	3.299	.088
Error(Treat*Dose)	Sphericity Assumed	3693.867	64	57.717		
	Greenhouse-Geisser	3693.867	41.499	89.011		
	Huynh-Feldt	3693.867	53.332	69.262		
	Lower-bound	3693.867	16.000	230.867		

SPSS Output 11



Levene's Test of Equality of Error Variances^a

	F	df1	df2	Sig.
Time Taken to Deliver Presents After 1 Sherry	.194	1	16	.666
Time Taken to Deliver Presents After 2 Sherries	2.279	1	16	.151
Time Taken to Deliver Presents After 3 Sherries	.325	1	16	.577
Time Taken to Deliver Presents After 4 Sherries	.909	1	16	.355
Time Taken to Deliver Presents After 5 Sherries	7.305	1	16	.016
Time Taken to Deliver Presents After 1 Mince Pie	1.693	1	16	.212
Time Taken to Deliver Presents After 2 Mince Pies	2.854	1	16	.111
Time Taken to Deliver Presents After 3 Mince Pies	1.294	1	16	.272
Time Taken to Deliver Presents After 4 Mince Pies	.933	1	16	.348
Time Taken to Deliver Presents After 5 Mince Pies	2.267	1	16	.152

Tests the null hypothesis that the error variance of the dependent variable is equal across groups.

a.

Design: Intercept+Helper

Within Subjects Design: Treat+Dose+Treat*Dose

Tests of Between-Subjects Effects

Measure: MEASURE_1

Transformed Variable: Average


Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Intercept	8964.605	1	8964.605	1989.065	.000
Helper	44.494	1	44.494	9.872	.006
Error	72.111	16	4.507		

SPSS Output 12

Estimates

Measure: MEASURE_1

Type of Helper	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
Elf	20.744	.708	19.244	22.245
Fairy	23.889	.708	22.389	25.389



SPSS Output 13

Figure 14

If you requested that SPSS display means for the Helper effect you should scan through your output and find the table in a section headed *Estimated Marginal Means*. SPSS Output 13 is a table of means for the main effect of Helper with the associated standard errors. This information is plotted in Figure 14. It is clear from this graph that delivery times were different, with elves delivering quicker than fairies.

The Effect of Treat


The first part of SPSS Output 11 tells us the effect of the type of treat consumed by the elves/fairies. For this effect sphericity wasn't an issue, so we look at the uncorrected *F*-ratios. You can request that SPSS produce means of the main effects (see Field, 2009; 2013) and if you do this, you'll find the table in SPSS Output 14 in a section headed *Estimated Marginal Means*. SPSS Output 14 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 15 uses this information to display the means for each condition. It is clear from this graph that mean delivery times were higher after sherry ($M = 25.99$) than after mince pies ($M = 18.64$). Therefore, sherry slowed down present delivery significantly compared to mince pies.

TIP



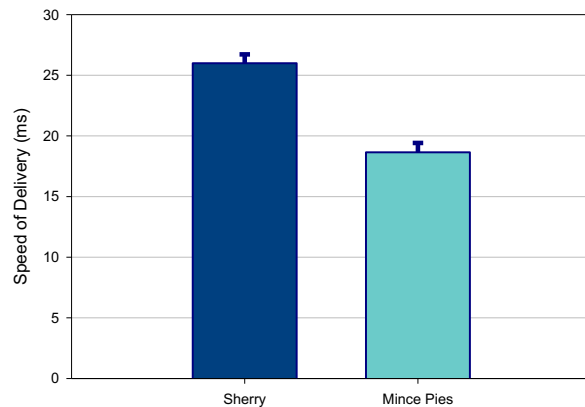
This effect should be reported as:

- There was a significant main effect of the type of treat, $F(1, 16) = 42.35, p < .001$
- This effect tells us that if we ignore the number of treats consumed and whether the helper was an elf or fairy, delivery times for presents were slower after one type of treat than after the other type.

Estimates

Measure: MEASURE_1

Treat	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	25.989	.735	24.431	27.547
2	18.644	.773	17.006	20.283




SPSS Output 14

Figure 15

The Effect of Dose

SPSS Output 11 also reports the effect of the number of treats consumed (dose) by the elves and fairies. This effect violated the assumption of sphericity and so we look at the corrected F -ratios. All of the corrected values are significant and we should report the Greenhouse-Geisser corrected values because they are the most conservative. You should report the sphericity data (Mauchly's test etc.) as explained in the first example. The effect itself could be reported as:

TIP



This effect violated sphericity so we report that (note the df):

- There was a significant main effect of the number of treats consumed, $F(2.33, 37.34) = 92.19, p < .001$.
- This effect tells us that if we ignore the type of treat that was consumed, and whether the helper was an elf or a fairy delivery times of presents were slower after consuming certain amounts of treats.

Note the degrees of freedom represent the Greenhouse-Geisser corrected values. 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.

If we requested means of the main effects (see Field, 2009, Chapter 13) then you'll see the table in SPSS Output 15, 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 16 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).



Estimates

Measure: MEASURE_1

Dose	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	10.528	.522	9.422	11.634
2	13.167	.495	12.117	14.216
3	21.972	.994	19.865	24.080
4	26.667	1.606	23.262	30.072
5	39.250	1.703	35.640	42.860

SPSS Output 15




Figure 16

The Treat × Helper Interaction

SPSS Output 11 indicated that the type of helper interacted in some way with the type of treat consumed.

TIP

→ We can report that ‘there was a significant interaction between the type of treat consumed and whether the helper was an elf or a fairy, $F(1, 16) = 7.76, p = .013$ ’.

→ This effect tells us that the effect of different treats on the delivery speed of presents was different in elves compared to fairies.

We can use the estimated marginal means to interpret this interaction (or we could have asked SPSS for a plot of Treat × Helper using the dialog box in Figure 13). The means and interaction graph (Figure 17 and SPSS Output 16) show the meaning of this result. The graph shows the average delivery times for elves and fairies for the two types of treat (regardless of how many were eaten). The graph clearly shows that for mince pies delivery times were virtually identical for elves and fairies (the bars are the same height). However, for after sherry, fairies took much longer than elves to deliver presents. In general this interaction seems to suggest that although elves and fairies were equally affected by mince pies, fairies were affected more by sherry than elves were.



4. Type of Helper * Treat

Measure: MEASURE_1

Type of Helper	Treat	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Elf	1	22.844	1.039	20.641	25.048
	2	18.644	1.093	16.327	20.962
Fairy	1	29.133	1.039	26.930	31.337
	2	18.644	1.093	16.327	20.962

SPSS Output 16

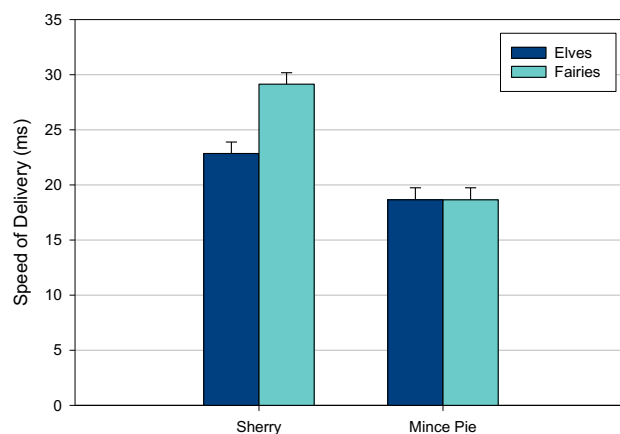



Figure 17



The Dose × Helper Interaction

SPSS Output 11 indicated that the type of helper interacted in some way with the number of treats consumed.

TIP



The interaction violated sphericity and so we report from the ANOVA table that:

- We can report that ‘there was a significant interaction between the Dose of the treat and whether the helper was an elf or a fairy, $F(2.33, 37.34) = 3.63, p = .030$.
- This effect tells us that the delivery times of elves and fairies was affected differently by the number of treats that they consumed.

We can use the estimated marginal means to interpret this interaction (or we could have asked SPSS for a plot of Dose × Helper using the dialog box in Figure 13). The means and interaction graph (Figure 18 and SPSS Output 17) show the meaning of this result. The graph shows the average delivery times for elves and fairies at each of the doses (ignoring which treat was consumed). The graph shows that delivery times are very similar for elves and fairies at doses 1 to 3, but at dose 4 and 5 fairies start to become quite a lot slower than elves. In general this interaction seems to suggest that although elves and fairies were equally affected at lower doses of treat, fairies were affected more than elves by high doses (4 or 5 doses).

5. Type of Helper * Dose

Measure: MEASURE_1

Type of Helper	Dose	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Elf	1	11.667	.738	10.103	13.231
	2	12.889	.700	11.404	14.373
	3	20.722	1.406	17.742	23.703
	4	24.000	2.272	19.185	28.815
	5	34.444	2.408	29.340	39.549
Fairy	1	9.389	.738	7.825	10.953
	2	13.444	.700	11.960	14.929
	3	23.222	1.406	20.242	26.203
	4	29.333	2.272	24.518	34.149
	5	44.056	2.408	38.951	49.160

SPSS Output 17

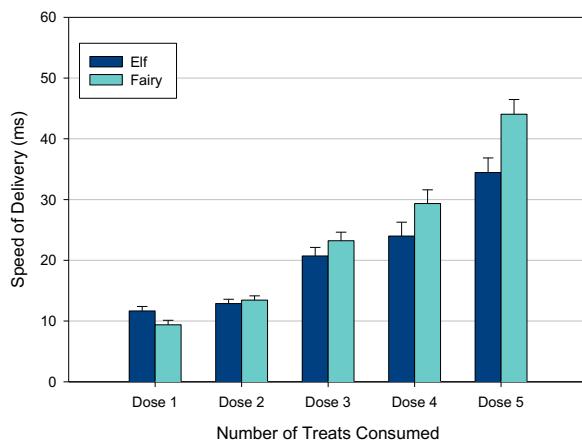



Figure 18

The Treat × Dose Interaction Effect

SPSS Output 11 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 18 (and these values are the same as in the table of descriptives).

TIP



The interaction violated sphericity and so we report that:

- There was a significant interaction between the type of treat consumed and the number of treats consumed, $F(2.59, 41.50) = 10.22, p < .001$.
- This effect tells us that the effect of consuming more treats was stronger for one of the treats than for the other.

We can use the means in SPSS Output 18 to plot an interaction graph, which is essential for interpreting the interaction. Figure 19 shows that the pattern of responding for the two treats is very similar for small doses (the bars are almost identical heights for 1 and 2 doses). 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 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,





2009, chapter 13). However, in general terms, Santa Claus should conclude that the number of treats consumed had a much greater effect in slowing down his helpers 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). However, at this stage we don't know whether this effect is found in all helpers or whether it differs in elves and fairies.

6. Treat * Dose

Measure: MEASURE_1

Treat	Dose	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	1	10.278	.853	8.469	12.086
	2	13.444	.683	11.996	14.893
	3	24.889	1.698	21.290	28.488
	4	32.722	2.108	28.253	37.192
	5	48.611	2.980	42.295	54.928
2	1	10.778	.607	9.490	12.065
	2	12.889	.744	11.311	14.467
	3	19.056	1.435	16.014	22.097
	4	20.611	1.963	16.449	24.773
	5	29.889	2.396	24.809	34.969

SPSS Output 18

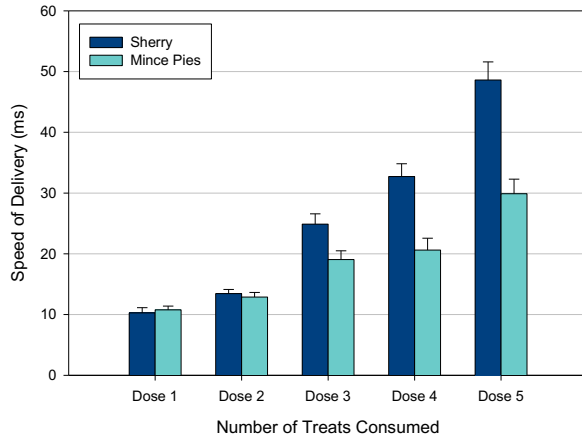


Figure 19

The Treat × Dose × Helper Interaction

SPSS Output 11 tells us that there is a significant three-way Treat × Dose × Helper interaction:

TIP

→ We can report that 'there was a significant interaction between the type of treat eaten, how many treats were eaten and whether the helper was an elf or a fairy, $F(2.59, 41.50) = 3.30, p = .035$ '.

→ The three-way interaction tells us whether the dose × treat interaction described above is the same for elves and fairies (i.e. whether the combined effect of the type of treat and their number of treats consumed was the same for elves as for fairies).

The nature of this interaction is shown up in Figure 20, which shows the helper by dose interaction for the two treats separately (sherry on the left and mince pies on the right). The means in this graph were taken from SPSS Output 19. First look at the Mince pies graph (right hand side). Hopefully it's clear that elves and fairies delivery times are relatively similar at all doses (the bars are more or less the same height). In other words the effect that dose has on delivery times is pretty similar in elves and fairies at all doses. Now, let's look at the Sherry graph (left hand side). The picture is very different here: for 1, 2 or even 3 doses of sherry the elves and fairies are affected to a similar degree (the bars are the same height approximately). However, at the fourth dose of sherry fairies slow down considerably compared to elves, and this effect is even more pronounced at 5 doses. As such, this interaction seems to suggest that elves and fairies are fairly comparable at all doses of mince pies and small doses of sherry (up to 3 doses), but at large doses of sherry, fairies slow down delivering presents substantially more than elves do.





7. Type of Helper * Treat * Dose

Measure: MEASURE_1

Type of Helper	Treat	Dose	Mean	Std. Error	95% Confidence Interval	
					Lower Bound	Upper Bound
Elf	1	1	11.333	1.206	8.776	13.891
		2	13.778	.966	11.729	15.826
		3	23.333	2.401	18.243	28.424
		4	26.444	2.982	20.124	32.765
		5	39.333	4.214	30.400	48.266
	2	1	12.000	.859	10.179	13.821
		2	12.000	1.053	9.769	14.231
		3	18.111	2.029	13.810	22.413
		4	21.556	2.777	15.669	27.442
		5	29.556	3.389	22.371	36.740
Fairy	1	1	9.222	1.206	6.665	11.780
		2	13.111	.966	11.063	15.159
		3	26.444	2.401	21.354	31.535
		4	39.000	2.982	32.679	45.321
		5	57.889	4.214	48.956	66.822
	2	1	9.556	.859	7.735	11.376
		2	13.778	1.053	11.546	16.009
		3	20.000	2.029	15.698	24.302
		4	19.667	2.777	13.780	25.553
		5	30.222	3.389	23.038	37.406

SPSS Output 19

Sherry

Mince Pies/Christmas Pudding


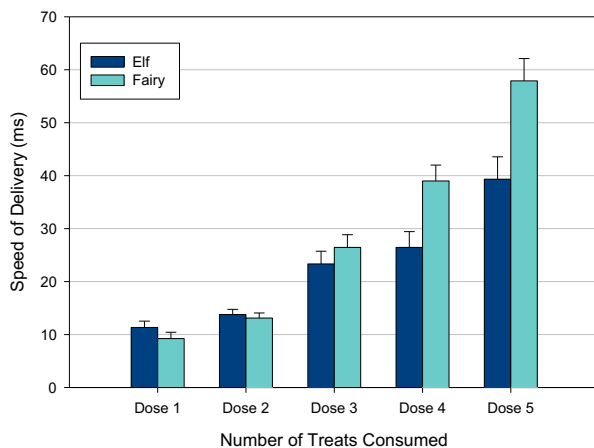


Figure 20

Conclusions

What should be clear from this handout is that when more than two independent variables are used in an ANOVA, it yields complex interaction effects that require a great deal of concentration to interpret (imagine interpreting a four-way interaction!). Therefore, it is essential to take a systematic approach to interpretation and plotting graphs is a particularly useful way to proceed.

Guided Example

There is evidence that attitudes towards stimuli can be changed using positive and negative imagery (e.g. Stuart, Shimp and Engle, 1987, but see Field and Davey, 1999). Some researchers were interested in answering two questions. On the

DISCOVERING STATISTICS



one hand, the government had funded them to look at whether negative imagery in advertising could be used to change attitudes towards alcohol. Conversely, an alcohol company had provided funding to see whether positive imagery could be used to improve attitudes towards alcohol. The scientists designed *two* studies to address both issues.

In the first study, participants viewed a total of 3 mock adverts over three sessions. In one session, they saw three adverts containing three different products with a negative image (a dead body with the slogan ‘drinking this product makes your liver explode’: The products were: (1) a brand of beer (Brain Death)); (2) a brand of wine (Dangleberry); and (3) a brand of water (Puritan). The sex of the participants was noted. Table 1 contains the data (each row represents a single subject). After each advert subjects were asked to rate the drinks on a scale ranging from –100 (dislike very much) through 0 (neutral) to 100 (like very much). The order of adverts was randomised. There are two independent variables in each experiment: the type of drink (beer, wine or water) and the sex of the participant (male or female).

Table 2: Data for Experiment 1

Experiment	Experiment 1: Negative Imagery		
Drink	Beer	Wine	Water
Male	6	–5	–14
	30	–12	–10
	15	–15	–16
	30	–4	–10
	12	–2	5
	17	–6	–6
	21	–2	–20
	23	–7	–12
	20	–10	–9
	27	–15	–6
Female	–19	–13	–2
	–18	–16	–17
	–8	–23	–19
	–6	–22	–11
	–6	–9	–10
	–9	–18	–17
	–17	–17	–4
	–12	–15	–4
	–11	–14	–1
	–6	–15	–1



- Enter the data into SPSS.
- Save the data onto a disk in a file called **drinkimagery.sav**.
- Conduct the appropriate analysis to see whether male’s and female’s attitudes to different drinks are differentially affected by negative imagery.

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? (i.e. and x by y what type of ANOVA).

Your Answer:

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

Your Answer:

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

Your Answer:

Report the main effect of sex in APA format. Is this effect significant and how would you interpret it? (You should also comment on the assumption of homogeneity of variance)



Your Answer:

Report the interaction effect between sex and type of drink in APA format. Is this effect significant and how would you interpret it?

Your Answer:

Answers to this guided question can be found on the module website in the file [answertomixedanovaguidedexample.pdf](#)

Unguided Example 1:

In a second experiment (a week later), the participants saw the same three brands, but this time presented with positive images (a sexy naked man or women—depending on the participant’s sex—and the slogan ‘drinking this product makes you a horny stud-muffin’). After each advert participants were asked to rate the drinks on the same scale.



- Analyse the data to see if imagery affects preferences for the different drinks, and whether these effects are different in men and women.
- Report the results in APA format.
- Is the prediction that the effect of positive imagery on preferences for drinks will differ in men and women supported?



Experiment	Experiment 2: Positive Imagery		
Drink	Beer	Wine	Water
Male	1	38	10
	43	20	9
	15	20	6
	10	28	20
	8	11	27
	17	17	9
	30	15	19
	34	27	12
	34	24	12
	26	23	21
Female	1	28	33
	7	26	23
	22	34	21
	30	32	17
	40	24	15
	15	29	13
	20	30	16
	9	24	17
	14	34	19
	15	23	29

Unguided Example 2

The data from the two previous examples could be combined and analysed in a three-way mixed ANOVA. Imagine that as well as positive and negative imagery, neutral imagery had been used. Participants viewed nine mock adverts over three sessions. In these adverts there were three products (a brand of beer, Brain Death, a brand of wine, Dangleberry, and a brand of water, Puritan). These could either be presented alongside positive, negative or neutral imagery. Over the three sessions, and nine adverts, each type of product was paired with each type of imagery. After each advert participants rated the drinks on the same scale as the previous two examples. The design, thus far, has two independent variables: the type of drink (beer, wine or water) and the type of imagery used (positive, negative or neutral). These two variables completely cross over, producing nine experimental conditions. Now imagine that I also took note of each person's sex. The data can be found in the file **MixedAttitude.sav**.



- Analyse the data to see if different types of imagery have different effects on preferences for the different drinks, and whether these effects are different in men and women.
- What analysis have you done?
- Report the results in APA format.
- Does the effect of imagery on drink preferences differ in men and women?

Answers are on the companion website of my book.

Unguided Example 3

Last week we saw a Speed dating example done on females (see last week's handout). Imagine this same experiment was also done on male participants (rating speed dates they had with women). The example from last week would now become a Mixed ANOVA (with sex as an additional variable). These data can be found in **LooksOrPersonality.sav**. Analyse these data with a three-way mixed ANOVA. Answers can be found in chapter 14 of my book (this is the data set I use in the book to explain mixed ANOVA).

Unguided Example 4

Text messaging is very popular but there are concerns that children will use this form of communication so much that they will not learn correct written English. One researcher conducted an experiment in which one group of children was encouraged to send text messages on their mobile phones over a six-month period. A second group was forbidden from sending text messages for the same period. To ensure that kids in this later group didn't use their phones, this group were given armbands that administered painful shocks in the presence of microwaves (like those emitted from phones). There were 50 different participants: 25 were encouraged to send text messages, and 25 were forbidden. The outcome was a score on a grammatical test (as a percentage) that was measured both before and after the experiment. The first independent variable was, therefore, text message use (text messengers versus controls) and the second independent variable was the time at which grammatical ability was assessed (before or after the experiment). The data are in the file **TextMessages.sav**.



- Analyse the data to see if text messaging affects grammatical ability.
- Report the results in APA format.
- Does text messaging reduce grammatical ability?

Answers are on the companion website of my book and some more detailed comments about can be found in Field & Hole (2003).

Multiple Choice Questions



Go to <https://studysites.uk.sagepub.com/field4e/study/mcqs.htm> and test yourself on the multiple choice questions for Chapter 15. If you get any wrong, re-read this handout (or Field, 2013, Chapter 15) and do them again until you get them all correct.

Terms of Use

This handout contains material from:

Field, A. P. (2013). *Discovering statistics using SPSS: and sex and drugs and rock 'n' roll (4th Edition)*. London: Sage.

This material is copyright Andy Field (2000-2016).

This document is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/), basically you can use it for teaching and non-profit activities but not meddle with it without permission from the author.