Chapter 16: Multivariate analysis of variance (MANOVA)

Smart Alex's Solutions

Task 1

> *A clinical psychologist decided to compare his patients against a normal sample. He observed 10 of his patients as they went through a normal day. He also observed 10 lecturers at the University of Sussex. He measured all participants using two dependent variables: how many chicken impersonations they did, and, how good their impersonations were (as scored out of 10 by an independent farmyard noise expert). The data are in the file* ***Chicken.sav****. Use MANOVA and discriminant function analysis to find out whether these variables could be used to distinguish manic psychotic patients from those without the disorder.*

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. It seems that manic psychotics and Sussex lecturers do pretty similar numbers of chicken impersonations (lecturers do slightly fewer actually, but they are of a higher quality).

Descriptive Statistics

	GROUP	Mean	Std. Deviation	N
QUALITY	Manic Psychosis	6.7000	1.05935	10
	Sussex Lecturers	7.6000	2.98887	10
	Total	7.1500	2.23077	20
QUANTITY	Manic Psychosis	12.1000	4.22821	10
	Sussex Lecturers	10.7000	4.37290	10
	Total	11.4000	4.24760	20

Output 1

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 three 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* is given as .000 (which is less than .05); hence, the covariance matrices are not equal (i.e. they are significantly different) and the assumption is broken. However, because group sizes are equal we can ignore this test because Pillai's trace should be robust to this violation (fingers crossed!).

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

Box's M	20.926
F	6.135
df1	3
df2	58320.000
Sig.	.000

Tests the null hypothesis that the observed covariance matrices of the dependent variables are equal across groups.

a. Design: Intercept+GROUP

Output 2

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 manic psychotics and Sussex lecturers differ along the two dimensions of quality and quantity of chicken impersonations. The column of real interest is the one containing the significance values of these *F*-ratios. For these data, all test statistics are significant with $p = .032$ (which is less than .05). From this result we should probably conclude that the groups do indeed differ in terms of the quality and quantity of their chicken impersonations; however, this effect needs to be broken down to find out exactly what's going on.

Multivariate Tests[b]

Effect		Value	F	Hypothesis df	Error df	Sig.
Intercept	Pillai's Trace	.919	96.201[a]	2.000	17.000	.000
	Wilks' Lambda	.081	96.201[a]	2.000	17.000	.000
	Hotelling's Trace	11.318	96.201[a]	2.000	17.000	.000
	Roy's Largest Root	11.318	96.201[a]	2.000	17.000	.000
GROUP	Pillai's Trace	.333	4.250[a]	2.000	17.000	.032
	Wilks' Lambda	.667	4.250[a]	2.000	17.000	.032
	Hotelling's Trace	.500	4.250[a]	2.000	17.000	.032
	Roy's Largest Root	.500	4.250[a]	2.000	17.000	.032

a. Exact statistic

b. Design: Intercept+GROUP

Output 3

Output 4 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. 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 for the quantity of chicken impersonations but has been broken for the quality of impersonations. This should dent our confidence in the reliability of the univariate tests to follow.

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

	F	df1	df2	Sig.
QUALITY	11.135	1	18	.004
QUANTITY	.256	1	18	.619

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

a. Design: Intercept+GROUP

Output 4

The ANOVA summary table for the dependent variables is shown in Output 5. The row of interest is that labelled *GROUP* (you'll notice that the values in this row are the same as for the row labelled *Corrected Model*: this is because the model fitted to the data contains only one independent variable: **group**). The row labelled *GROUP* contains an ANOVA summary table for quality and quantity of chicken impersonations, respectively. The values of p indicate that there was a non-significant difference between groups in terms of both (p is greater than .05 in both cases). The multivariate test statistics led us to conclude that the groups *did* differ in terms of the quality and quantity of their chicken impersonations yet the univariate results contradict this!

Tests of Between-Subjects Effects

Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	QUALITY	4.050[a]	1	4.050	.806	.381
	QUANTITY	9.800[b]	1	9.800	.530	.476
Intercept	QUALITY	1022.450	1	1022.450	203.360	.000
	QUANTITY	2599.200	1	2599.200	140.497	.000
GROUP	QUALITY	4.050	1	4.050	.806	.381
	QUANTITY	9.800	1	9.800	.530	.476
Error	QUALITY	90.500	18	5.028		
	QUANTITY	333.000	18	18.500		
Total	QUALITY	1117.000	20			
	QUANTITY	2942.000	20			
Corrected Total	QUALITY	94.550	19			
	QUANTITY	342.800	19			

a. R Squared = .043 (Adjusted R Squared = -.010)

b. R Squared = .029 (Adjusted R Squared = -.025)

Output 5

We don't need to look at contrasts because the univariate tests were non-significant (and in any case there were only two groups and so no further comparisons would be necessary), and instead, to see how the dependent variables interact, we need to carry out a discriminant function analysis (DFA).

Wilks' Lambda

Test of Function(s)	Wilks' Lambda	Chi-square	df	Sig.
1	.667	6.893	2	.032

Output 6

The initial statistics from the DFA tell us that there was only one variate (because there are only two groups) and this variate is significant (Output 6). Therefore, the group differences shown by the MANOVA can be explained in terms of *one* underlying dimension.

Standardized Canonical Discriminant Function Coefficients

	Function
	1
QUALITY	1.859
QUANTITY	-1.829

Output 7

The standardized discriminant function coefficients (Output 7) tell us the relative contribution of each variable to the variates. Both quality and quantity of impersonations have similar-sized coefficients indicating that they have equally strong influence in discriminating the groups. However, they have the opposite sign, which suggests that that group differences are explained by the difference between the quality and quantity of impersonations.

Functions at Group Centroids

GROUP	Function 1
Manic Psychosis	-.671
Sussex Lecturers	.671

Unstandardized canonical discriminant functions evaluated at group means

Output 8

The variate centroids for each group (Output 8) confirm that variate 1 discriminates the two groups because the manic psychotics have a negative coefficient and the Sussex lecturers have a positive one. There won't be a combined-groups plot because there is only one variate.

Overall we could conclude that manic psychotics are distinguished from Sussex lecturers in terms of the difference between the pattern of results for quantity of impersonations compared to quality. If we look at the means we can see that manic psychotics produce slightly more impersonations than Sussex lecturers (but remember from the non-significant univariate tests that this isn't sufficient, alone, to differentiate the groups), but the lecturers produce impersonations of a higher quality (but again remember that quality alone is not enough to differentiate the groups). Therefore, although the manic psychotics and Sussex lecturers produce similar numbers of impersonations of similar quality (see univariate tests), if we combine the quality and quantity we can differentiate the groups.

Task 2

A news story claimed that children who lie would become successful citizens (http://bit.ly/ammQNT). I was intrigued because although the article cited a lot of well-conducted work by Dr. Khang Lee that shows that children lie, I couldn't find anything in that research that supported the journalist's claim that children who lie become successful citizens. Imagine a Huxley-esque parallel universe in which the government was stupid enough to believe the contents of this newspaper story and decided to implement a systematic programme of infant conditioning. Some infants were trained not to lie, others were bought up as normal, and a final group was trained in the art of lying. Thirty years later, they collected data on how successful these children were as adults. They measured their ***salary****, and two indices out of 10 (10 = as successful as could possibly be, 0 = better luck in your next life) of how successful their* ***family*** *and* ***work*** *life is. The data are in* ***Lying.sav****. Use MANOVA and discriminant function analysis to find out whether lying really does make you a better citizen.*

Access the main MANOVA dialog box by selecting Analyze General Linear Model ▸ Multivariate... . Select the three dependent variables from the variables list (i.e., **Salary**, **Family Life** and **Work Life**) and drag them to the *Dependent Variables* box (or click on). Select **Lying** from the variables list and drag it (or click on) to the *Fixed Factor(s)* box. Your completed dialog box should look like Figure .

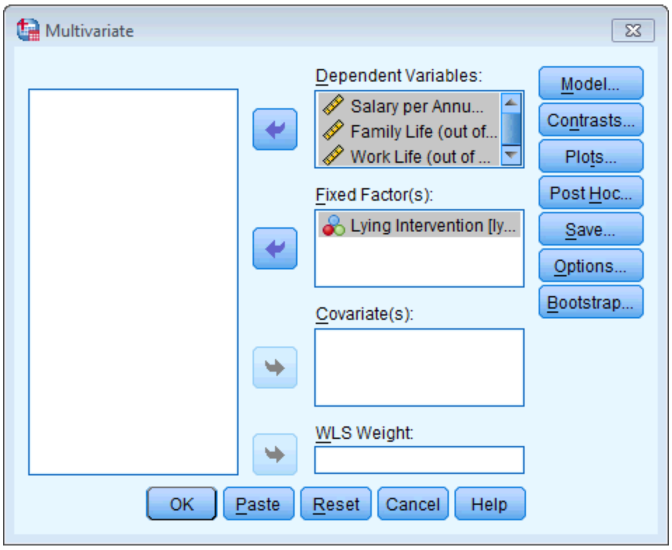

Figure 1

The default way to follow up a MANOVA is to look at individual univariate ANOVAs for each dependent variable. 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 group to the lying encouraged group. The lying encouraged 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 (see Figure).

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... . For the purposes of this example, I suggest selecting two of my usual recommendations: REGWQ and Games–Howell (see Figure). Once you have selected *post hoc* tests return to the main dialog box.

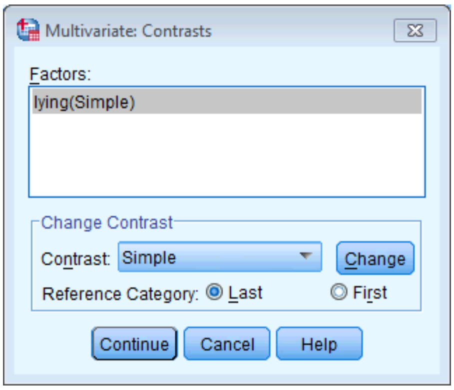

Figure 2

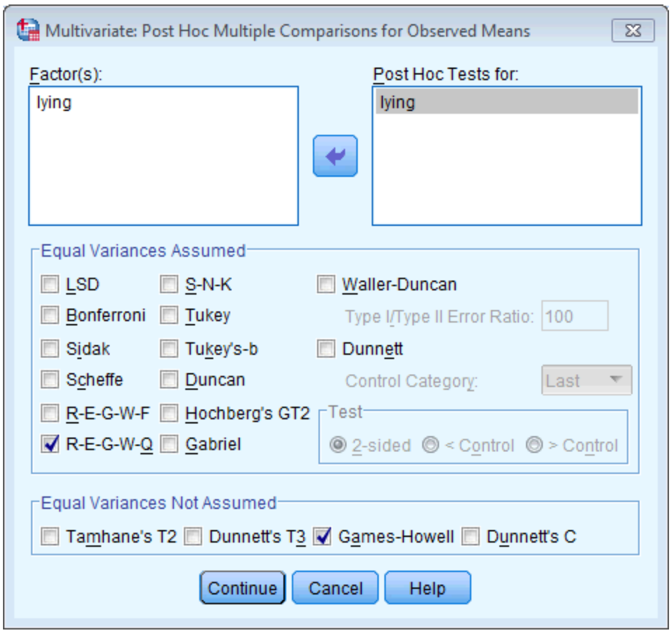

Figure 3

To access the *options* dialog box, click on Options... in the main dialog box and select the boxes as I have in Figure .

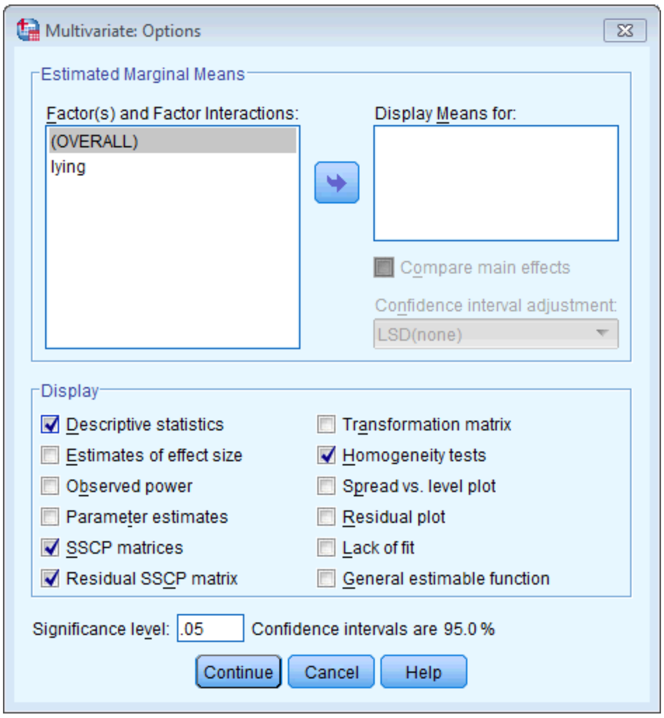

Figure 4

Descriptive Statistics

	Lying Intervention	Mean	Std. Deviation	N
Salary per Annum	Lying Prevented	28256.29	8349.676	14
	Normal Parenting	30918.86	5719.191	14
	Lying Encouraged	35411.57	8110.985	14
	Total	31528.90	7891.008	42
Family Life (out of 10)	Lying Prevented	6.36	2.898	14
	Normal Parenting	3.93	2.369	14
	Lying Encouraged	3.21	1.968	14
	Total	4.50	2.743	42
Work Life (out of 10)	Lying Prevented	4.14	2.179	14
	Normal Parenting	6.21	2.694	14
	Lying Encouraged	6.64	2.951	14
	Total	5.67	2.791	42

Output 9

Output contains the group means and standard deviations for each dependent variable in turn, split by the independent variable **Lying Intervention**. The means show that children encouraged to lie landed the best and highest-paid jobs, but had the worst family success compared to the other two groups. Children who were trained not to lie had great family lives but not so great jobs compared to children who were brought up to lie and children who experienced normal parenting. Finally, children who were in the normal parenting group (if that exists!) were pretty middle of the road compared to the other two groups.

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

Box's M	15.037
F	1.112
df1	12
df2	7371.000
Sig.	.345

Tests the null hypothesis that the observed covariance matrices of the dependent variables are equal across groups.

a. Design: Intercept + lying

Output 10

Output shows Box's test of the assumption of equality of covariance matrices. This statistic is non-significant, p = .345 (which is greater than .05), hence the covariance matrices are roughly equal as assumed.

Output shows the main table of results. The column of real interest is the one containing the significance values of the F-ratios. For these data, Pillai's trace (p = .002), Wilks's lambda (p = .001), Hotelling's trace (p < .001) and Roy's largest root (p < .001) all reach the criterion for significance at the .05 level. Therefore, we can conclude that the type of lying intervention had a significant effect on success later on in life. The nature of this effect is not clear from the multivariate test statistic: it tells us nothing about which groups differed from which, or about whether the effect of lying intervention was on work life, family life, salary, or a combination of all three. To determine the nature of the effect, a discriminant analysis would be helpful, but for some reason SPSS provides us with univariate tests instead.

Multivariate Tests[a]

Effect		Value	F	Hypothesis df	Error df	Sig.
Intercept	Pillai's Trace	.957	275.252[b]	3.000	37.000	.000
	Wilks' Lambda	.043	275.252[b]	3.000	37.000	.000
	Hotelling's Trace	22.318	275.252[b]	3.000	37.000	.000
	Roy's Largest Root	22.318	275.252[b]	3.000	37.000	.000
lying	Pillai's Trace	.478	3.979	6.000	76.000	.002
	Wilks' Lambda	.536	4.513[b]	6.000	74.000	.001
	Hotelling's Trace	.839	5.036	6.000	72.000	.000
	Roy's Largest Root	.807	10.219[c]	3.000	38.000	.000

a. Design: Intercept + lying

b. Exact statistic

c. The statistic is an upper bound on F that yields a lower bound on the significance level.

Output 11

Output 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. Levene's test should be non-significant for all dependent variables if the assumption of homogeneity of variance has been met. We can see here that the assumption has been met ($p > .05$ in all cases), which strengthens the case for assuming that the multivariate test statistics are robust.

Output contains an ANOVA summary table for each of the dependent variables. The *F*-ratios for each univariate ANOVA and their significance values are listed in the columns labelled *F* and *Sig*. These values are *identical* to those obtained if one-way ANOVA was conducted on each dependent variable independently. As such, MANOVA offers only *hypothetical* protection of inflated Type I error rates: there is no real-life adjustment made to the values obtained.

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

	F	df1	df2	Sig.
Salary per Annum	.527	2	39	.594
Family Life (out of 10)	1.251	2	39	.297
Work Life (out of 10)	.996	2	39	.379

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

a. Design: Intercept + lying

Output 1

Tests of Between-Subjects Effects

Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	Salary per Annum	366202116[a]	2	183101058	3.265	.049
	Family Life (out of 10)	76.000[b]	2	38.000	6.374	.004
	Work Life (out of 10)	50.048[c]	2	25.024	3.624	.036
Intercept	Salary per Annum	4.175E+10	1	4.175E+10	744.604	.000
	Family Life (out of 10)	850.500	1	850.500	142.665	.000
	Work Life (out of 10)	1348.667	1	1348.667	195.324	.000
lying	Salary per Annum	366202116	2	183101058	3.265	.049
	Family Life (out of 10)	76.000	2	38.000	6.374	.004
	Work Life (out of 10)	50.048	2	25.024	3.624	.036
Error	Salary per Annum	2.187E+9	39	56071437.5		
	Family Life (out of 10)	232.500	39	5.962		
	Work Life (out of 10)	269.286	39	6.905		
Total	Salary per Annum	4.430E+10	42			
	Family Life (out of 10)	1159.000	42			
	Work Life (out of 10)	1668.000	42			
Corrected Total	Salary per Annum	2.553E+9	41			
	Family Life (out of 10)	308.500	41			
	Work Life (out of 10)	319.333	41			

a. R Squared = .143 (Adjusted R Squared = .100)

b. R Squared = .246 (Adjusted R Squared = .208)

c. R Squared = .157 (Adjusted R Squared = .113)

Output 2

The values of p in Output indicate that there was a significant difference between intervention groups in terms of salary ($p = .049$), family life ($p = .004$), and work life ($p = .036$).

We should conclude that the type of intervention had a significant effect on the later success of children. However, this effect needs to be broken down to find out exactly what's going on.

Contrast Results (K Matrix)

Lying Intervention Simple Contrast[a]			Dependent Variable: Salary per Anum	Family Life (out of 10)	Work Life (out of 10)
Level 1 vs. Level 3	Contrast Estimate		-7155.286	3.143	-2.500
	Hypothesized Value		0	0	0
	Difference (Estimate - Hypothesized)		-7155.286	3.143	-2.500
	Std. Error		2830.231	.923	.993
	Sig.		.016	.002	.016
	95% Confidence Interval for Difference	Lower Bound	-12879.967	1.276	-4.509
		Upper Bound	-1430.604	5.009	-.491
Level 2 vs. Level 3	Contrast Estimate		-4492.714	.714	-.429
	Hypothesized Value		0	0	0
	Difference (Estimate - Hypothesized)		-4492.714	.714	-.429
	Std. Error		2830.231	.923	.993
	Sig.		.120	.444	.668
	95% Confidence Interval for Difference	Lower Bound	-10217.396	-1.152	-2.437
		Upper Bound	1231.967	2.581	1.580

a. Reference category = 3

Output 3

Looking at Output , we can see that when we compare children who were prevented from lying (level 1) with those who were encouraged to lie (level 3), there were significant differences in salary (p = .016), family success (p = .002) and work success (p = .016). Looking back at the means, we can see that children who were trained to lie had significantly higher salaries, significantly better work lives but significantly less successful family lives when compared to children who were prevented from lying.

When we compare children who experienced normal parenting (level 2) with those who were encouraged to lie (level 3), there were no significant differences between the three life success outcome variables (p > .05 in all cases).

In my opinion discriminant analysis is the best method for following up a significant MANOVA (see the book chapter) and we will do this next.

To access the main dialog box select Analyze Classify ▸ Discriminant...; it lists 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.

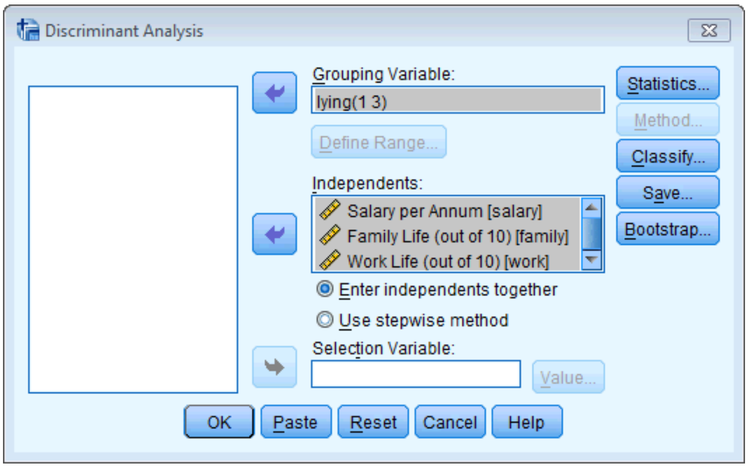

Figure 5

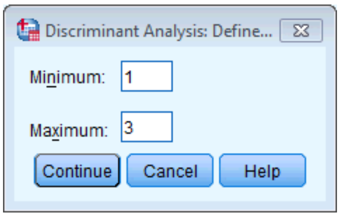

Figure 6

Your completed dialog boxes should look like Figure and 6 – see the book chapter for other options that are worth setting.

Output shows the covariance matrices for separate groups. These matrices are made up of the variances of each dependent variable for each group. 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 lying prevented group, all the dependent variables are positively related, so as one of the variables increases (e.g., success at work), the other two variables (family life and salary) increase also. In the normal parenting group, success at work is positively related to both family success and salary. However, salary and family success are negatively related, so as salary increases family success decreases and vice versa. Finally, in the lying encouraged group, salary has a positive relationship with both work success and family success, but success at work is negatively related to family success. It is important to note that these matrices don't tell us about the substantive importance of the relationships because they are unstandardized – they merely give a basic indication.

Covariance Matrices

Lying Intervention		Salary per Annum	Family Life (out of 10)	Work Life (out of 10)
Lying Prevented	Salary per Annum	69717086.8	2268.659	2263.956
	Family Life (out of 10)	2268.659	8.401	3.714
	Work Life (out of 10)	2263.956	3.714	4.747
Normal Parenting	Salary per Annum	32709144.9	-2999.242	1597.110
	Family Life (out of 10)	-2999.242	5.610	2.170
	Work Life (out of 10)	1597.110	2.170	7.258
Lying Encouraged	Salary per Annum	65788080.7	7075.022	7575.066
	Family Life (out of 10)	7075.022	3.874	-.841
	Work Life (out of 10)	7575.066	-.841	8.709

Output 4

Output shows the initial statistics from the discriminant analysis. First we are told the eigenvalues for each variate. These eigenvalues are converted into percentage of variance accounted for, and the first variate accounts for 96.1% of variance compared to the second variate, which accounts for only 3.9%. This table also shows the canonical correlation, which we can square to use as an effect size (just like R^2, which we have encountered in regression).

Eigenvalues

Function	Eigenvalue	% of Variance	Cumulative %	Canonical Correlation
1	.807[a]	96.1	96.1	.668
2	.033[a]	3.9	100.0	.178

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

Output 5

Wilks' Lambda

Test of Function(s)	Wilks' Lambda	Chi-square	df	Sig.
1 through 2	.536	23.698	6	.001
2	.968	1.220	2	.543

Output 17

Output shows the significance tests of both variates ('1 through 2' in the table), and the significance after the first variate has been removed ('2' in the table). So, effectively we test the model as a whole, and then peel away variates one at a time to see whether what's left is significant. In this case with two variates we get only two steps: the whole model, and then the model after the first variate is removed (which leaves only the second variate). When both variates are tested in combination Wilks's lambda has the same value (.536), degrees of freedom (6) and significance value (.001) as in the MANOVA (see Output). The important point to note from this table is that the two variates significantly discriminate the groups in combination (p = .001), but the second variate alone is non-significant, p = .543. Therefore, the group differences shown by the MANOVA can be explained in terms of *two* underlying dimensions in combination.

Output and 19 are the most important for interpretation. The coefficients in these tables tell us the relative contribution of each variable to the variates.

Standardized Canonical Discriminant Function Coefficients

	Function	
	1	2
Salary per Annum	.399	.936
Family Life (out of 10)	−.844	.226
Work Life (out of 10)	.620	−.526

Output 18

Structure Matrix

	Function	
	1	2
Family Life (out of 10)	−.635*	.197
Work Life (out of 10)	.477*	−.285
Salary per Annum	.421	.860*

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

Output 19

If we look at variate 1 first, family life has the opposite effect to work life and salary (work life and salary have positive relationships with this variate, whereas family life has a negative relationship). Given that these values (in both tables) can vary between 1 and −1, we can also see that family life has the strongest relationship, work life also has a strong relationship, whereas salary has a relatively weaker relationship to the first variate. The first variate, then, could be seen as one that differentiates family life from work life and salary (it affects family life in the opposite way to salary and work life). Salary has a very strong positive relationship to the second variate, family life has only a weak positive relationship and work life has a medium negative relationship to the second variate. This tells us that this variate represents something that affects salary and to a lesser degree family life in a different way than work life. Remembering that ultimately these variates are used to differentiate groups, we could say that the first variate differentiates groups by some factor that affects family differently than work and salary, whereas the second variate differentiates groups on some dimension that affects salary (and to a small degree family life) and work in different ways.

We can also use a combined-groups plot. This graph plots the variate scores for each person, grouped according to the experimental condition to which that person belonged. The graph (Figure) tell us that (look at the big squares) variate 1 discriminates the lying prevented group from the lying encouraged group (look at the horizontal distance between these centroids). The second variate differentiates the normal parenting group from the lying prevented and lying encouraged groups (look at the vertical distances), but this difference is not as dramatic as for the first variate. Remember that the variates significantly discriminate the groups in combination (i.e., when both are considered).

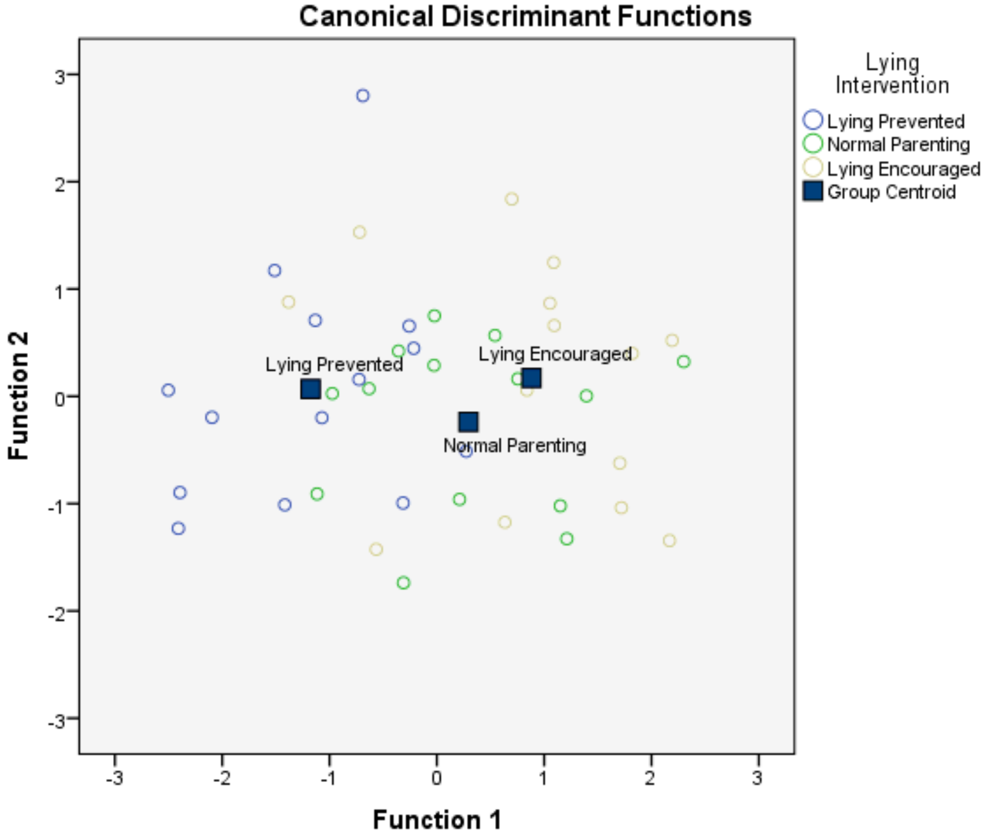

Figure 7

Reporting results

- ✓ Using Pillai's trace, there was a significant effect of lying on future success, $V = 0.48$, $F(6, 76) = 3.98$, $p = .002$. Separate univariate ANOVAs on the outcome variables revealed significant effects of lying on salary $F(2, 39) = 3.27$, $p = .049$, family, $F(2, 39) = 6.37$, $p = .004$ and work $F(2, 39) = 3.62$, $p = .036$.

- ✓ The MANOVA was followed up with discriminant analysis, which revealed two discriminant functions. The first explained 96.1% of the variance, canonical $R^2 = .45$, whereas the second explained only 3.9%, canonical $R^2 = .03$. In combination these discriminant functions significantly differentiated the lying intervention groups, $\Lambda = .536$, $\chi^2(6) = 23.70$, $p = .001$, but removing the first function indicated that the second function did not significantly differentiate the intervention groups, $\Lambda = .968$, $\chi^2(2) = 1.22$, $p = .543$. The correlations between outcomes and the discriminant functions revealed that salary loaded more highly onto the second function ($r = .94$) than the first (r = .40); family life loaded more highly on the first function (r = –.84) than the second function (r = .23); work life loaded fairly evenly onto both functions but in opposite directions (r = .62 for the first function and r = –.53 for the second). The discriminant function plot showed that the first function discriminated the lying intervention group from the lying prevented group, and the second function differentiated the normal parenting group from the two interventions.

Task 3

*I was interested in whether students' knowledge of different aspects of psychology improved throughout their degree (**Psychology.sav**). I took a sample of first years, second years and third years and gave them five tests (scored out of 15) representing different aspects of psychology: **Exper** (experimental psychology such as cognitive and neuropsychology); **Stats** (statistics); **Social** (social psychology); **Develop** (developmental psychology); **Person** (personality). (1) Determine whether there are overall group differences along these five measures. (2) Interpret the scale-by-scale analyses of group differences. (3) Select contrasts that test the hypothesis that second and third years will score higher than first years on all scales. (4) Select* post hoc *tests and compare these results to the contrasts. (5) Carry out a discriminant function analysis including only those scales that revealed group differences for the contrasts. Interpret the results*

Output 20 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.

Descriptive Statistics

	Gorup	Mean	Std. Deviation	N
Experimental Psychology	1st Year	5.6364	2.1574	11
	2nd Year	5.5000	1.5916	16
	3rd Year	7.0000	2.1213	13
	Total	6.0250	2.0062	40
Statistics	1st Year	7.5455	3.5599	11
	2nd Year	8.6875	2.3866	16
	3rd Year	10.4615	3.0988	13
	Total	8.9500	3.1211	40
Social Psychology	1st Year	10.3636	2.7303	11
	2nd Year	8.5625	2.8040	16
	3rd Year	8.7692	1.6408	13
	Total	9.1250	2.5236	40
Personality	1st Year	10.6364	3.3248	11
	2nd Year	8.4375	1.9990	16
	3rd Year	8.3846	2.3993	13
	Total	9.0250	2.6745	40
Developmental	1st Year	11.0000	2.6458	11
	2nd Year	8.8750	1.7078	16
	3rd Year	8.7692	3.0319	13
	Total	9.4250	2.5908	40

Output 6

Output 21 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 three 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 = .06 (which is greater than .05); hence, the covariance matrices are roughly equal and the assumption is tenable.

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

Box's M	54.241
F	1.435
df1	30
df2	3587
Sig.	.059

Tests the null hypothesis that the observed covariance matrices of the dependent variables are equal across groups.

a. Design: Intercept+GROUP

Output 7

Output 22 shows the main table of results. For our purposes, the group effects are of interest because they tell us whether or not the scores from different areas of psychology differ across the three years of the degree programme. The column of real interest is the one containing the significance values of these F-ratios. For these data, Pillai's trace (p =.02), Wilks's lambda (p = .012), Hotelling's trace (p =.007) and Roy's largest root (p =.01) all reach the criterion for significance of the .05 level. From this result we should probably conclude that the profile of knowledge across different areas of psychology does indeed change across the three years of the degree. The nature of this effect is not clear from the multivariate test statistic.

Multivariate Tests[c]

Effect		Value	F	Hypothesis df	Error df	Sig.
Intercept	Pillai's Trace	.960	159.166[a]	5.000	33.000	.000
	Wilks' Lambda	.040	159.166[a]	5.000	33.000	.000
	Hotelling's Trace	24.116	159.166[a]	5.000	33.000	.000
	Roy's Largest Root	24.116	159.166[a]	5.000	33.000	.000
GROUP	Pillai's Trace	.510	2.330	10.000	68.000	.020
	Wilks' Lambda	.522	2.534[a]	10.000	66.000	.012
	Hotelling's Trace	.853	2.730	10.000	64.000	.007
	Roy's Largest Root	.773	5.255[b]	5.000	34.000	.001

a. Exact statistic

b. The statistic is an upper bound on F that yields a lower bound on the significance level.

c. Design: Intercept+GROUP

Output 8

Output 23 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. 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.

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

	F	df1	df2	Sig.
Experimental Psychology	1.311	2	37	.282
Statistics	.746	2	37	.481
Social Psychology	2.852	2	37	.071
Personality	2.440	2	37	.101
Developmental	2.751	2	37	.077

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

a. Design: Intercept+GROUP

Output 9

The ANOVA summary table for the dependent variables is shown in Output 24. The row of interest is that labelled *GROUP*, which contains an ANOVA summary table for each of the areas of psychology. The values of *p* indicate that there was a non-significant difference between student groups in terms of all areas of psychology ($p > .05$ in each case). The multivariate test statistics led us to conclude that the student groups *did* differ significantly across the types of psychology, yet the univariate results contradict this (again ... I really should stop making up data sets that do this!).

Tests of Between-Subjects Effects

Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	Experimental Psychology	18.430[a]	2	9.215	2.461	.099
	Statistics	52.504[b]	2	26.252	2.967	.064
	Social Psychology	23.584[c]	2	11.792	1.941	.158
	Personality	39.415[d]	2	19.708	3.044	.060
	Developmental	37.717[e]	2	18.859	3.114	.056
Intercept	Experimental Psychology	1428.058	1	1428.058	381.378	.000
	Statistics	3093.775	1	3093.775	349.637	.000
	Social Psychology	3330.118	1	3330.118	548.129	.000
	Personality	3273.395	1	3273.395	505.575	.000
	Developmental	3562.212	1	3562.212	588.250	.000
GROUP	Experimental Psychology	18.430	2	9.215	2.461	.099
	Statistics	52.504	2	26.252	2.967	.064
	Social Psychology	23.584	2	11.792	1.941	.158
	Personality	39.415	2	19.708	3.044	.060
	Developmental	37.717	2	18.859	3.114	.056
Error	Experimental Psychology	138.545	37	3.744		
	Statistics	327.396	37	8.849		
	Social Psychology	224.791	37	6.075		
	Personality	239.560	37	6.475		
	Developmental	224.058	37	6.056		
Total	Experimental Psychology	1609.000	40			
	Statistics	3584.000	40			
	Social Psychology	3579.000	40			
	Personality	3537.000	40			
	Developmental	3815.000	40			
Corrected Total	Experimental Psychology	156.975	39			
	Statistics	379.900	39			
	Social Psychology	248.375	39			
	Personality	278.975	39			
	Developmental	261.775	39			

a. R Squared = .117 (Adjusted R Squared = .070)
b. R Squared = .138 (Adjusted R Squared = .092)
c. R Squared = .095 (Adjusted R Squared = .046)
d. R Squared = .141 (Adjusted R Squared = .095)
e. R Squared = .144 (Adjusted R Squared = .098)

Output 10

We don’t need to look at contrasts because the univariate tests were non-significant, and instead, to see how the dependent variables interact, we need to carry out a DFA.

Wilks' Lambda

Test of Function(s)	Wilks' Lambda	Chi-square	df	Sig.
1 through 2	.522	22.748	10	.012
2	.926	2.710	4	.608

Output 25

The initial statistics from the DFA (Output 25) tell us that only one of the variates is significant (the second variate is non-significant, p = 0.608). Therefore, the group differences shown by the MANOVA can be explained in terms of *one* underlying dimension.

Standardized Canonical Discriminant Function Coefficients

	Function	
	1	2
Experimental Psychology	.367	.789
Statistics	.921	-.081
Social Psychology	-.353	.319
Personality	-.260	.216
Developmental	-.618	.013

Output 26

The standardized discriminant function coefficients (Output 26) tell us the relative contribution of each variable to the variates. Looking at the first variate, it's clear that statistic has the greatest contribution to the first variate. Most interesting is that on the first variate, statistics and experimental psychology have positive weights, whereas social, developmental and personality have negative weights. This suggests that the group differences are explained by the difference between experimental psychology and statistics compared to other areas of psychology.

Functions at Group Centroids

	Function	
Gorup	1	2
1st Year	-1.246	.186
2nd Year	9.789E-02	-.333
3rd Year	.934	.252

Unstandardized canonical discriminant functions evaluated at group means

Output 27

The variate centroids for each group tell us that variate 1 discriminates the first years from second and third years because the first years have a negative value whereas the second and third years have positive values on the first variate (Output 27).

The relationship between the variates and the groups is best illuminated using a combined-groups plot. Figure 8 plots the variate scores for each person, grouped according to the year of their degree. In addition, the group centroids are indicated, which are the average variate scores for each group. The plot for these data confirms that variate 1 discriminates the first years from subsequent years (look at the horizontal distance between these centroids).

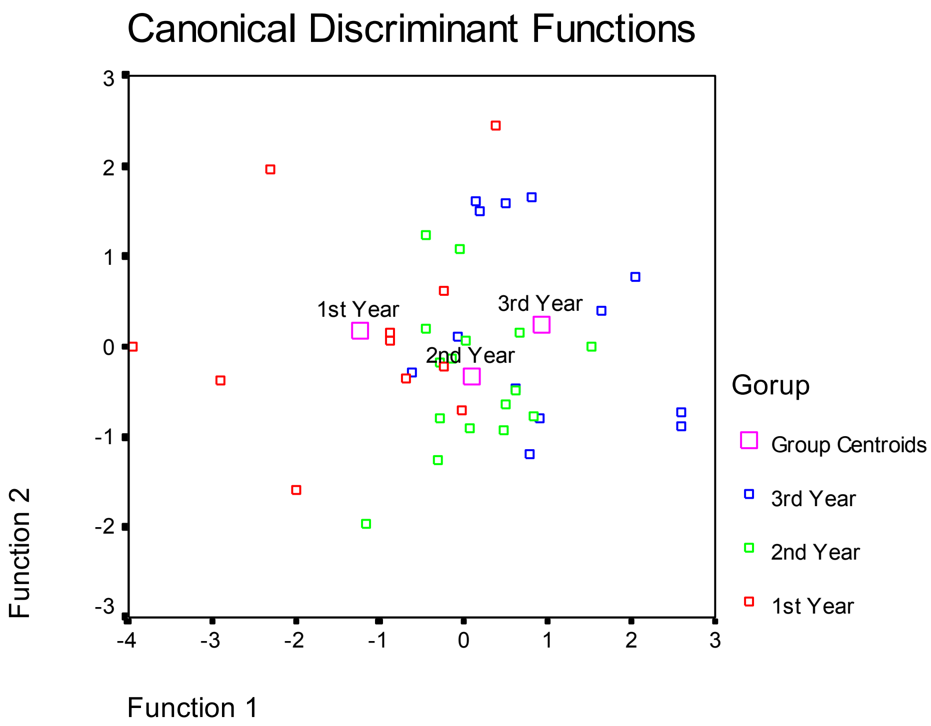

Figure 8

Overall we could conclude that different years are discriminated by different areas of psychology. In particular, it seems as though statistics and aspects of experimentation (compared to other areas of psychology) discriminate between first-year undergraduates and subsequent years. From the means, we could interpret this as first years struggling with statistics and experimental psychology (compared to other areas of psychology) but with their ability improving across the three years. However, for other areas of psychology, first years are relatively good but their abilities decline over the three years. Put another way, psychology degrees improve only your knowledge of statistics and experimentation.☺