MANOVA using SPSS

MANOVA

The Main Analysis

Load the data in the file **OCD.sav**. Note that there are three columns: one column is a coding variable for the **group** variable (I used the codes CBT = 1, BT = 2, NT = 3), and the remaining two columns contain the scores for each dependent variable respectively. Once the data have been entered, access the main MANOVA dialog box by using the **Analyze⇒General Linear Model⇒Multivariate...** menu path.

The ANOVAs (and various multiple comparisons) carried out after the main MANOVA are identical to running separate ANOVA procedures in SPSS for each of the dependent variables. Hence, the main dialog box and options for MANOVA are very similar to the factorial ANOVA procedure. The main difference to the main dialog box is that the space labelled *Dependent Variables* has room for several variables. Select the two dependent variables from the variables list (that is **actions** and **thoughts**) and transfer them to the *Dependent Variables* box by clicking on . Select **group** from the variables list and transfer it to the *Fixed Factor(s)* box by clicking on . There is also a box in which you can place covariates. For this analysis there are no covariates; however you can apply the principles of ANCOVA to the multivariate case and conduct multivariate analysis of covariance (MANCOVA). Once you have specified the variables in the analysis, you can select any of the other dialog boxes by clicking the buttons on the right-hand side.

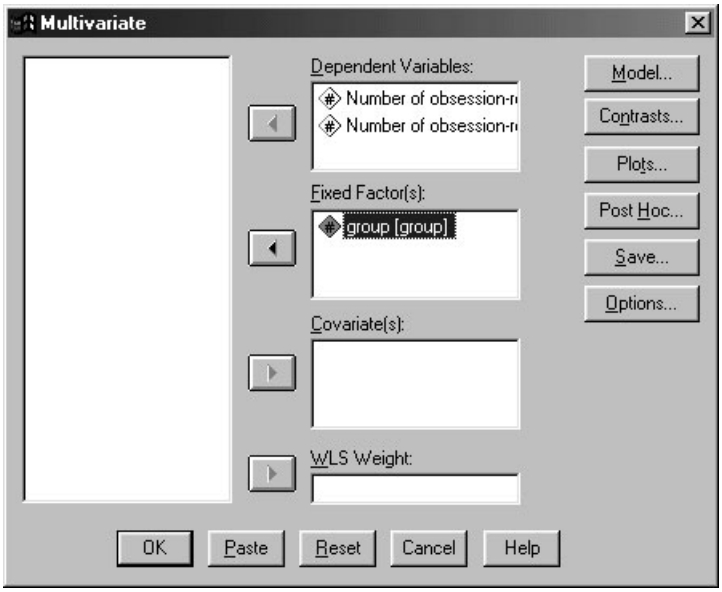

Figure 1: Main dialog box for MANOVA

Multiple Comparisons in MANOVA

The default way to follow up a MANOVA is to look at individual univariate ANOVAs for each dependent variable. For these tests, SPSS has the same options as in the univariate ANOVA procedure (see Field, 2000, Chapter 7). The Contrasts... button opens a dialog box for specifying one of several standard contrasts for the independent variable(s) in the analysis. For this example it makes sense to use a *simple* contrast that compares each of the experimental groups to the no-treatment control group. The no-treatment control group was coded as the last category (it had the highest code in the data editor), so we need to select the group

variable and change the contrast to a simple contrast using the last category as the reference category (Figure 2). For more details about contrasts see Field (2000, section 7.1.4). Instead of running a contrast, we could carry out *post hoc* tests on the independent variable to compare each group to all other groups. To access the *post hoc* tests dialog box click on Post Hoc... . The dialog box is the same as that for factorial ANOVA (see Field, 2000). For the purposes of this example, I suggest selecting two of my usual recommendations: REGWQ and Games-Howell. Once you have selected *post hoc* tests return to the main dialog box.

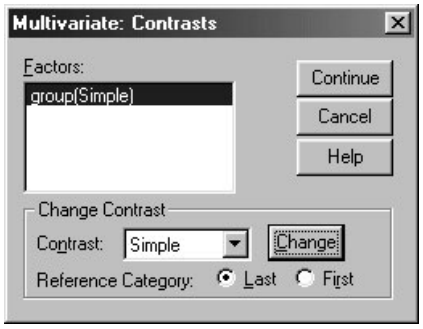

Figure 2: Contrasts for independent variable(s) in MANOVA

Ouput from MANOVA

Preliminary Analysis and Testing Assumptions

SPSS Output 1 shows an initial table of descriptive statistics that is produced by clicking on the descriptive statistics option in the *options* dialog box. This table contains the overall and group means and standard deviations for each dependent variable in turn and it is clear from the means that subjects had many more obsession-related thoughts than behaviours.

Descriptive Statistics

	group	Mean	Std. Deviation	N
Number of obsession-related thoughts	CBT	13.40	1.90	10
	BT	15.20	2.10	10
	No Treatment Control	15.00	2.36	10
	Total	14.53	2.21	30
Number of obsession-related behaviours	CBT	4.90	1.20	10
	BT	3.70	1.77	10
	No Treatment Control	5.00	1.05	10
	Total	4.53	1.46	30

SPSS Output 1

SPSS Output 2 shows Box's test of the assumption of equality of covariance matrices. This statistic tests the null hypothesis that the variance-covariance matrices are the same in all groups. Therefore, if the matrices are equal (and therefore the assumption of homogeneity is met) this statistic should be *non-significant*. For these data $p = 0.18$ (which is greater than 0.05): hence, the covariance matrices are roughly equal and the assumption is tenable. If the value of Box's test was significant ($p < 0.05$) then the covariance matrices are significantly different and so the homogeneity assumption would have been violated.

The effect of violating this assumption is unclear. Hakstian et al. (1979) report that Hotelling's T^2 is robust in the two-group situation when sample sizes are equal. As a general rule of thumb, if sample sizes are equal then disregard Box's test, because it is highly unstable, and assume Hotelling's and Pillai's statistics to be robust (see Field, 2000). However, if group sizes are different, then robustness cannot be assumed (especially if Box's test is significant at $p < 0.001$). The more dependent variables you have measured, and the greater the differences in sample sizes, the more distorted the probability values produced by SPSS become. Tabachnick and Fidell (1996), therefore, suggest that if the larger samples produce greater variances and

covariances then the probability values will be conservative (and so significant findings can be trusted). However, if it is the smaller samples that produce the larger variances and covariances then the probability values will be liberal and so significant differences should be treated with caution (although non-significant effects can be trusted). As such, Box's test need only really be examined when sample sizes differ: it should not be trusted when multivariate normality cannot be assumed (or is in question), and the variance-covariance matrices for samples should be inspected to assess whether the printed probabilities are likely to be conservative or liberal. In the event that you cannot trust the printed probabilities, there is little you can do except equalize the samples by randomly deleting cases in the larger groups (although with this loss of information comes a loss of power).

Bartlett's test of sphericity tests whether the assumption of sphericity has been met and is useful only in univariate repeated measures designs because MANOVA does not require this assumption.

Box's Test of Equality of Covariance Matrices[a]

Box's M	9.959
F	1.482
df1	6
df2	18169
Sig.	.180

a. Design: Intercept+GROUP

Bartlett's Test of Sphericity[a]

Likelihood Ratio	.042
Approx. Chi-Square	5.511
df	2
Sig.	.064

a. Design: Intercept+GROUP

SPSS Output 2

MANOVA Test Statistics

SPSS Output 3 shows the main table of results. For our purposes, the group effects are of interest because they tell us whether or not the therapies had an effect on the OCD clients. You'll see that SPSS lists the four multivariate test statistics. In the next column these values are transformed into an *F*-ratio with two degrees of freedom. The column of real interest, however, is the one containing the significance values of these *F*-ratios. For these data, Pillai's trace ($p = 0.049$), Wilks's lambda ($p = 0.050$) and Roy's largest root ($p = 0.020$) all reach the criterion for significance of 0.05. However, Hotelling's trace ($p = 0.051$) is non-significant by this criterion. This scenario is interesting, because the test statistic we choose determines whether or not we reject the null hypothesis that there are no between-group differences. However, given what we know about the robustness of Pillai's trace when sample sizes are equal, we might be well advised to trust the result of that test statistic, which indicates a significant difference. Interestingly, this example highlights the additional power associated with Roy's root (you should note how this statistic is considerably more significant than all others) when the test assumptions have been met.

From this result we should probably conclude that the type of therapy employed had a significant effect on OCD. The nature of this effect is not clear from the multivariate test statistic. First, it tells us nothing about which groups differed from which, and second it tells us nothing about whether the effect of therapy was on the obsession-related thoughts, the obsession-related behaviours, or a combination of both. To determine the nature of the effect, SPSS provides us with univariate tests.

Multivariate Tests[a]

Effect		Value	F	Hypothesis df	Error df	Sig.
Intercept	Pillai's Trace	.983	745.230[c]	2.000	26.000	.000
	Wilks' Lambda	.017	745.230[c]	2.000	26.000	.000
	Hotelling's Trace	57.325	745.230[c]	2.000	26.000	.000
	Roy's Largest Root	57.325	745.230[c]	2.000	26.000	.000
GROUP	Pillai's Trace	.318	2.557	4.000	54.000	.049
	Wilks' Lambda	.699	2.555[c]	4.000	52.000	.050
	Hotelling's Trace	.407	2.546	4.000	50.000	.051
	Roy's Largest Root	.335	4.520	2.000	27.000	.020

a. Design: Intercept+GROUP

b. Computed using alpha = .05

c. Exact statistic

SPSS Output 3

Univariate Test Statistics

SPSS Output 4 initially shows a summary table of Levene's test of equality of variances for each of the dependent variables. These tests are the same as would be found if a one-way ANOVA had been conducted on each dependent variable in turn (see Field, 2000, chaper 7). Levene's test should be non-significant for all dependent variables if the assumption of homogeneity of variance has been met. The results for these data clearly show that the assumption has been met. This finding not only gives us confidence in the reliability of the univariate tests to follow, but also strengthens the case for assuming that the multivariate test statistics are robust.

The next part of the output contains the ANOVA summary table for the dependent variables. The row of interest is that labelled *GROUP*. The row labelled *GROUP* contains an ANOVA summary table for each of the dependent variables, and values are given for the sums of squares for both actions and thoughts. The row labelled *Error* contains information about the residual sums of squares and mean squares for each of the dependent variables. The row labelled *Corrected Total* contains the values of the total sums of squares for each dependent variable. The important parts of this table are the columns labelled *F* and *Sig*. in which the *F*-ratios for each univariate ANOVA and their significance values are listed. The values associated with the univariate ANOVAs conducted after the MANOVA are *identical* to those obtained if one-way ANOVA was conducted on each dependent variable. This fact illustrates that MANOVA offers only hypothetical protection of inflated type I error rates: there is no real-life adjustment made to the values obtained.

The values of p in SPSS Output 4 indicate that there was a non-significant difference between therapy groups in terms of both obsession-related thoughts ($p = 0.136$) and obsession-related behaviours ($p = 0.080$). These two results should lead us to conclude that the type of therapy has had no significant effect on the levels of OCD experienced by clients. Those of you that are still awake may have noticed something odd about this example: the multivariate test statistics led us to conclude that therapy had had a significant impact on OCD, yet the univariate results indicate that therapy has not been successful. Before reading any further, have a think about why this anomaly has occurred.

The reason for the anomaly in these data is simple: the multivariate test takes account of the correlation between dependent variables and so for these data it has more power to detect group differences. With this knowledge in mind, the univariate tests are not particularly useful for interpretation, because the groups differ along a combination of the dependent variables. To see how the dependent variables interact we need to carry out a discriminant function analysis, which will be described in a later section.

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

	F	df1	df2	Sig.
Number of obsession-related thoughts	.076	2	27	.927
Number of obsession-related behaviours	1.828	2	27	.180

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

a. Design: Intercept+GROUP

Tests of Between-Subjects Effects

Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	Number of obsession-related thoughts	19.467[b]	2	9.733	2.154	.136
	Number of obsession-related behaviours	10.467[c]	2	5.233	2.771	.080
Intercept	Number of obsession-related thoughts	6336.533	1	6336.533	1402.348	.000
	Number of obsession-related behaviours	616.533	1	616.533	326.400	.000
GROUP	Number of obsession-related thoughts	19.467	2	9.733	2.154	.136
	Number of obsession-related behaviours	10.467	2	5.233	2.771	.080
Error	Number of obsession-related thoughts	122.000	27	4.519		
	Number of obsession-related behaviours	51.000	27	1.889		
Total	Number of obsession-related thoughts	6478.000	30			
	Number of obsession-related behaviours	678.000	30			
Corrected Total	Number of obsession-related thoughts	141.467	29			
	Number of obsession-related behaviours	61.467	29			

a. Computed using alpha = .05

b. R Squared = .138 (Adjusted R Squared = .074)

c. R Squared = .170 (Adjusted R Squared = .109)

SPSS Output 4

Contrasts

Earlier I suggested carrying out a *simple* contrast that compares each of the therapy groups to the no-treatment control group. SPSS Output 5 shows the results of these contrasts. The table is divided into two sections conveniently labelled *1st vs. 3rd* and *2nd vs. 3rd* where the numbers correspond to the coding of the group variable (i.e. 1 represents the lowest code used in the data editor and 3 the highest). If you coded the group variable using the same codes as I did, then these contrasts represent CBT vs. NT and BT vs. NT respectively. Each contrast is performed on both dependent variables separately and so they are identical to the contrasts that would be obtained from a univariate ANOVA. The table provides values for the contrast estimate, and the hypothesized value (which will always be zero because we are testing the null hypothesis that the difference between groups is zero). The observed estimated difference is then tested to see whether it is significantly different from zero based on the standard error. A 95% confidence interval is produced for the estimated difference.

The first thing that you might notice is that SPSS does not produce an exact significance value for the contrast: so how can we tell whether the group differences are significant? The simple answer is to look at the confidence interval. Field (2000, chapters 4, 5 and 6) explains that a 95% confidence interval tells us the values of the difference between groups between which 95% of samples will fall. If these boundaries cross zero (i.e. the lower is a minus number and the upper a positive value), then this tells us that within our 95% of samples, a good

proportion of samples will have group differences of zero (i.e. there will be no difference between the groups). Therefore, we cannot be confident that the observed group difference is meaningful because a different sample would have given us no group difference. If, however, the confidence interval does not cross zero (i.e. both values are positive or negative), then we can be confident that we would find a difference between the groups in 95% of samples taken from the same population. As such, we can be confident that genuine group differences exist. The take-home message here is that if the confidence interval includes zero then the contrast is non-significant; if the confidence interval does not include zero then we can say that the contrast is significant at $p < 0.05$. For these data all confidence intervals include zero (the lower bounds are negative whereas the upper bounds are positive) and so no contrasts are significant. This was expected because the univariate ANOVAs were both non-significant and so we would not expect there to be group differences.

Contrast Results (K Matrix)

group Simple Contrast[a]			Dependent Variable: Number of obsession-related thoughts	Dependent Variable: Number of obsession-related behaviours
1st vs. 3rd	Contrast Estimate		-1.600	-1.000E-01
	Hypothesized Value		0	0
	Difference (Estimate - Hypothesized)		-1.600	-1.000E-01
	Std. Error		.951	.615
	95% Confidence Interval for Difference	Lower Bound	-4.144	-1.745
		Upper Bound	.944	1.545
2nd vs. 3rd	Contrast Estimate		.200	-1.300
	Hypothesized Value		0	0
	Difference (Estimate - Hypothesized)		.200	-1.300
	Std. Error		.951	.615
	95% Confidence Interval for Difference	Lower Bound	-2.344	-2.945
		Upper Bound	2.744	.345

a. Reference category = 3

SPSS Output 5

Following Up MANOVA with Discriminant Analysis

A significant MANOVA can be followed up using either univariate ANOVA or discriminant analysis. In the example in this chapter, the univariate ANOVAs were not a useful way of looking at what the multivariate tests showed because the relationship between dependent variables is obviously having an effect. However, these data were designed especially to illustrate how the univariate ANOVAs should be treated cautiously and in real life a significant MANOVA is likely to be accompanied by at least one significant ANOVA. However, this does not mean that the relationship between dependent variables is not important, and it is still extremely important to investigate the nature of this relationship. Discriminant analysis is the best way to achieve this, and I strongly recommend that you follow up a MANOVA with both univariate tests and discriminant analysis if you want to fully understand your data.

Discriminant analysis is quite straightforward in SPSS: to access the main dialog box simply follow the menu path **Analyze⇒Classify⇒Discriminant...**(see Figure 3). The main dialog box will list the variables in the data editor on the left-hand side and provides two spaces on the right: one for the group variable and one for the predictors. In discriminant analysis we look to see how we can best separate (or discriminate) a set of groups using several predictors (so it is a little like logistic regression but where there are several groups rather than two). It might be confusing to think of actions and thoughts as independent variables (after all, they were dependent variables in the MANOVA!) which is why I refer to them as predictors.

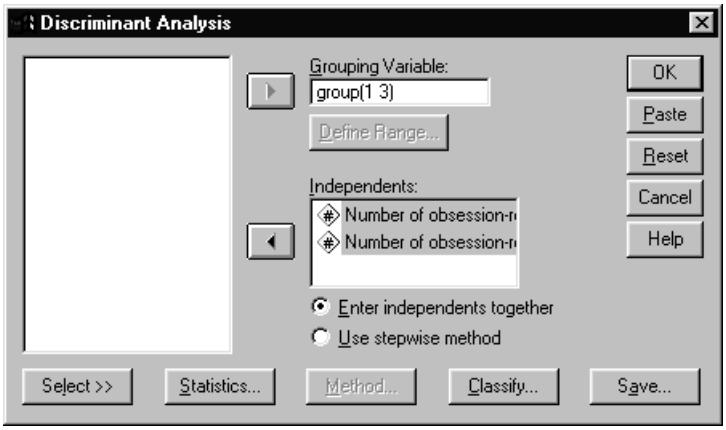

Figure 3: Main dialog box for discriminant analysis

To run the analysis, select the variable **group** and transfer it to the box labelled *Grouping Variable* by clicking on ▸. Once this variable has been transferred, the Define Range... button will become active and you should click this button to activate a dialog box in which you can specify the value of the highest and lowest coding values. Once you have specified the codings used for the grouping variable, you should select the variables **actions** and **thoughts** and transfer them to the box labelled *Independents* by clicking on ▸. There are two options available to determine how the predictors are entered into the model. The default is that both predictors are entered together and this is the option we require (because in MANOVA the dependent variables are analysed simultaneously). It is possible to enter the dependent variables in a stepwise manor and if this option is selected the Method... button becomes active, which opens a dialog box for specifying the criteria upon which predictors are entered.

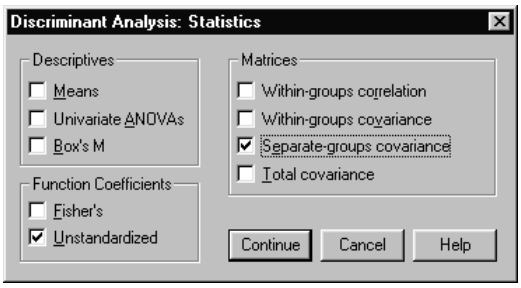

Figure 4: Statistics options for discriminant analysis

For the purpose of following up MANOVA, we need only be concerned with the remaining options. Click on Statistics... to activate the dialog box in Figure 4. This dialog box allows us to request group means, univariate ANOVAs and Box's test of equality of covariance matrices, all of which have already been provided in the MANOVA output (so we need not ask for them again). There is also an option to display a separate-groups covariance matrix, which can be useful for gaining insight into the relationships between dependent variables for each group (this matrix is something that the MANOVA procedure doesn't display and I recommend selecting it). Finally, we can ask for a total covariance matrix, which displays covariances and variances of the dependent variables overall. Another useful option is to select *Unstandardized* function coefficients. This option will produce the unstandardized βs for each variate.When you have finished with this dialog box, click on Continue to return to the main dialog box.

If you click on Classify... you will access the dialog box in Figure 5. In this dialog box there are several options available. First, you can select how prior probabilities are determined: if your group sizes are equal then you should leave the default setting as it is; however, if you have an unbalanced design then it is beneficial to base prior probabilities on the observed group sizes. The default option for basing the analysis on the within-group covariance matrix is fine

(because this is the matrix upon which the MANOVA is based). You should also request a combined-groups plot, which will plot the variate scores for each subject grouped according to the therapy they were given. The separate-groups plots show the same thing but using different graphs for each of the groups; when the number of groups is small it is better to select a combined plot because they are easier to interpret. The remaining options are of little interest when using discriminant analysis to follow up MANOVA. The only option that is useful is the summary table, which provides an overall gauge of how well the discriminant variates classify the actual subjects. When you have finished with the options click on Continue to return to the main dialog box.

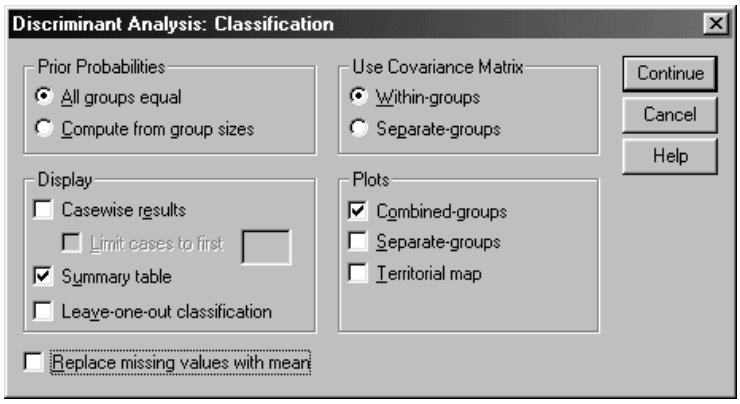

Figure 5: Discriminant analysis classification options

Output from the Discriminant Analysis

SPSS Output 6 shows the covariance matrices for separate groups (selected in Figure 4). These matrices are made up of the variances of each dependent variable for each group. The covariances are obtained by taking the cross-products between the dependent variables for each group and dividing each by 9—the degrees of freedom, $N-1$ (where N is the number of observations). The values in this table are useful because they give us some idea of how the relationship between dependent variables changes from group to group. For example, in the CBT group behaviours and thoughts have virtually no relationship because the covariance is almost zero. In the BT group thoughts and actions are positively related, so as the number of behaviours decrease, so does the number of thoughts. In the NT condition there is a negative relationship, so if the number of thoughts increases then the number of behaviours decrease. It is important to note that these matrices don't tell us about the substantive importance of the relationships because they are unstandardized (see Field, 2000, Chapter 3), they merely give a basic indication.

Covariance Matrices

group		Number of obsession-related behaviours	Number of obsession-related thoughts
CBT	Number of obsession-related behaviours	1.433	4.444E-02
	Number of obsession-related thoughts	4.444E-02	3.600
BT	Number of obsession-related behaviours	3.122	2.511
	Number of obsession-related thoughts	2.511	4.400
No Treatment Control	Number of obsession-related behaviours	1.111	-1.111
	Number of obsession-related thoughts	-1.111	5.556

SPSS Output 6

SPSS Output 7 shows the initial statistics from the discriminant analysis. At first we are told the eigenvalues for each variate and you should note that the values correspond to the values of the diagonal elements of the matrix HE^{-1}. These eigenvalues are converted into percentage of variance accounted for, and the first variate accounts for 82.2% of variance compared to the second variate, which accounts for only 17.8%. The next part of the output shows Wilks's lambda which has the same value (0.699), degrees of freedom (4) and significance value (0.05) as in the MANOVA (see SPSS Output 3). The important point to note from this table is that only one of the variates is significant (the second variate is non-significant, $p = 0.173$). Therefore, the group differences shown by the MANOVA can be explained in terms of *one* underlying dimension.

Eigenvalues

Function	Eigenvalue	% of Variance	Cumulative %	Canonical Correlation
1	.335[a]	82.2	82.2	.501
2	.073[a]	17.8	100.0	.260

a. First 2 canonical discriminant functions were used in the analysis.

Wilks's Lambda

Test of Function(s)	Wilks's Lambda	Chi-square	df	Sig.
1 through 2	.699	9.508	4	.050
2	.932	1.856	1	.173

SPSS Output 7

The tables in SPSS Output 8 are the most important for interpretation. The first table shows the standardized discriminant function coefficients for the two variates. These values are standardized versions of the values in the eigenvectors (see lecture). If you recall that the variates can be expressed in terms of a linear regression equation, the standardized discriminant function coefficients are equivalent to the standardized betas in regression. Hence, these coefficients tell us the relative contribution of each variable to the variates. It is clear from the size of the values for these data that the number of obsessive behaviours has a greater contribution to the first variate than the number of thoughts, but that the opposite is true for variate 2. Also, remembering that standardized beta coefficients vary within ± 1, it is noteworthy that both variables have a large contribution to the first variate (i.e. they are both important) because their values are quite close to 1 and -1 respectively. Bearing in mind that only the first variate is important, we can conclude that it is necessary to retain both dependent variables in the set of discriminators (because their standardized weights are of a similar magnitude). The fact that one dependent variable has a negative weight and one a positive weight indicates that group differences are explained by the difference between dependent variables.

Standardized Canonical Discriminant Function Coefficients

	Function	
	1	2
Number of obsession-related behaviours	.829	.584
Number of obsession-related thoughts	-.713	.721

Structure Matrix

	Function	
	1	2
Number of obsession-related behaviours	.711*	.703
Number of obsession-related thoughts	-.576	.817*

Pooled within-groups correlations between discriminating variables and standardized canonical discriminant functions
Variables ordered by absolute size of correlation within function.

*. Largest absolute correlation between each variable and any discriminant function

SPSS Output 8

Another way of looking at the relationship between dependent variables and discriminant variates is to look at the structure matrix, which gives the canonical variate correlation coefficients. These values are comparable to factor loadings (see Field, Chapter 11) and indicate the substantive nature of the variates. Bargman (1970) argues that when some dependent variables have high canonical variate correlations while others have low ones then the ones with high correlations contribute most to group separation. As such they represent the relative contribution of each dependent variable to group separation (see Bray and Maxwell, 1985, pp. 42–45). We are again interested only in the first variate (because the second was non-significant) and looking at the structure matrix we can conclude that the number of behaviours was slightly more important in differentiating the three groups (because 0.711 is greater than 0.576). However, the number of thoughts is still very important because the value of the correlation is quite large. As with the standardized weights, the fact that one dependent variable has a positive correlation, whereas the other has a negative one, indicates that group separation is determined by the difference between the dependent variables.

Canonical Discriminant Function Coefficients

	Function	
	1	2
Number of obsession-related behaviours	.603	.425
Number of obsession-related thoughts	-.335	.339
(Constant)	2.139	-6.857

Unstandardized coefficients

Functions at Group Centroids

	Function	
group	1	2
CBT	.601	-.229
BT	-.726	-.128
No Treatment Control	.125	.357

Unstandardized canonical discriminant functions evaluated at group means

SPSS Output 9

The next part of the output (SPSS Output 9) tells us first the canonical discriminant function coefficients, which are the unstandardized versions of the standardized coefficients described above. These values are less useful than the standardized versions, but do demonstrate from where the standardized versions come. The next table gives the values of the variate centroids

for each group. The centroids are simply the mean variate scores for each group. For interpretation we should look at the sign of the centroid (positive or negative), and from these data it looks as if variate 1 discriminates the BT group from the other two (notably the CBT group because the difference between centroids is greatest for these groups). The second variate (which was non-significant) seems to discriminate the NT group from the two experimental groups (but not significantly so).

The relationship between the variates and the groups is best illuminated using a combined-groups plot (selected using the dialog box in Figure 5). This graph plots the variate scores for each subject, grouped according to the experimental condition to which that subject belonged. In addition, the group centroids are indicated which are the average variate scores for each group. Figure 6 shows this plot for the OCD data, and what is clear from the position of the centroids (the big circles labelled with the group initials) is that variate 1 discriminates the BT group from the CBT (look at the horizontal distance between these centroids). The second variate does not differentiate any groups: we know this already because it was non-significant, but the plot shows that the vertical distances between group centroids is very small which indicates no group separation on this variate.

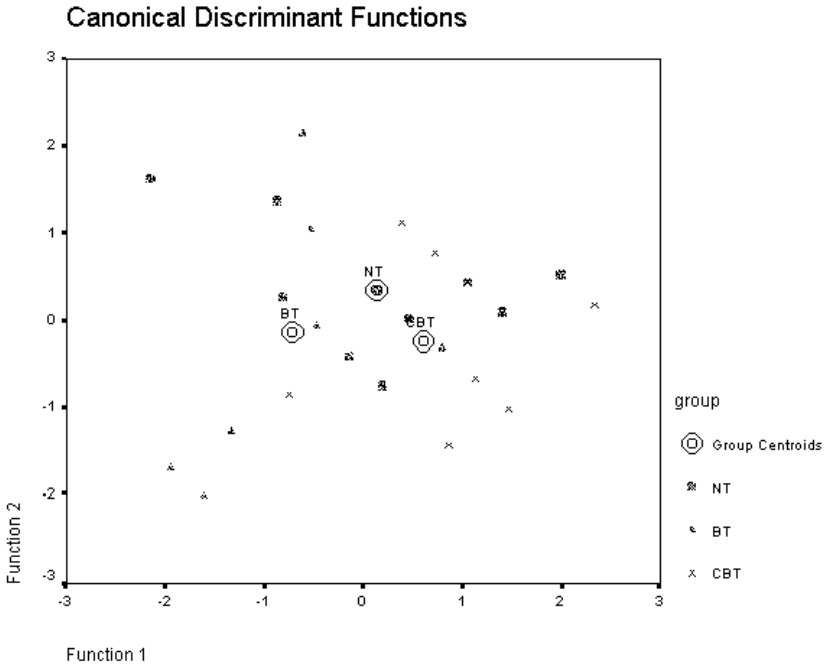

Figure 6: Combined-groups plot

Reference

This handout is an edited version of chapter 10 from:

Field, A. P. (2000). *Discovering statistics using SPSS for Windows: advanced techniques for the beginner*. London: Sage publications.

So, it's copyright etc. etc.

Some other Questions to Try

Question 1

A clinician noticed that a number of his manic psychotic patients did chicken impersonations in public. He wondered whether this behaviour could be used to diagnose this disorder and so decided to compare his patients against a normal sample. He observed 5 of his patients as they went through a normal day and also observed five control subjects who were postgraduate students at the University of Sussex. He measured them along two dependent variables: first, how many chicken impersonations they did in the streets of Brighton over the course of a day, and second, how good their impersonations were (as scored out of 10 by an independent farmyard noise expert). The data are in Table 1:

Group	DV1 (Quality of Impersonation)	DV2 (Number of Impersonations)
Manic Psychosis	3	9
Manic Psychosis	5	15
Manic Psychosis	5	15
Manic Psychosis	4	13
Manic Psychosis	1	8
Sussex Postgrad	8	13
Sussex Postgrad	4	9
Sussex Postgrad	4	7
Sussex Postgrad	2	7
Sussex Postgrad	9	15

Table 1

1. Carry out the appropriate test to find out whether there are any differences between the manic psychosis group and the Sussex postgraduates in terms of the quality and quantity of their chicken impersonations.
2. Comment on the important parts of the output and interpret the results.
3. There is something strange about the results - what is it? Why do you think this has happened?

Question 2

I was interested in whether students' knowledge of different aspects of psychology improved throughout their degree. I took a sample of first years, second years and third years and gave them 5 tests (scored out of 15) representing different aspects of psychology.

- **Exper**: Experimental Psychology (Cognitive, Neuropsychology etc.**)**
- **Stats**: Statistics
- **Social:** Social Psychology
- **Develop:** Developmental Psychology
- **Person:** Personality

Your task:

Carry out an appropriate *general* analysis to determine whether there are overall group differences along these 5 measures.

Look at the scale-by-scale analyses of group differences produced in the output and interpret the results accordingly.

Select *contrasts* that test the hypothesis that second and third years will score higher than first years on all scales.

Select tests that compare all groups to each other—briefly compare these results with the contrasts.

Carry out a separate analysis in which you test whether a combination of the measures can successfully discriminate the groups (comment only briefly on this analysis). Include only those scales that revealed group differences for the *contrasts*. How do the results help you to explain the findings of your initial analysis?

GROUP	EXPER	STATS	SOCIAL	DEVELOP	PERSON
1st Year	4.00	12.00	12.00	12.00	11.00
	10.00	15.00	14.00	14.00	13.00
	6.00	5.00	7.00	8.00	12.00
	6.00	6.00	9.00	8.00	6.00
	4.00	5.00	9.00	13.00	13.00
	7.00	8.00	10.00	9.00	10.00
	3.00	2.00	7.00	9.00	6.00
	8.00	9.00	14.00	15.00	15.00
	6.00	7.00	10.00	10.00	9.00
	3.00	6.00	14.00	14.00	15.00
	5.00	8.00	8.00	9.00	7.00
2nd Year	6.00	10.00	8.00	10.00	12.00
	6.00	6.00	8.00	8.00	8.00
	5.00	11.00	13.00	10.00	11.00
	2.00	4.00	7.00	7.00	7.00
	8.00	7.00	5.00	9.00	7.00
	6.00	6.00	5.00	7.00	7.00
	5.00	8.00	7.00	8.00	6.00
	5.00	10.00	6.00	8.00	10.00
	8.00	10.00	12.00	12.00	10.00
	4.00	9.00	10.00	8.00	7.00
	5.00	10.00	11.00	9.00	10.00
	4.00	9.00	10.00	9.00	8.00
	7.00	10.00	7.00	7.00	8.00
	4.00	7.00	8.00	9.00	5.00
	7.00	14.00	14.00	13.00	11.00
	6.00	8.00	6.00	8.00	8.00
3rd Year	5.00	11.00	6.00	9.00	10.00
	10.00	4.00	8.00	5.00	9.00
	7.00	13.00	8.00	8.00	12.00
	6.00	8.00	10.00	8.00	9.00
	5.00	12.00	10.00	14.00	10.00
	9.00	11.00	11.00	12.00	10.00
	10.00	9.00	9.00	12.00	9.00
	6.00	9.00	7.00	9.00	7.00
	6.00	14.00	8.00	8.00	4.00
	9.00	12.00	9.00	9.00	5.00
	5.00	14.00	8.00	6.00	7.00
	9.00	13.00	12.00	11.00	11.00
	4.00	6.00	8.00	3.00	6.00