The SAGE Encyclopedia of Educational Research, Measurement, and Evaluation

Multivariate Analysis of Variance

Contributors: David W. Stockburger
Edited by: Bruce B. Frey
Book Title: The SAGE Encyclopedia of Educational Research, Measurement, and Evaluation
Chapter Title: "Multivariate Analysis of Variance"
Pub. Date: 2018
Access Date: March 2, 2018
Publishing Company: SAGE Publications, Inc.
City: Thousand Oaks,
Print ISBN: 9781506326153
Online ISBN: 9781506326139
DOI: http://dx.doi.org/10.4135/9781506326139.n456
Print pages: 1120-1126

Multivariate analysis of variance (MANOVA) is an extension of univariate analysis of variance (ANOVA) in which the independent variable is some combination of group membership but there is more than one dependent variable. MANOVA is often used either when the researcher has correlated dependent variables or, instead of a repeated-measures ANOVA, to avoid the sphericity assumption. While MANOVA has the advantage of providing a single, more powerful test of multiple dependent variables, it can be difficult to interpret the results.

For example, a researcher might have a large data set of information from a high school about their former students. Each student can be described using a combination of two factors: gender (male or female) and whether they graduated from high school (yes or no). The researcher wishes to analyze and make decisions about the statistical significance of the main effects and interaction of the factors using a simultaneous combination of interval predictor variables such as GPA, attendance, degree of participation in various extracurricular activities (e.g., band, athletics), weekly amount of screen time, and family income.

Put in a broader statistical context, MANOVA is a special case of canonical correlation and is closely related to discriminant function analysis (DFA). DFA predicts group membership based on multiple interval measures and can be used after a MANOVA to assist in the interpretation of the results.

This entry explains MANOVA by first reviewing the underlying theory of univariate ANOVA and then demonstrating how MANOVA extends ANOVA by using the simplest case of two dependent variables. After the rationale of the analysis is understood, it can be extended to more than two dependent variables but is difficult to present visually. In that case, matrix algebra provides a shorthand method of mathematically presenting the analysis.

Univariate ANOVA

In univariate ANOVA, the independent variable is some combination of group membership and a single interval-dependent variable. The data can be visualized as separate histograms for each group, as seen in Figure 1, with four groups of 20 observations each.

Figure 1 Histogram of four groups

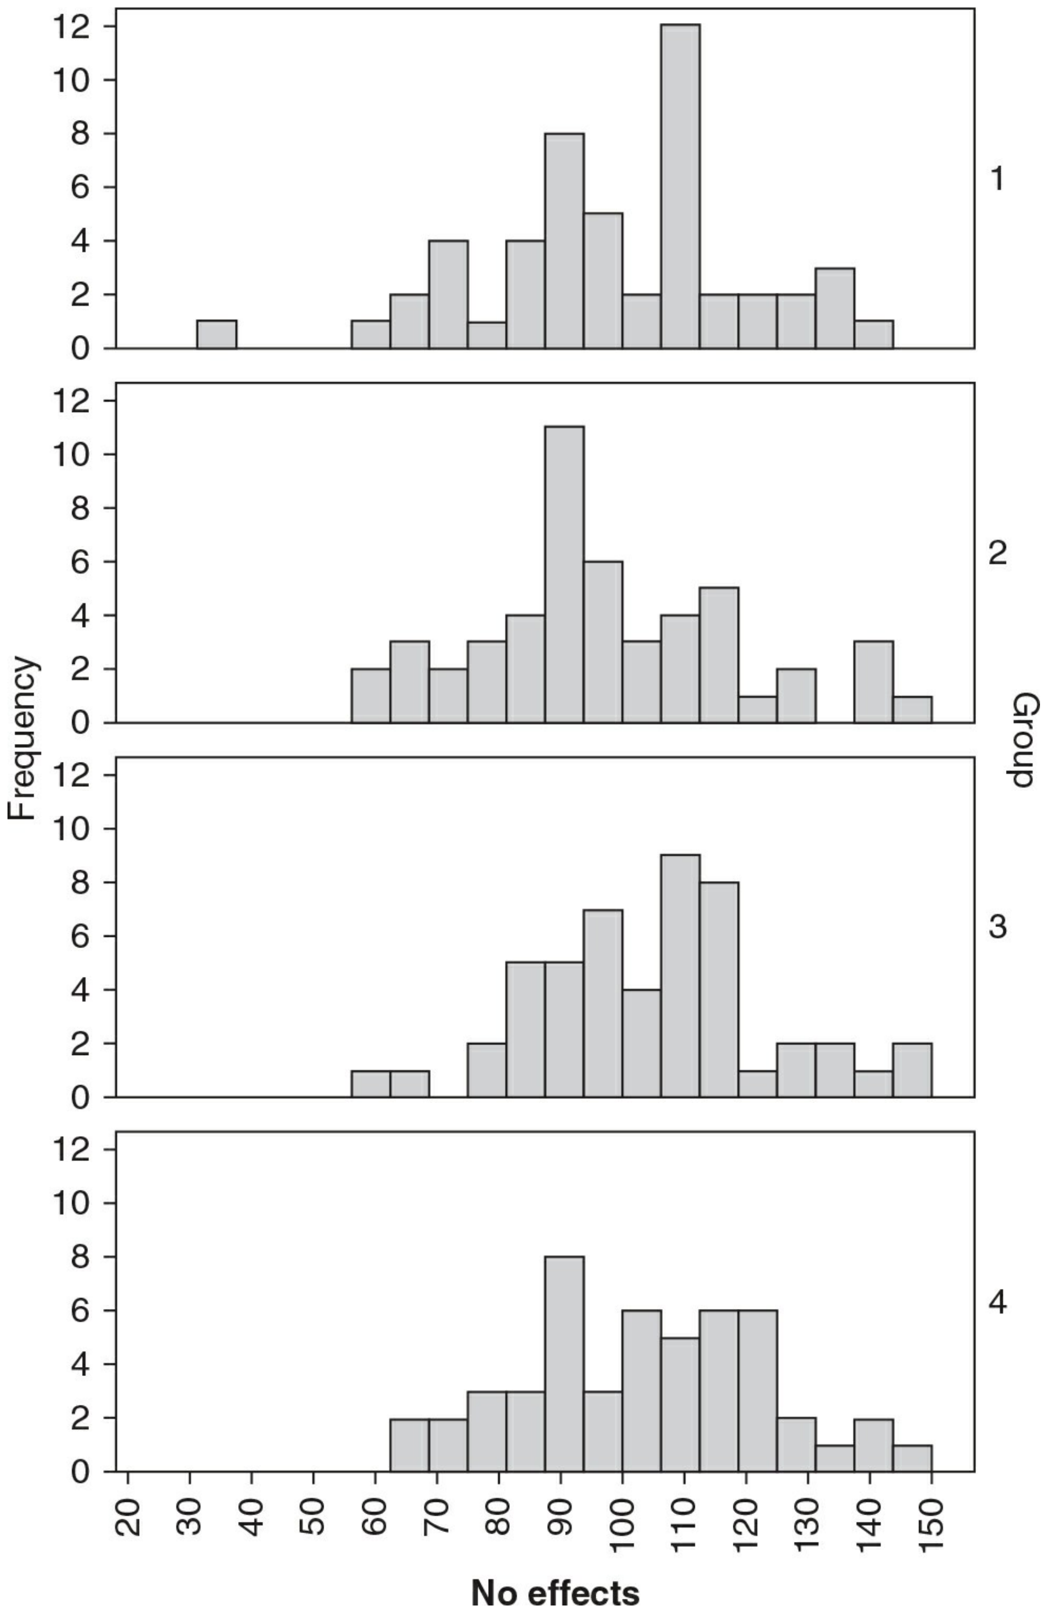

The ratio of the variability between the means of the groups relative to the variability within groups is fundamental to ANOVA. This is done by modeling the sampling distribution of each group with a normal curve model, assuming that both the separate sample means estimate μ and σ is equal in all groups and estimated by a formula using a weighted mean of the sample variances. The assumption of identical within-group variability is called the *homogeneity of* *variance* assumption. The model of the previous data is illustrated in Figure 2.

Figure 2 Normal curve model of four groups

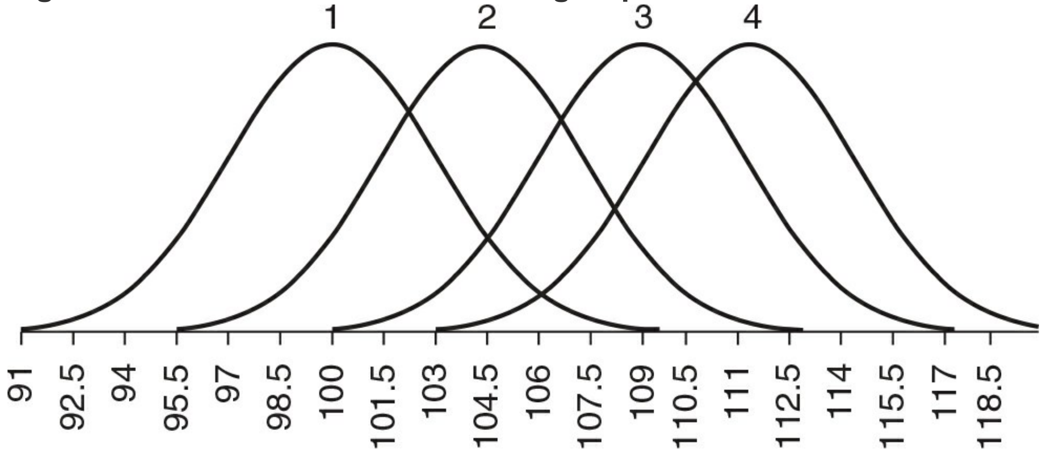

From this model, two estimates of σ^2 are computed. The first, mean square between (MS_B), uses the variability of the means, and the second, mean square within (MS_W), uses the estimate of combined variability within the groups. A computed statistic, called *F*, is the ratio of the two variance estimates:

$$F = \mathrm{MS_B}/\mathrm{MS_W}.$$

The distribution of the *F* statistic is known, given the assumptions of the model are correct. If the computed *F* ratio is large relative to what would be expected by chance, then real effects can be inferred; that is, the means of the groups are significantly different from each other. The between variability can be partitioned using contrasts to account for the structure of group membership, with separate main effects, interactions, and nested main effects, among others, being tested using the ANOVA procedure.

MANOVA

MANOVA and ANOVA have similar independent variables, but in MANOVA there are two or more dependent variables. Although the computations involved in MANOVA are much more complicated and best understood using matrix operations, the basic concept is similar to the univariate case. This will be illustrated by first examining one of the simpler cases of MANOVA, with four groups and two dependent variables. The extension to more groups and dependent variables, while not illustrated, can be inferred from this case.

Four Groups and Two Dependent Variables

The data for four groups and two dependent variables can be illustrated using a scatterplot (see Figure 3). The paired means for each group are called *centroids*, and in matrix algebra terminology together they constitute a vector of length equal to the number of groups. Three of the four standard statistics used in hypothesis testing in MANOVA compare the variability of the centroids to the within-group variability. To do this, they model the dependent variables with a multivariate normal distribution. In a multivariate normal distribution, all univariate distributions will be normal, but having univariate normal distributions for all variables does not guarantee a multivariate normal distribution. In addition, all groups are assumed to have similar variance/covariance matrices, which corresponds to the homogeneity of variance assumption in univariate ANOVA. The bivariate normal model of the sampling distribution of data shown in Figure 3 is presented in Figure 4.

Figure 3 Scatterplot of four groups

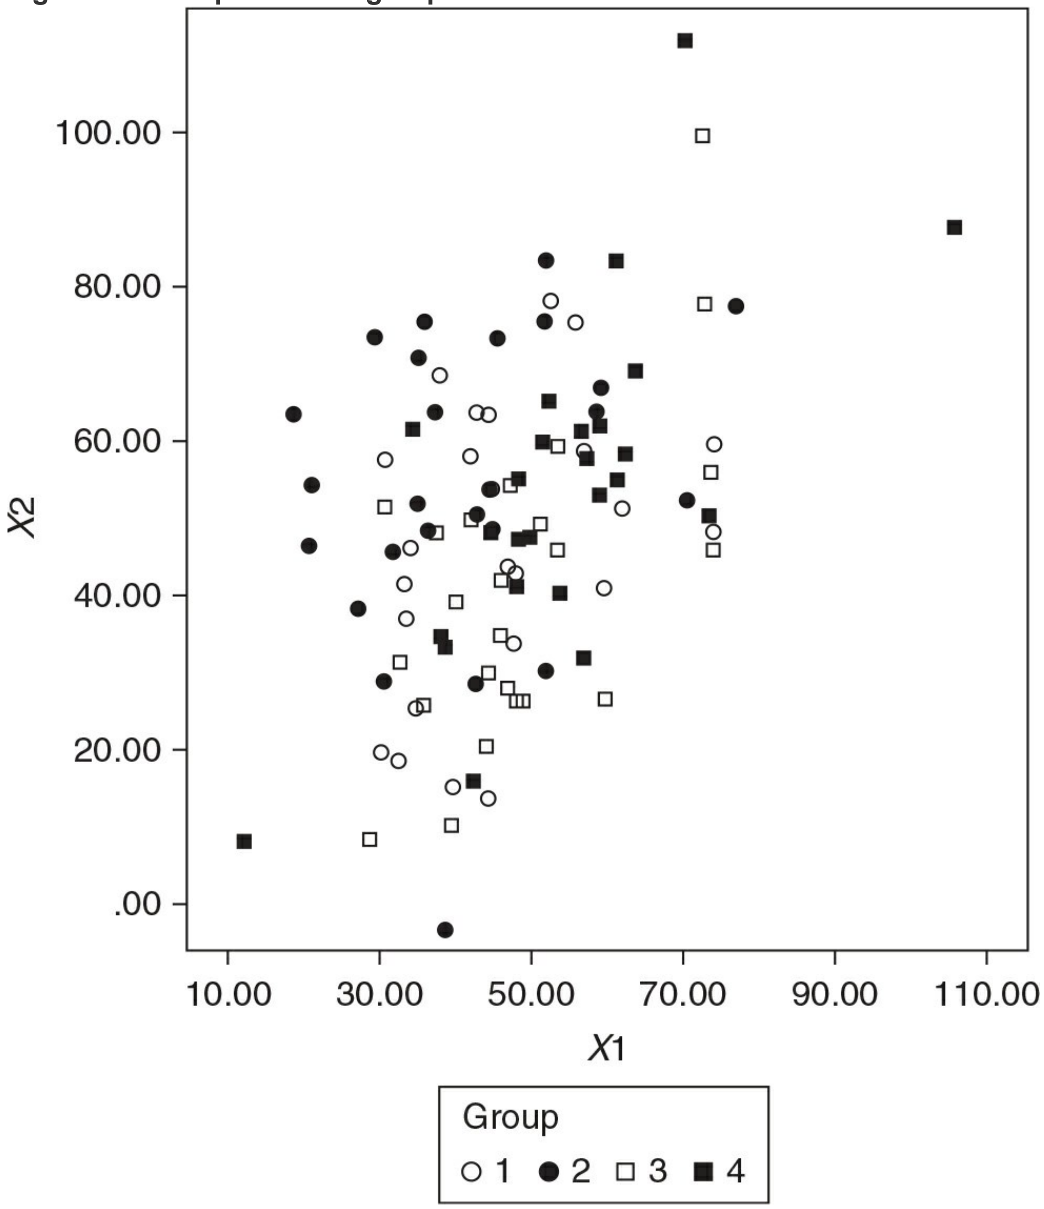

Figure 4 Multivariate normal curve model of four groups

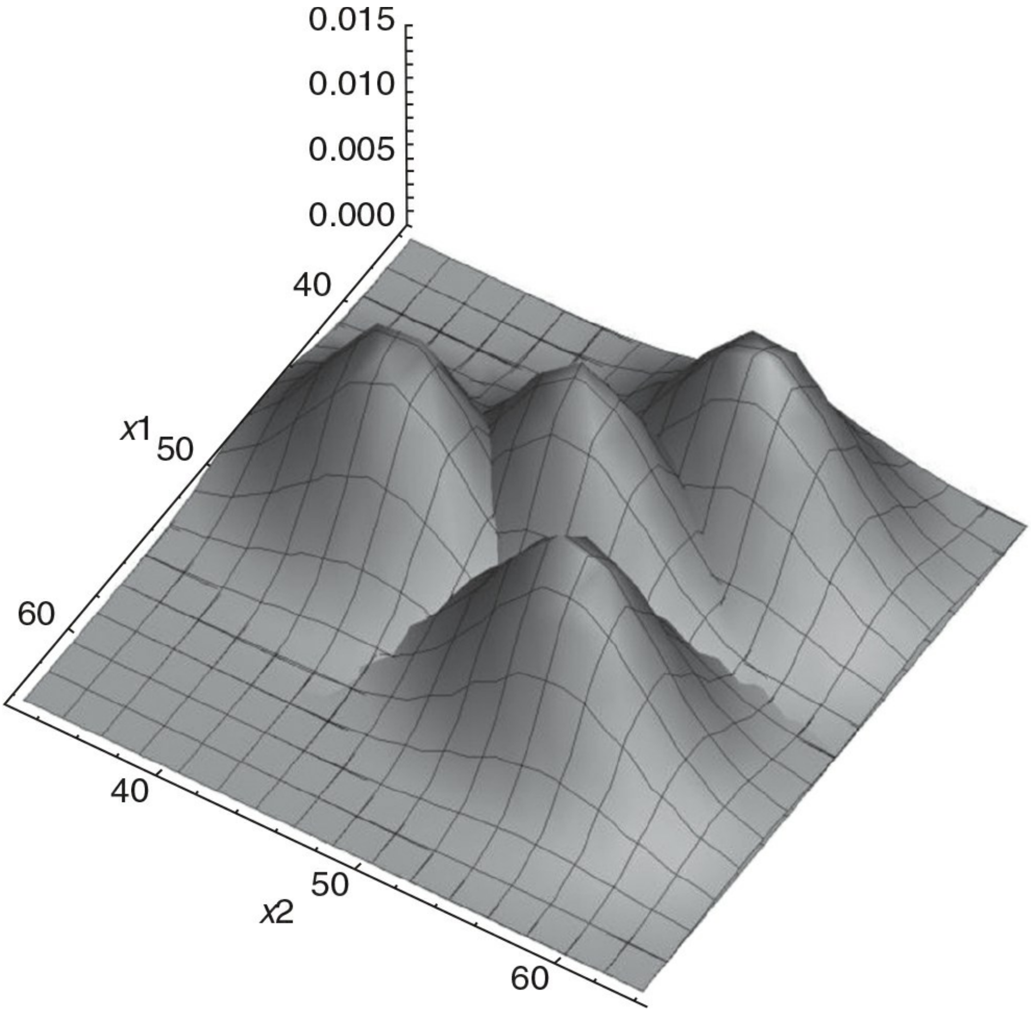

Having data that meet the equal variance/covariance matrix assumption ensures that all individual bivariate normal distributions have the same shape and orientation.

The default SPSS MANOVA output for the example data is shown in Figure 5. The focus of the analysis is on the four "sig" levels of the group effect. Three of the four, Pillai's trace, Wilks's λ, and Hotelling's trace, estimate the ratio of the variability between centroids and the within variability of the separate bivariate normal distributions. They do so in slightly different ways, but given fairly equal and large group Ns, will generate a sig level within a few thousands of each other. The interpretation of these three sig levels is that in combination, the means of dependent measures significantly differentiate between the groups. As in univariate ANOVA, the between variability can be partitioned using contrasts to account for the structure of group membership with separate main effects, interactions, and nested main effects, among others.

Figure 5 MANOVA Output Using SPSS

Multivariate Tests						
Effect		Value	F	Hypothesis df	Error df	Sig.
Group	Pillai's Trace	0.21	3.745	6	192	0.002
	Wilks's Lambda	0.798	3.781[b]	6	190	0.001
	Hotelling's Trace	0.244	3.817	6	188	0.001
	Roy's Largest Root	0.195	6.246[c]	3	96	0.001

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

The fourth default statistic, Roy's largest root, takes a different approach to multivariate hypothesis testing. The data matrix is rotated (transformed using linear transformations) such that the variance between groups is maximized and the variance within groups is minimized. Figure 6 illustrates the rotation of the means in the example data, with the dark solid line showing the rotation. Roy's largest root is computed as a univariate ANOVA on the first extracted root and should be interpreted in light of this transformation. The *F* statistic for Roy's largest root will always be equal to or greater than the largest univariate ANOVA *F* statistic when there are only two dependent variables because if one or more of the dependent measures failed to add any discriminating ability beyond the other dependent measures, the transformation weight for those factors would be zero. Thus, the significance of Roy's largest root will always be equal to or smaller than the smallest of the significance levels. For the example data, the first root was extracted using DFA and saved as a variable to allow comparison with analyses.

Figure 6 Extraction of largest root

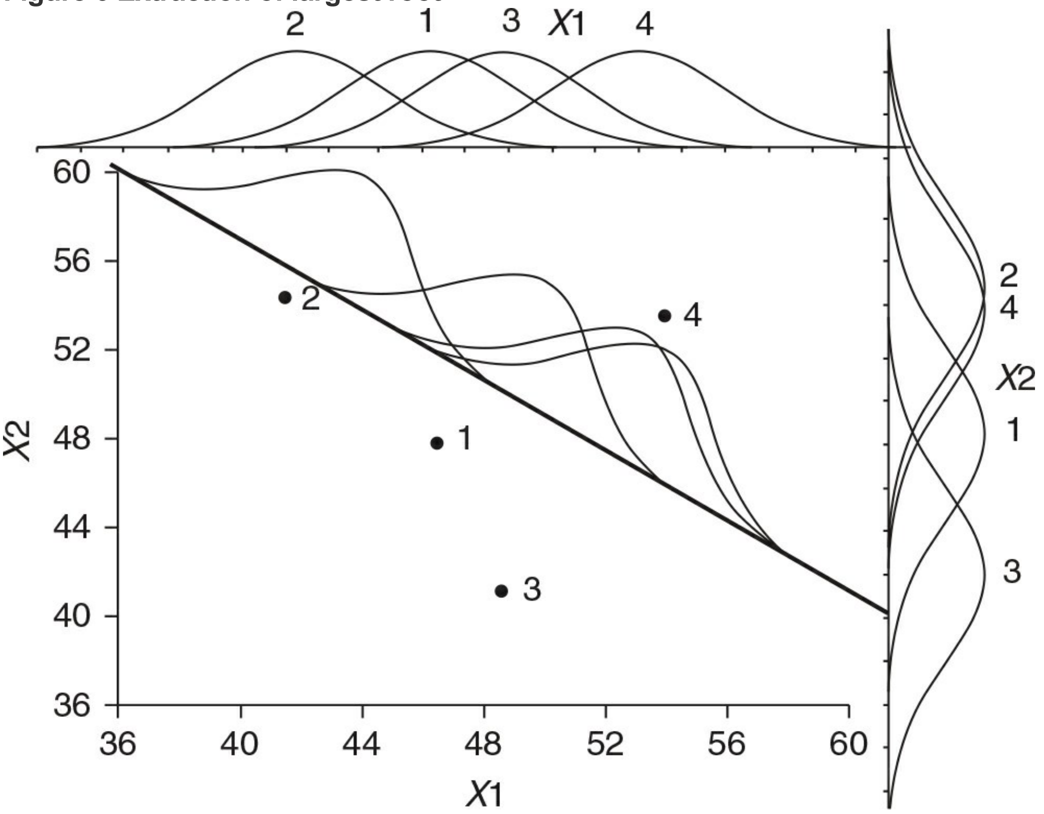

With multivariate dependent measures, another option is to perform a principal component analysis (PCA) on the dependent measures and then do a univariate ANOVA on the first extracted factor, much like Roy's largest root does on the first extracted root in DFA. In PCA, the first orthogonal factor has the greatest variance. This analysis was performed on the example data to compare its results with the others.

In order to interpret the results of MANOVA, univariate ANOVAs are often done to observe how the individual variables contribute to the variability. The results of univariate ANOVAs are presented in Figure 7 for X1, X2, DFA largest root, and the first factor in the PCA.

Figure 7 Univariate ANOVAs

ANOVA Table						
		Sum of Squares	df	Mean Square	F	Sig.
X1	Between Groups	1983.411	3	661.137	3.192	0.027
	Within Groups	19885.616	96	207.142		
	Total	21869.027	99			
X2	Between Groups	2828.229	3	942.743	2.286	0.084
	Within Groups	39587.943	96	412.374		
	Total	42416.172	99			
DFA Principle Root	Between Groups	18.738	3	6.246	6.246	0.001
	Within Groups	96	96	1		
	Total	114.738	99			
PCA	Between Groups	4.731	3	1.577	1.606	0.193
	Within Groups	94.269	96	0.982		
	Total	99	99			

It is interesting to note that the MANOVA statistics all provided a smaller significance level than either of the two dependent measures individually. The univariate ANOVA on the DFA largest root was identical to Roy's largest root result presented in Figure 5. The PCA analysis had the largest significance level and was not statistically significant. The bottom line was that in this case MANOVA appeared to be more powerful than the individual univariate ANOVAs and that PCA did not appear to be a viable alternative.

Power Analysis of MANOVA With Three Groups and Two Dependent Measures

Power estimates for the various MANOVA statistics can be obtained by using simulated data. Figure 8 shows the estimated power of three simulations of 100 observations each and α set at .05. In the first case with a cell size of 10, X1 was generated using a random normal distribution and X2 was set equal to X1, with additional random normal error and small group effects added. That the effects were small relative to the random error can be seen in the low power (.15) observed for the univariate *F* test of the X2 variable. The power for X1 is greater than expected by chance. Pillai's trace, Wilks's λ, and Hotelling's trace all showed a moderate and equal increase in power over the individual univariate power estimates. Roy's largest root showed the greatest power at .45.

Figure 8 Power analysis

	Cell N	Pillai's Trace	Wilks's Lambda	Hotelling's Trace	Roy's Largest Root	Univariate F	
						X1	X2
X1 and X2 correlated with only X2 with effects	19	0.23	0.23	0.23	0.45	0.1	0.15
X1 and X2 correlated with only X2 with effects	100	0.83	0.83	0.83	0.95	0.06	0.55

The second analysis was similar to the first except that cell size was increased to 100. Similar results to the first analysis were found, with all power estimates except for X1 much larger than the case with the smaller cell size. Both of these simulations might be more appropriate for an analysis of covariance in which the variability of the first variable could be factored out before the second variable was analyzed.

The third analysis used a cell size of 50 and uncorrelated X1 and X2 variables, except they were each constructed with similar small effect added. Individually, the variables had power estimates of .38 and .43, respectively, but in combination, Pillai's trace, Wilks's λ, and Hotelling's trace all showed a substantial increase in power. Roy's largest root showed the greatest power at .87. Although this example is hardly a definitive power analysis, it makes a fairly strong argument that performing a MANOVA over multiple univariate ANOVAs results in a fairly significant increase in power.

MANOVA With Three or More Dependent Measures

MANOVA with three or more dependent measures provides a challenge in visualization and interpretation. Basically, the procedure is an extension of the simpler case of two variables but with a greater number of centroids for each group. MANOVA works by comparing the variability of the different centroids to the variability within cells. It requires the assumption of a multivariate normal distribution of the variables with equal variance/covariance matrices for each cell. Violation of these assumptions is likely to lead to a reduction in the power of the analysis.

If statistical significance is found for an effect in MANOVA using Pillai's trace, Wilks's λ, or Hotelling's trace, it means that the centroids of the dependent variables are different for the different levels of the independent variable relative to the within variability. For three dependent variables, it is possible to create a three-dimensional visualization of the centroids and by rotating the vector get a reasonable understanding of the results. Beyond that, interpretation of results becomes problematic. Another caution, as in any multivariate analysis, is that when the measures are highly correlated, collinearity may generate strange results.

If statistical significance is found for an effect in MANOVA using Roy's largest root, univariate ANOVA of the computed principal root can provide an interpretation of the results. In addition, an analysis of the linear transformation that is used to create the principal root can provide additional information, clarifying the results.

In terms of power in MANOVA, it seems reasonable to extend the limited power analysis just presented to the more complicated situation. Generally, that would mean that the power of MANOVA is greater than the individual univariate analyses. If statistical significance is found in a MANOVA, it does not necessarily mean that any of the univariate analyses will be significant. With respect to the increase in power in the case of Roy's largest root, however, all bets are off in that if a DFA reveals more than one significant root, the power of analyzing only the principal root will be reduced.

Because of the difficulty in interpreting a MANOVA, it is recommended to use the technique to develop a deeper understanding of a data set only after a thorough understanding of the simpler, univariate data has been achieved. Rather than starting from the complicated analysis and working backward, start with the simple analysis and use the more complicated analysis to test hypotheses about multivariate relationships within the data.

Limitations

MANOVA provides an extension of univariate ANOVA to simultaneously test for effects over two or more dependent variables. In general, it delivers greater power than multiple univariate tests and its assumptions of similar variance/covariance matrices for all cells are less onerous than the sphericity assumption necessary for repeated-measures ANOVA.

Although it has the advantage of generating output that is similar to ANOVA, difficulty of interpretation is MANOVA's greatest limitation. Statistical significance in MANOVA shows that group means are different for different levels of the independent variable. With two and possibly three dependent measures, visual presentation allows the researcher some tools for analysis, but beyond that, if statistical significance is found, the researcher knows something is going on but is generally unsure of what it is.

Another limitation is the requirement that the dependent variables be a multivariate normal distribution with equal variance/covariance matrices for each cell. MANOVA is fairly robust with respect to this assumption when cell sizes are fairly large and approximately equal otherwise exploration of the reasonableness of this assumption is required.

See alsoAnalysis of Covariance; Analysis of Variance; Canonical Correlation; Discriminant Function Analysis; Normal Distribution; Power; Variance

David W. Stockburger
http://dx.doi.org/10.4135/9781506326139.n456
10.4135/9781506326139.n456

Further Readings

Johnson, R. A., & Wichern, D. W. (1982). Applied multivariate statistical analysis (
3rd ed.
). Upper Saddle River, NJ: Prentice Hall.
Pedhazur, E. J. (1973). Multiple regression in behavioral research explanation and prediction (
3rd ed.
). Fort Worth, TX: Holt, Rinehart and Winston.
van de Geer, J. P. (1971). Introduction to multivariate analysis for the social sciences. San Francisco, CA: W. H. Freeman.