

Multivariate Analysis Using the MIXED Procedure

S. Paul Wright, University of Tennessee, Knoxville, TN

Abstract

The MIXED procedure, already widely used for fitting mixed effects and repeated measures models, is also a valuable tool for multivariate analysis. Capabilities of MIXED which are lacking in standard multivariate procedures include: (1) MIXED uses observations having incomplete responses; (2) MIXED handles non-standard (e.g., multiple design) multivariate models; (3) MIXED handles non-standard (e.g., non-separable) multivariate hypotheses; (4) MIXED allows for structured covariance matrices. This paper presents a series of examples, with emphasis on points (1) and (2) above, to demonstrate multivariate linear model analysis using MIXED. Topics covered include likelihood ratio tests, Wald tests, and model selection using information criteria and their corresponding *R*-squares. Other potential applications, such as principle component analysis, discriminant analysis, canonical correlation, and structural equation modeling will be described briefly. The audience should be familiar with standard multivariate linear model analysis as done by GLM. Some exposure to MIXED would be helpful.

1 Introduction

Since there are plenty of routines available for multivariate analysis, one may wonder why one would even consider using the MIXED procedure for multivariate analysis. Here are several reasons, listed in order of their importance, most important first. (1) Most multivariate procedures discard an entire observation if it contains any missing data; MIXED does not. MIXED uses observations that have incomplete response data. (2) MIXED can analyze non-standard models such as multiple design multivariate models (the seemingly unrelated regressions models of econometrics) in which each response variable can have a different set of explanatory variables. (3) MIXED can test non-standard hypotheses such as multivariate linear model hypotheses that cannot be written in the standard "separable" form of $H_0: \mathbf{LBM}=\mathbf{D}$ (see the *SAS/STAT® User's Guide, Version 6, Fourth Edition, Volume 1*, p. 17). (4) MIXED allows the use of structured covariance matrices (which are especially useful in repeated measures problems but usually less so in multivariate problems) along with a choice of methods for estimating them.

2 The Multivariate Linear Model

The standard multivariate linear model is commonly written as $\mathbf{Y} = \mathbf{XB} + \mathbf{E}$. \mathbf{Y} is an n -by- r matrix of r response variables measured on n subjects; \mathbf{X} is an n -by- p matrix of explanatory variables; \mathbf{B} is a p -by- r matrix of regression coefficients; and \mathbf{E} is an n -by- r "error" matrix whose rows

are iid normal with mean $\mathbf{0}$ and covariance matrix Σ . Here is a simple example with two responses and one explanatory variable (in addition to an intercept term) measured on three subjects.

$$\begin{pmatrix} Y_{1_1} & Y_{2_1} \\ Y_{1_2} & Y_{2_2} \\ Y_{1_3} & Y_{2_3} \end{pmatrix} = \begin{pmatrix} 1 & X_1 \\ 1 & X_2 \\ 1 & X_3 \end{pmatrix} \begin{pmatrix} \beta_{01} & \beta_{02} \\ \beta_{11} & \beta_{12} \end{pmatrix} + \begin{pmatrix} \epsilon_{1_1} & \epsilon_{2_1} \\ \epsilon_{1_2} & \epsilon_{2_2} \\ \epsilon_{1_3} & \epsilon_{2_3} \end{pmatrix}.$$

I refer to this standard format as the "wide" format. To use MIXED, \mathbf{Y} , \mathbf{B} and \mathbf{E} are rewritten as vectors and \mathbf{X} is rearranged accordingly, resulting in what I call the "tall" format.

$$\begin{pmatrix} Y_{1_1} \\ Y_{2_1} \\ Y_{1_2} \\ Y_{2_2} \\ Y_{1_3} \\ Y_{2_3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & X_1 & 0 \\ 0 & 1 & 0 & X_1 \\ 1 & 0 & X_2 & 0 \\ 0 & 1 & 0 & X_2 \\ 1 & 0 & X_3 & 0 \\ 0 & 1 & 0 & X_3 \end{pmatrix} \begin{pmatrix} \beta_{01} \\ \beta_{02} \\ \beta_{11} \\ \beta_{12} \end{pmatrix} + \begin{pmatrix} \epsilon_{1_1} \\ \epsilon_{2_1} \\ \epsilon_{1_2} \\ \epsilon_{2_2} \\ \epsilon_{1_3} \\ \epsilon_{2_3} \end{pmatrix},$$

or in matrix notation, $\tilde{\mathbf{y}} = \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} + \tilde{\boldsymbol{\epsilon}}$ where $\tilde{\mathbf{y}} = \text{vec}(\mathbf{Y}')$, $\tilde{\boldsymbol{\beta}} = \text{vec}(\mathbf{B}')$, $\tilde{\boldsymbol{\epsilon}} = \text{vec}(\mathbf{E}')$, and $\tilde{\mathbf{X}} = \mathbf{X} \otimes \mathbf{I}$. $\tilde{\mathbf{X}}$ is an N -by- P matrix with $N=n^*r$ and $P=p^*r$. Furthermore, the multivariate linear hypothesis $\mathbf{LBM}=\mathbf{D}$ becomes, in "tall" format, $\mathbf{C}\tilde{\boldsymbol{\beta}}=\mathbf{d}$ where $\mathbf{d} = \text{vec}(\mathbf{D}')$ and $\mathbf{C} = \mathbf{L} \otimes \mathbf{M}'$. Note that standard multivariate tests *must* be in the "separable" $\mathbf{L} \otimes \mathbf{M}'$ form, but the \mathbf{C} matrix in MIXED is not limited to such hypotheses.

3 An Example

Brown and Beerstecher (1951) reported results of chemical analyses of 45 urine samples (identified with values 1 to 45 in a variable named **ID** in the programs below) from young adult males classified as either underweight or overweight. In analyzing this data set, Smith, et al. (1962), Seber (1984), and Morrison (1992) used four weight groups (variable **G** with values 1, 2, 3, 4): lighter underweight ($n_1=12$), heavier underweight ($n_2=14$), lighter overweight ($n_3=11$) and heavier overweight ($n_4=8$). There were eleven response variables of which only three will be used here: **Y1** = pH, **Y3** = pigment creatinine, **Y4** = phosphate (mg/ml). There were also two covariates: **X1** = volume of sample in ml, and **X2** = (specific gravity - 1) times 1000. This example will be used to demonstrate multivariate analysis of variance (MAOV), multivariate analysis of covariance (MACV), and multivariate regression.

4 MAOV using GLM and MIXED

The following SAS code reads the data in its original "wide" format, does a standard MAOV using GLM, rearranges the data to "tall" format and does the analysis using MIXED. This example shows that the basic "trick" to fitting multivariate models in MIXED is to have a CLASS variable (named VAR in this paper) to identify the response. The VAR variable, in this example, generates three design matrix columns corresponding to three intercept terms, one for each response. Therefore, the NOINT option is used on the MODEL statement to prevent MIXED from generating another, unnecessary intercept column. In general, VAR is crossed with each other effect in the model to create the \tilde{X} matrix as the examples below show. The SOLUTION or S option on the MODEL statement causes the estimated regression coefficients to be printed. The REPEATED statement specifies an unstructured covariance matrix among the three responses. Other aspects of MIXED's input and output are discussed below.

```
data wide;
  infile "urine.dat";
  id+1;
  input G y1-y11 x1 x2;
  keep id G y1 y3 y4 x1 x2;
proc glm data=wide;
  class G;
  model y1 y3 y4 = G / nouni;
  manova h=G;
data tall;
  set wide;
  var='Y1'; y=y1; output;
  var='Y3'; y=y3; output;
  var='Y4'; y=y4; output;
  keep id G var y x1 x2;
*** #1: MAOV model using ML ***;
proc mixed data=tall method=ML IC;
  class var G;
  model y= var*G / noint notest s;
  repeated var / subject=id type=un;
*** #2: Null model using ML ***;
proc mixed data=tall method=ML IC;
  class var;
  model y= var / noint notest s;
  repeated var / subject=id type=un;
*** #3: MAOV: REML, CONTRAST ***;
proc mixed data=tall method=REML;
  class var G;
  model y= var*G / noint notest s;
  repeated var / subject=id type=un;
  contrast "GROUP"
    var*G 1 0 0 -1,
    var*G 0 1 0 -1,
    var*G 0 0 1 -1,
    var*G 0 0 0 0 1 0 0 -1,
    var*G 0 0 0 0 0 1 0 -1,
    var*G 0 0 0 0 0 0 1 -1,
    var*G 0 0 0 0 0 0 0 1 0 0 -1,
    var*G 0 0 0 0 0 0 0 0 1 0 -1,
    var*G 0 0 0 0 0 0 0 0 0 1 -1 /
    chisq;
```

```
make 'contrast' out=showit;
proc print data=showit; run;
```

Here are selected portions from the output of the first MIXED procedure above. For more complete samples of MIXED's input and output, see Littell, et al. (1996) or see *SAS/STAT Software: Changes and Enhancements through Release 6.12*.

Model Fitting Information for Y

Observations	135.0000
Log Likelihood	-173.072
-2 Log Likelihood	346.1446

Information Criteria

Better	Parms	...	AIC	...	BIC
Smaller	18	...	382.1	...	434.4

Solution for Fixed Effects

Effect	VAR	G	Estimate	...
VAR*G	Y1	1	5.666667	...
VAR*G	Y1	2	5.450000	...
VAR*G	Y1	3	5.500000	...
VAR*G	Y1	4	5.325000	...
VAR*G	Y3	1	15.891667	...
VAR*G	Y3	2	17.821429	...
VAR*G	Y3	3	16.354545	...
VAR*G	Y3	4	11.912500	...
VAR*G	Y4	1	1.787500	...
VAR*G	Y4	2	1.564286	...
VAR*G	Y4	3	1.881818	...
VAR*G	Y4	4	3.231250	...

CONTRAST Statement Results

Source	NDF	DDF	ChiSq	F	...
GROUP	9	45	32.42	3.60	...

To save space the GLM output is not shown, but for future reference, the multivariate tests of the GROUP (G) effect in GLM include Wilks' Lambda ($\Lambda=0.52588539$, $F=3.1923$ with 9 and 95.07 degrees of freedom, p -value=.00204) and the Hotelling-Lawley trace ($U=0.79084991$, $F=3.3099$ with 9 and 113 d.f., p -value=.00128). These two tests are also available from MIXED, but not directly. At present, multivariate tests for multivariate hypotheses are not available in MIXED. (Multivariate tests for within-subjects effects in repeated measures models are available via HLM and HLPS options on the REPEATED statement.) The NOTEST option on the MODEL statement suppresses the usual tests which might be of interest in a repeated measures model but not in a multivariate model.

The estimates reported for this model ("Solution for Fixed Effects") are simply the cell means. To obtain a likelihood ratio test for GROUP, it is necessary to fit two models by maximum likelihood (METHOD=ML on the PROC statement): the restricted model ("Null" model in this case) and the unrestricted model ("MAOV" model). Let M

= -2 times the maximized log-likelihood with subscripts R and U for restricted and unrestricted models. Then $M_U=346.1446$ (shown above under "Model Fitting Information") and $M_R=375.0649$ (from the output of Model #2, not shown). The likelihood ratio chi-squared statistic is $M_R - M_U = 28.9203$ with 9 degrees of freedom since the MAOV model has 18 parameters (12 cell means plus 6 variances and covariances) while the null model has only 9 parameters (3 means plus 6 (co)variances). Wilks' Lambda is $\Lambda = \exp[-(M_R - M_U)/n] = 0.525885$ (the same as given by GLM) where $n=45$ is the number of subjects. GLM reports Rao's F approximation for Λ which is generally preferred to the chi-squared approximation. Its formula is given in the *SAS/STAT User's Guide*, page 18.

The CONTRAST statement in MIXED Model #3 above produces (by default) an F statistic which, in this case, has the value $F=3.60$ with 9 and 45 degrees of freedom and p -value=.0019. The CHISQ option on the CONTRAST statement requests printing of the corresponding chi-squared statistic which is simply the F multiplied by its numerator degrees of freedom of 9, i.e., chi-squared=32.42. This chi-squared is the Wald test statistic for the Group effect. The MAKE statement was used to store the CONTRAST results and print them in order to display more decimal places. Kleinbaum (1973) showed that U , the Hotelling-Lawley trace statistic, is simply the Wald chi-squared divided the error degrees of freedom, $n-p$. (Recall that $p=\text{rank}(X)$ or, in this example, the number of groups, namely 4.) In the example, $U=32.424846 / 41 = .79085$, the same as given by GLM. The F test reported by MIXED is *not* the same as that reported by GLM for the Hotelling-Lawley trace though it is asymptotically equivalent. In small samples it is generally better to use the Pillai-Sampson (1959) approximation used by GLM or, even better, the McKeon (1974) F approximation. Formulas are given in the *SAS/STAT User's Guide*, page 19 and in Seber (1984).

5 The MAOV with Missing Data

To dramatize the difference between GLM and MIXED in handling missing data, in each of the 45 multivariate observations, one of the three response values was randomly set to missing. Thus the total number of response values was reduced from $N=135$ to $N=90$. GLM cannot analyze the data *at all* since there are now *zero* complete observations! The MIXED procedures run *exactly as before* with no change in the SAS code. The likelihood ratio and Wald chi-squared statistics may be calculated as before. And with small samples, the Rao and McKeon F approximations almost certainly give more accurate results than the chi-squared tests, but their distributions in the case of missing data need more investigation. These F approximations require a value for " n ". Using $n=45$ (the total number of subjects) in the formulas may tend to make the tests too "liberal" since the formulas assume complete data ($N = n*r = 135$). A reasonable approach is to use $n' = N/r = 30$. This approach has been shown to work well in repeated measures problems (Wright and Wolfinger, 1996). For the current example with $n'=30$, $M_U=219.7317$ and $M_R=244.6180$ for a likelihood ratio chi-squared of 24.8863

on 9 degrees of freedom. $\Lambda=0.4362485$, and Rao's $F=2.6427$ with 9 and 58.56 d.f. and p -value=.01221. The Wald chi-squared is 26.037854. McKeon's $F=2.5979$ with 9 and 34.67 d.f. and p -value=.02088.

6 Additional Models

A convenient way to choose the "best" model out of a large set of models is to use a model selection criterion rather than to rely on hypothesis tests. A number of popular criteria are of the general form

$$GIC = M + d * C$$

where GIC stands for Generalized Information Criterion. Recall that M is -2 times the maximized log-likelihood. " d " is the number of estimated parameters; " C " is the "cost", the penalty paid for each parameter estimated. Smaller values of GIC indicate better models. The IC option on the PROC MIXED statement produces a table with four varieties of GIC . One line of such a table, with two of the four criteria, is shown in the output above. The PARMS column shows the value of " d ". In the example shown, there are 18 parameters, 12 regression coefficients (cell means) plus 6 elements in the covariance matrix (not shown). $C=2$ for Akaike's criterion (AIC) and $C=\log N$ for Schwarz's Bayesian information criterion (BIC) where " \log " is the natural logarithm. For multivariate models, a better choice for BIC would probably be *to use* $C=\log n'$ (Kass & Raftery, 1995); this version of GIC is labeled " n' - BIC " in the tables below. These values for the Null and MAOV models, and other models discussed below, are given in Table 1. These models are for the complete data ($N=135$), but a similar table could be constructed for the incomplete data ($N=90$).

Table 1. Model Selection Criteria.

Model	d	M	AIC	n'-BIC	BIC
Null	9	375.1	393.1	409.3	419.2
MAOV	18	346.1	382.1	414.6	434.4
MACV	24	305.9	353.9	397.2	423.6
MACV2	18	325.3	361.3	393.8	413.6
MREG	18	314.7	350.7	383.2	403.0
MDM1	17	318.8	352.8	383.5	402.2
MDM2	16	318.9	350.9	379.8	397.4
MDM3	14	320.7	348.7	374.0	389.4

An alternate way to express these model selection criteria is in terms of "percent improvement" or "proportional reduction in error" relative to a base model which is usually the Null model. The result is a multivariate analog of the univariate R -square. A general formula is

$$R^2_{GIC} = 1 - \exp[(GIC - GIC_0) / N]$$

where GIC_0 is the GIC value for the base model and N is the total number of response values ($N=n*r$ if there is no missing data). In the univariate case ($r=1$) with $C=0$, this is just the ordinary R -square (Magee, 1990). In the multivariate case with $C=0$, this is (usually) $\hat{\gamma}_3$ of Cramer and Nicewander (1979), labeled A_λ in Rencher (1995).

When $C > 0$, it becomes an "adjusted" R -square for which larger values represent better models. A possible advantage to using the R -square representation is that one may discover that the "best" model is none too good. Table 2 shows the R -square values for the models in Table 1. A negative R -square simply means that a model is worse than the Null model. The models themselves are described below.

Table 2. Model Selection R-Squares.

Model	d	C=0	AIC	n'-BIC	BIC
Null	9	.0	.0	.0	.0
MAOV	18	.1928	.0777	-.0403	-.1194
MACV	24	.4010	.2520	.0857	-.0330
MACV2	18	.3082	.2096	.1084	.0407
MREG	18	.3608	.2696	.1761	.1135
MDM1	17	.3408	.2579	.1740	.1185
MDM2	16	.3402	.2681	.1962	.1491
MDM3	14	.3314	.2800	.2306	.1982

The Null and MAOV models were given above. MACV is the full, parallel slopes, multivariate analysis of covariance model. The CLASS and MODEL statements are

```
class var G;
model y= var*G var*X1 var*X2 /
      noint notest s;
```

All other MIXED statements are unchanged. This model has 24 parameters: 12 intercept terms (for 3 responses in 4 groups), 3 slopes for X1 (one for each response), 3 slopes for X2, and 6 covariance matrix elements. Including the two covariates in the model improves it considerably (relative to MAOV) according to AIC (with $C=2$), n' - BIC ($C=\log[45]=3.81$), and BIC ($C=\log[135]=4.91$; but according to BIC , the MACV model is still worse than the Null model.

In the original analysis of Brown and Beerstecher (1951) there were only two groups, underweight ($G=1, 2$) and overweight ($G=3, 4$). The MACV2 row in Table 1 and Table 2 shows the results for this model. This is the best model so far according to n' - BIC and BIC but not AIC . AIC is known for its tendency to favor models with more parameters. The CLASS and MODEL statements for MACV2 are

```
class var G34;
model y= var*G34 var*X1 var*X2 /
      noint notest s;
```

The variable G34 was created with the DATA step statement

```
G34=(G in (3,4));
```

Since the four groups represent four levels of a quantitative variable ("obesity"), another possibility is to treat "Group" as quantitative. To avoid confusion (and to permit the use of "Group" as both categorical and quantitative in subsequent models), a new variable named "XG" was created (in a DATA step) having the same values

as G (namely 1, 2, 3, 4). The analysis is now a multivariate regression. The CLASS and MODEL statements are

```
class var;
model y= var var*XG var*X1 var*X2 /
      noint notest s;
```

This model fits three intercepts (one for each level of VAR), three slopes for XG, three for X1 and three for X2 (plus the six elements of the covariance matrix). Results are labeled MREG in Tables 1 and 2. This model is the best so far, by any of the three criteria.

The remaining models all attempt to improve the fit by allowing each response variable to have a different set of explanatory variables. These are "multiple design" multivariate models ("seemingly unrelated regressions" models in econometrics); they are labeled "MDM" in Tables 1 and 2. To specify these models in MIXED, the VAR variable is replaced by individual 0-1 dummy variables, one for each response variable. The variable names are V1, V3 and V4. Below is the SAS code for the MDM1 model, including the DATA step to create the dummy variables. In this model, response Y1 is modeled as unrelated to the explanatory variables and so has only an intercept term in the model. Response Y3 has a different intercept in each of the four groups but the same slopes (for X1 and X2) in all groups; i.e., Y3 has a standard, parallel slopes, analysis of covariance model. Y4 has a multiple regression model with an intercept and three slopes (for XG, X1, and X2).

```
data tall; set tall;
V1=(var='Y1');
V3=(var='Y3');
V4=(var='Y4');
proc mixed data=tall method=ML IC;
class G;
model y= V1 V3*G V3*X1 V3*X2
        V4 V4*XG V4*X1 V4*X2 /
      noint notest s;
```

Model MDM1 is slightly better or slightly worse than MREG, depending on which criterion is used. In MDM1, the coefficient for $V3*X2$ (not shown) is essentially zero. Dropping this term from the model gives model MDM2.

```
model y= V1 V3*G V3*X1
        V4 V4*XG V4*X1 V4*X2 /
      noint notest s;
```

Finally, the four Y3 intercepts fall into two groups of roughly equal value, one for $G = 1$ or 2 , another for $G = 3$ or 4 . This suggests replacing $V3*G$ with $V3*G34$, giving model MDM3.

```
model y= V1 V3*G34 V3*X1
        V4 V4*XG V4*X1 V4*X2 /
      noint notest s;
```

Model MDM3 is best according to all three criteria. Here are the estimated regression coefficients ("Solution for Fixed Effects") for MDM3.

Effect	G34	Estimate	Std Error	...
V1		5.497778	0.043828	...
V3*G34	0	11.925395	1.213271	...
V3*G34	1	9.063030	1.341848	...
V3*X1		0.019679	0.004187	...
V4		0.946152	0.636977	...
V4*XG		0.514166	0.085652	...
V4*X1		-0.005331	0.000988	...
V4*X2		0.048612	0.017455	...

As an aid to interpretation, these results can be written as four different univariate models: [1] for Y1; [2] for Y3 in Groups 1 and 2; [3] for Y3 in Groups 3 and 4; and [4] for Y4. These four models are:

- [1] $Y1 = 5.498$,
- [2] $Y3 = 11.925 + 0.020 \cdot X1$,
- [3] $Y3 = 9.063 + 0.020 \cdot X1$,
- [4] $Y4 = 0.946 + 0.514 \cdot XG - 0.005 \cdot X1 + 0.049 \cdot X2$.

Many other models could be tried, including analysis of covariance models with non-parallel slopes and models with unequal covariance matrices in different groups (using the GROUP option on the REPEATED statement). It is hoped that the examples given here are sufficient to indicate the general approach which can be adapted to many different situations.

7 Other Multivariate Procedures

Space does not permit demonstration of how MIXED can be used to enhance other multivariate procedures, but a few "hints" may be helpful. Many multivariate procedures in SAS will accept as input a covariance matrix (a TYPE=COV data set) as an alternative to raw data. Some examples are PRINCOMP, FACTOR, CANCORR, CANDISC, DISCRIM, and CALIS. All these procedures, if given raw data, will omit any observation that contains a missing value. Instead, MIXED can be used to estimate the covariance matrix using incomplete as well as complete observations. This matrix can be put into TYPE=COV form using a DATA step or the IML procedure. Where a value for the number of multivariate observations is needed, n' can be used. Many of these procedures can also produce "scores" (factor scores, canonical variates, etc.) which are useful for plotting and other purposes. For this, raw data is indispensable. One possibility worth exploring is to use the PREDICTED option of the MODEL statement in MIXED (version 6.12) to obtain empirical BLUP estimates of the missing response values for use in calculating "scores" (perhaps using the SCORE procedure). Another instance in which raw data is indispensable is in classifying observations in discriminant analysis. The DISCRIM procedure rejects any incomplete observations, but the formulas DISCRIM uses for classification (see the SAS/STAT User's Guide, page 680) apply just as well to incomplete observations, without requiring that estimates be inserted for missing values, by choosing the submatrix of Σ that corresponds to the nonmissing data. MIXED could be used to estimate Σ , and IML could be used to implement the classification formulas.

References

Brown, J. D. & Beerstecher, E. (1951). Metabolic Patterns of Underweight and Overweight Individuals. University of Texas Publication, Biochemical Institute Studies IV,

Cramer, E. M. & Nicewander, W. A. (1979). Some Symmetric, Invariant Measures of Multivariate Association. *Psychometrika* **44**, 43-54.

Kass, R. E. & Raftery, A. E. (1995). Bayes Factors. *JASA* **90**, 773-795.

Kleinbaum, D. G. (1973). Testing Linear Hypotheses in Generalized Multivariate Models. *Communications in Statistics* **1**, 433-457.

Littell, R. C., Milligan, G. A., Stroup, W. W. & Wolfinger, R. D. (1996). *SAS System for Mixed Models*. Cary, NC: SAS Institute Inc.

Magee, L. (1990). R^2 Measures Based on Wald and Likelihood Ratio Joint Significance Tests. *The American Statistician* **44**, 250-253.

McKeon, J. J. (1974). F Approximations to the Distribution of Hotelling's T^2 . *Biometrika* **61**, 381-383.

Morrison, D. F. (1992). *Multivariate Statistical Methods*. New York: McGraw-Hill.

Pillai, K. C. S. & Sampson, P. (1959). On Hotelling's Generalization of T^2 . *Biometrika* **46**, 160-168.

Rencher, A. C. (1995). *Methods of Multivariate Analysis*. New York: John Wiley & Sons.

SAS Institute Inc. (1990). *SAS/STAT User's Guide, Version 6, Fourth Edition, Volume 1*. Cary, NC: SAS Institute Inc.

SAS Institute Inc. (1997). *SAS/STAT Software: Changes and Enhancements through Release 6.12*. Cary, NC: SAS Institute Inc.

Seber, G. A. F. (1984). *Multivariate Observations*. New York: John Wiley & Sons.

Smith, H., Gnanadesikan, R. & Hughes, J. B. (1962). Multivariate Analysis of Variance (MANOVA). *Biometrics* **18**, 22-41.

Wright, S. P. & Wolfinger, R. D. (1996). Repeated Measures Analysis using MIXED Models: Some Simulation Results. In Gregoire, T. G., Brillinger, D. R., Diggle, P. J., Russek-Cohen, E., Warren, W. G., and Wolfinger, R. D, *Modelling Longitudinal and Spatially Correlated Data*. New York: Springer-Verlag.

SAS and SAS/STAT are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. © indicates USA registration.

Author Information:

S. Paul Wright
 328 SMC - Statistics Department
 University of Tennessee
 Knoxville, TN 37995-0532
 email: wrightsp@utk.edu