

Mixed-Up Mixed Models: Things That Look Like They Should Work But Don't, and Things That Look Like They Shouldn't Work But Do

Robert M. Hamer, Ph.D., UMDNJ Robert Wood Johnson Medical School, Piscataway NJ
P.M. Simpson, University of Arkansas for Medical Sciences, Little Rock, AR

Introduction

In recent years, the use of mixed models in fitting data in the biomedical sciences, social sciences, economics, and business, has become more widespread. Part of the increase in the use of such models, apart from their inherent utility, is because software to fit such models has become increasingly available. SAS Institute's contribution to the mixed model software is PROC MIXED. However, because mixed models are more complex and more flexible than the general linear model, the potential for confusion and errors is higher. This paper outlines some confusion that may occur when data analysts experienced in the use PROC GLM to analyze data for both the fixed-effects and mixed-effects models use PROC MIXED to analyze data. Littell, Milliken, Stroup, and Wolfinger (1996) is a very good reference on mixed models in the context of using PROC MIXED, and Milliken and Johnson (1992, 1989, and in press) are good general references on experimental design, including mixed models.

Mixed Models

In this section, we define both the general linear model and the mixed general linear model (called hereafter the mixed model). We then give several examples, and demonstrate the use of both PROC GLM and PROC MIXED in the analysis of such data.

The General Linear Model consists of

$$Y = X\beta + \epsilon$$

Where Y is the n -by- 1 matrix of response variables, X is the n -by- k design or model matrix, β is the k -by- 1 matrix of parameters, and ϵ is the n -by- 1 error matrix. We assume, for consistency with the mixed model, that p , the number of columns in Y and β equals 1; that is, this is a univariate general linear model.

In this model, all effects, which are represented by columns of X , are assumed to be *fixed effects*; that is, we assume that the levels of the factors represent either levels we manipulate, or are measured without error, and are the only levels in which we are interested. This is the model that PROC GLM fits. When we use PROC GLM,

the CLASS and MODEL statements jointly specify the columns of Y and X , which imply the contents of β and ϵ .

Effects or factors for which we cannot make these assumptions are called *random effects*. Randomized blocks designs are typically mixed models, because the blocks are typically considered to have been a random sample of all possible blocks which we might have considered, and thus we are interested in generalizing beyond these particular blocks. Repeated measures designs are typically mixed models because usually we have at least one within-subjects or repeated factor, and we wish to generalize beyond the particular subjects we have used.

When we fit a mixed model with PROC GLM, we are taking advantage of the fact that for parameter estimation, for many models, the parameter estimates obtained when a factor is fixed have the same values as those obtained when a factor is random. However, it is frequently the case that the standard errors of those estimates are wrong (which we cannot fix in PROC GLM), and that hypothesis tests involving these factors often require that their F-statistics use something other than the mean square for error in the denominator. PROC GLM has facilities for controlling the choice of denominator in an F-statistic.

The Mixed Model consists of

$$Y = X\beta + Z\gamma + \epsilon$$

where y , X , β , and ϵ have the same meanings as in the general linear model, but for the fixed effects only (and therefore have different dimensions), Z is an n -by- m design or model matrix for the random effects, and γ is an m -by- 1 matrix of parameters for the random effects (although the meaning of the parameters for the random effects differs slightly from the meaning of parameters for fixed effects).

There are two things to remember here:

1. For a given mixed design (e.g., a design with one between-subjects factor and one within-subjects factor) the columns that PROC MIXED places in X and Z are exactly the same columns that PROC GLM would place in (fixed effects) X . It is just that in PROC MIXED these columns are distributed between two matrices, and hence, parameter estimates are

distributed correspondingly between two matrices. (There is actually a small technical exception to this statement, but we will pretend here that this exception does not exist.)

- If you use PROC MIXED to fit a fixed effects design of the type that is appropriate to fit with PROC GLM, you should get the same results, within rounding error, if you have balanced data.

Choice of Procedure: GLM or MIXED

Why would we be interested in using PROC MIXED for mixed effect models, rather than using PROC GLM? When we have random factors in a model, the following occurs when we use PROC GLM to analyze the data:

- GLM may frequently get the standard errors of the parameter estimates wrong, as a consequence of not really knowing that a factor is random, rather than fixed. This in turn means that some of the standard errors of the least squares means may be wrong. This may have implications for hypothesis testing, for estimation of differences between least squares means, and placing confidence intervals on estimates of least squares means. PROC MIXED gets the standard errors right.
- The estimation process that PROC GLM uses is least squares. We have to calculate variance components for random effects by hand, using estimates produced by PROC GLM. Under some circumstances, these estimates may be negative. PROC MIXED gets these variance component estimates right.
- Since PROC GLM does not really understand random effects and mixed models, it tests all hypotheses by using an F statistic formed by dividing the mean square for an effect by the mean square error. In mixed models, the resulting F tests often do not test the hypothesis that one wants to test. One must examine expected mean squares in order to determine which mean square to use in the denominator to test a particular effect. Sometimes, there is no appropriate mean square to use in a denominator, and one must synthesize one as a function of several mean squares (if one can). PROC MIXED does not construct its tests in this manner, and does not encounter this problem.
- The basic model that PROC GLM fits assumes that the errors are all independently and identically distributed. When we use PROC GLM to fit mixed models, we assume that all variance components are independently and (within each variance component) identically distributed. We can in a sense finesse the independence portion by assuming a very simple structure of dependence among errors, but that very simple structure is often wrong, especially with repeated measures data. PROC MIXED does not assume that the errors or variance components are independently and (within variance components) identically distributed. PROC MIXED does not need to make this assumption because it incorporates the

structure and relationships among the errors and variance components into the model. This opens new potential in our modeling. We no longer need to assume something about the errors (that they are all independently and identically distributed); we can model them, estimate parameters for them, test hypotheses about them, and place confidence intervals on them.

Example 1

Consider a two-factor, fixed effects, between-subjects design, of the type appropriate to analyze using PROC GLM. Suppose the two factors are called (imaginatively) "A" and "B". In this example, we have 2 levels of A, 3 levels of B, and 10 observations per cell. We have:

```
proc glm data=one;
  class a b;
  model y=a b a*b;
  lsmeans a b a*b;
run;
proc mixed data=one;
  class a b;
  model y=a b a*b;
  lsmeans a b a*b;
run;
```

(Remark: in this example, and in all following examples, we use Type III tests, which are appropriate in most circumstances. This choice has been the subject of many presentations and publications and is not without some controversy.)

Using some fictional data, we obtain the following output from PROC GLM:

Source	F Value	Pr > F
a	0.85	0.3601
b	3.66	0.0322
a*b	0.22	0.8010

With the same data, PROC MIXED gives us:

Effect	Num DF	Den DF	F Value	Pr > F
a	1	54	0.85	0.3601
b	2	54	3.66	0.0322
a*b	2	54	0.22	0.8010

Note that for a balanced fixed effects design, although the methods used to solve for parameter estimates and test hypotheses are different, the estimates and tests are the same. The standard errors of the parameter estimates are identical, and thus, the estimates for least squares

means, and standard errors of the least squares means are identical.

Comment 1: When the data are unbalanced, it cannot be guaranteed, even for entirely fixed effects designs, that the estimates and hypothesis tests will be the same. This is due to the differences in estimation methods.

Example 2

Consider a two-factor, mixed effects, repeated measures design, with one between-subjects factor (called GROUP), and one within-subjects factor (called TIME). This is a split-plot type structure, with subjects playing the role of whole plots, and the levels of time playing the role of the subplots. In a real split plot design, we would randomly assign treatments to subplots, thus ensuring that except for treatment effects, the subplots (corresponding to the levels of the factor we are calling TIME) are all independent of each other. In a repeated measures design, even without considering the effects of the within-subjects treatments, the occasions, or levels of TIME, are probably not independent of each other. If our experimental units are living organisms, for example, it is usually the case that things measured closely together in time correlate more highly than things measured farther apart in time. More generally, repeated measurements made on the same subjects are nearly always correlated in some way.

Remark: In this dataset, all the effects of interest are assumed to be fixed effects. That is, the two groups we have are assumed to be the only groups to which we wish to generalize from our samples, and the four time points are the only time points to which we wish to generalize. The only random effect is the SUBJECT(GROUP) effect.

There are several ways we can analyze these data with PROC GLM, and several ways we can analyze them with PROC MIXED. (For this example, GROUP has 2 levels, and TIME has 4 levels.)

The oldest way to fit this model is the classic univariate approach to repeated measures, in which we fit it as a split plot type design, with subjects playing the role of the whole plots, the subjects randomly assigned to levels of GROUP, the levels of time playing the role of the subplots. To fit this with PROC GLM we use:

```
proc glm data=twob;
  class group time subject;
  model response = group time group*time
                subject (group);
  test h=group e=subject (group);
```

This analysis assumes that the correlations among the levels of time are a constant (sphericity). It was necessary

to use a TEST statement, because the default test of the GROUP effect uses the mean square error as the denominator of the F statistic, and examination of expected mean squares would show that that is the wrong mean square to use. This analysis is not capable of testing whether the sphericity assumption is met, and not capable of performing either of the common corrections for violations of the sphericity assumption.

We can obtain the same analysis using PROC GLM with the REPEATED statement:

```
proc glm data=twoa;
  class group;
  model time1-time4=group;
  repeated time / printe;
  run;
```

Note several differences between the SAS code in this example and the previous one:

1. In the original example, which used data in the SAS dataset called "twob," the data were stored with one observation per subject-occasion, and there was a variable called TIME which distinguished for which level of time a particular value for the dependent variable, RESPONSE. In other words, for these data, which had 2 groups, 20 subjects per group, and 4 time points per subject, the SAS dataset "twoa" had $2 \times 20 \times 4 = 160$ observations.
2. In the second example, using the REPEATED statement, the same data were stored differently in the "twoa" dataset, with one observation per subject, and the values for the 4 time points stored as separate variables, TIME1-TIME4. The REPEATED statement produces *exactly* the same split plot type analysis, with *exactly* the same results as did the earlier analysis, and additionally produces other analyses and output. Most notable are that it produces a multivariate approach to repeated measures, a partial correlation matrix among the levels of the within-subject factors, from which it constructs a test of whether the sphericity assumption is met, and a multivariate approach to repeated measures. We will not discuss the multivariate approach here, except to note that it may be a reasonable approach to some repeated measures analyses.

We can also analyze the same data with PROC MIXED. Again, this is an analysis in which the only random effect is the SUBJECT(GROUP) effect. It turns out that several specifications for PROC MIXED will produce identical analyses:

```

proc mixed data=twob;
  class group time subject;
  model response = group time
                    group*time;
  random subject(group);
run;

proc mixed data=twob;
  class group time subject;
  model response = group time
                    group*time;
  random subject(group) / type=cs;
run;

proc mixed data=twob;
  class group time subject;
  model response = group time
                    group*time;
  repeated / subject=subject(group)
            type=cs;
run;

```

The first program is the closest specification to that using PROC GLM using a split plot type approach. In PROC GLM, the MODEL statement generates a model matrix, X , with appropriate columns for a constant, group, time, the interaction, and subjects(group). In the first PROC MIXED, we are generating exactly the same columns with the MODEL and RANDOM statements; it is just that generating the subjects(group) columns using the RANDOM statement places those columns in the Z matrix. The results are identical to those produced by PROC GLM for balanced data, except possibly for differences due to the estimation methods used in the two procedures. That should rarely occur.

The second program is identical to the first, except for the use of the `type=cs` option on the RANDOM statement. This produces identical results to those obtained by PROC GLM and by the first PROC MIXED program with two exceptions:

1. By using the `type=cs` option on the random statement, we told PROC MIXED to fit an extra parameter, corresponding to the variance component for subjects(group). Because this extra parameter corresponds exactly to a parameter we are already fitting, the estimate was exactly zero, with no degrees of freedom.
2. This produced an error message both in the PROC MIXED output and in the SAS log, stating that the

```

Convergence criteria met but final
hessian is not positive definite.

```

This message occurred because we attempted to fit the same parameter twice. PROC MIXED set the second parameter estimate to zero, proceeded normally, and

arrived at the correct answers. Thus, the mere fact that a message about a non-positive-definite hessian is generated might not mean that you have a problem. We will see later that this message can occur on other occasions.

Comment 2: A message about a non-positive-definite hessian is *not necessarily indicative* of something about which you should worry. This message *may* arise because you have overspecified the model in a way that PROC MIXED can handle.

Comment 3: Take care in specifying a covariance structure among the random effects. It may not be necessary, because you may already have a parameter for it.

The third program uses a different PROC MIXED statement, the REPEATED statement.

Comment 4: There is a REPEATED statement available in PROC GLM, in PROC MIXED, (and in PROC CATMOD, although we don't discuss PROC CATMOD here). **These REPEATED statements are three different statements. They have different syntaxes, they specify different things, they are used in different procedures, and they have, in general, different effects. Do not confuse them. They work differently from each other.** What they do have in common is that they all tell their procedures how to fit repeated measures models.

Finally, we have a fourth PROC MIXED program:

```

proc mixed data=twob;
  class group time subject;
  model response = group time
                    group*time;
  repeated / subject=subject(group);
run;

```

This program produces a different set of answers than did the other three PROC MIXED programs. This program did not fit a variance component for the subjects (group), but treated it as the residual. This changed all the other tests.

Comment 5: In a split plot type repeated measures design, the default covariance structure that PROC MIXED fits when you use the RANDOM statement to specify the within-subjects factor is compound symmetry, while the default covariance structure that PROC MIXED uses when you use the REPEATED statement to specify the subjects factor, within which the repeated measures occur, is not, but rather is the variance component structure, which assumes that the repeated measures are independent of one another.

What is Sphericity, Anyway?

When analyzing a repeated measures design using a split-plot type approach, it is often desirable to test whether the assumptions of the analysis are met; that is, whether the covariance structure is plausibly that which you assume. There are at least five ways in which the literature states the assumption you must make: *compound symmetry, sphericity, circularity, a Huynh-Feldt (H-F) covariance structure, and a Type H covariance structure*. Type H, H-F, and circularity are the same condition, and sphericity of an orthogonal decomposition of the covariance matrix follows if the covariance matrix is circular, so these four names are really the four names for the same assumption. Compound symmetry is a stronger condition, a special case of H-F. So we only need to define two forms for the covariance matrix to cover the common possibilities. If we assume, for example, four repeated measures, we have:

Compound Symmetry :

$$\begin{bmatrix} \sigma^2 + \sigma_1 & \sigma_1 & \sigma_1 & \sigma_1 \\ \sigma_1 & \sigma^2 + \sigma_1 & \sigma_1 & \sigma_1 \\ \sigma_1 & \sigma_1 & \sigma^2 + \sigma_1 & \sigma_1 \\ \sigma_1 & \sigma_1 & \sigma_1 & \sigma^2 + \sigma_1 \end{bmatrix}$$

Huynh-Feldt:

$$\begin{bmatrix} \sigma_1^2 & \frac{\sigma_1^2 + \sigma_2^2}{2 - \lambda} & \frac{\sigma_1^2 + \sigma_3^2}{2 - \lambda} & \frac{\sigma_1^2 + \sigma_4^2}{2 - \lambda} \\ \frac{\sigma_2^2 + \sigma_1^2}{2 - \lambda} & \sigma_2^2 & \frac{\sigma_2^2 + \sigma_3^2}{2 - \lambda} & \frac{\sigma_2^2 + \sigma_4^2}{2 - \lambda} \\ \frac{\sigma_3^2 + \sigma_1^2}{2 - \lambda} & \frac{\sigma_3^2 + \sigma_2^2}{2 - \lambda} & \sigma_3^2 & \frac{\sigma_3^2 + \sigma_4^2}{2 - \lambda} \\ \frac{\sigma_4^2 + \sigma_1^2}{2 - \lambda} & \frac{\sigma_4^2 + \sigma_2^2}{2 - \lambda} & \frac{\sigma_4^2 + \sigma_3^2}{2 - \lambda} & \sigma_4^2 \end{bmatrix}$$

As stated, H-F, circularity, Type H, and sphericity on an orthogonal decomposition are the same.

Comment 5: PROC GLM tests for sphericity of an orthogonal decomposition of the covariance matrix when you use the PRINTE option. The null hypothesis is that the covariance structure fits the sphericity structure, and the alternative is that it fails to fit. Thus, rejection of the null hypothesis is rejection of the assumption that the data meet the sphericity assumption. **However, the model that GLM fits actually assumes compound symmetry, a slightly different although very similar assumption. It is often impossible to differentiate between them in a set of real data.**

Comment 6: To make things more complex, the corrected p-values presented by GLM are based on

changes to the degrees of freedom, and do not actually change the model fit, whether one chooses to consider the Greenhouse-Geisser corrected p-value, which deflates the *df* for the extent to which compound symmetry is violated, or the Huynh-Feldt corrected p-value, which deflates the *df* for the extent to which the H-F conditions are violated.

When using PROC GLM to analysis such a design, the test of sphericity is done using the PRINTE option on the REPEATED statement. This tests whether restricting the covariance structure to the Huynh-Feldt structure above fits significantly worse than an unrestricted covariance structure, that is, allowing arbitrary correlations among the repeated measures:

```
proc glm data=twoa;
  class group;
  model time1-time4=group;
  repeated time / printe;
  title 'Repeated Measures with
        Repeated Statement';
run;
```

Among the information that the PRINTE option causes to be printed is Mauchly's test of sphericity.

Comment 7: If the transformation you specify (or, by not specifying a transformation, allow to default) on the REPEATED statement is not an orthogonal transformation, PROC GLM gives you two tests. **The test that in fact tests whether the covariance matrix of interest meets the H-F conditions is the test of sphericity on orthogonal components of the transformed variates.**

You can obtain an equivalent test using PROC MIXED. However, there are several caveats about doing this. First, you must remember that the sphericity test that PROC GLM performs tests the H-F conditions. Therefore, if you wish to test the same conditions using PROC MIXED, you must test whether the covariance matrix among the repeated measures meets the H-F conditions. Second, you must do this using some hand calculations or SAS programming outside of PROC MIXED. The way to test whether the data meet the H-F criteria is to test, using a likelihood ratio test, whether the fit of a model in which the covariance matrix is constrained to the H-F conditions, is significantly worse than the fit of a model in which the covariance matrix is unconstrained. You must do this using two runs of PROC MIXED, take the difference between the two -2 log likelihoods you obtain, and compare that difference to a Chi-Square with *df* equal to the difference between the two *df* in the two competing models:

```

proc mixed data=twob;
  class group time subject;
  model response = group time
                group*time;
  repeated / subject = subject(group)
                type=hf;
  title 'Mixed with Random statement
        with HF conditions';
run;
proc mixed data=twob;
  class group time subject;
  model response = group time
                group*time;
  repeated / subject = subject(group)
                type=un;
  title 'Mixed with Random Statement
        with UN conditions';
run;

```

When PROC GLM performed Mauchly's test on an orthogonal decomposition of the covariance matrix on our example data, it obtained a $\chi^2 = 5.74$, with 5 df (nonsignificant). When we ran the above two PROC MIXED procedures, we obtain the following -2 log likelihoods: 450.3234 (with 5 parameters) and 444.3780 (with 10 parameters). The difference between them is 5.9415 (distributed under the null hypothesis as a χ^2 with 10-5=5 df; nonsignificant).

Comment 8: When attempting to construct a test of sphericity using PROC MIXED, parallel to the one constructed by PROC GLM, remember to compare an H-F covariance structure to an unrestricted covariance structure.

Comment 9: These tests that GLM and MIXED produce will not be exactly numerically the same, but will be close.

Comment 10: You can use exactly this procedure to test the difference between any two covariance structures. For example, suppose you wanted to know if H-F fit significantly better than CS. You would fit the model twice, once with TYPE=CS, once with TYPE=HF, take the difference between the two -2 log likelihoods, and compare it to a χ^2 with *df* equal to the difference in the numbers of parameters.

Comment 11: Fitting different covariance structures. If your experience with repeated measures designs is constrained to traditional split-plot-type approaches, or to use of a multivariate approach, the idea of fitting different covariance structures may be new. When you used a multivariate approach, you actually fit an unrestricted covariance structure. That is, the correlation matrix among the repeated measures was not constrained to any particular form; each correlation was estimated. When you use PROC MIXED, the methodology exists to constrain the covariance matrix in many ways, that is, to fit parameters for the correlations among the repeated measures. The more parameters you fit, the more degrees of freedom you use up, and the less powerful

your other tests become. The correlation structure needs to be considered a priori.

Comment 12: There are many other covariance structures you may fit. Most of them are probably not very appropriate to most repeated measures designs. The structures we think are reasonable for most repeated measures designs are the following:

1. Variance Components: Use this structure when you feel, except for treatment effects, that the repeated measures are all of equal variance and are uncorrelated.
2. Compound Symmetry: Use this structure when you feel that, except for treatment effects, all the repeated measures are equally correlated, and have equal variance. Recall that sphericity of the correlation matrix implies compound symmetry of the covariance matrix, and that it is sphericity of the correlation matrix that GLM tests.
3. Huynh-Feldt: Use this structure when you feel that, except for the treatment effects, the repeated measures may have different variances, and the correlations may be proportional to the averages of the two variances. The Huynh-Feldt structure is more flexible than compound symmetry, and although it fits a few more parameters, it does not fit an excessive number. Compound symmetry is a special case of the Huynh-Feldt structure.
4. Auto-Regressive Lag 1: Use this structure when you feel that, except for treatment effects, the repeated measures may have equal variances, but the correlation between the repeated measures, except for treatment effects, may be proportional to the distance between the repeated measures. This is not an unusual situation to have occur in repeated measures, over time, on living organisms. In that case, this would be saying that the farther apart in time two measures are taken, the lower their correlation. In such repeated measures designs, this structure should probably be considered.

Comment 13: Examining different covariance structures. There are at least two popular strategies to fitting and examining statistical models, and subject areas differ with respect to commonly used strategies. One strategy, commonly used in the analysis of designed experiments, is to decide at the time you design the experiment the model you will fit and the analysis you will perform. The other is to fit and refit many models to the same data, with the goal of finding the best-fitting model (in some sense). Since the mixed model allows you to fit a new class of parameters (the covariance structure parameters), this model offers the opportunity for even more iterative model fitting. You should not blindly fit all possible covariance structures, and chose the one that appears to fit best. This greatly risks incorporating chance into the model parameter estimates. On the other hand, if you fit the wrong covariance structure, your tests of fixed effects are wrong. In general, you should allow your knowledge of the experiment, experimental units, and repeated

measures to lead you to judicious modeling of the covariance structures.

Comment 14: There are closely related *heterogeneous* and *banded* versions of many of these structures. Heterogeneous means that the variances of the repeated measures may not be constrained to be equal, and banded means that after a specified distance between two repeated measures, the correlations drop to zero. These may be reasonable.

Covariates in Mixed Models

Frequently in both fixed effect and mixed models, one may have covariates. For example, in a repeated measures design, with one between-subjects factor, and one within-subjects factor (just the example we have used in the last section), we might have a covariate. We use the term covariate in the same sense as it is often used in linear models discussions, that is, a continuous predictor variable, often viewed as a nuisance variable, for which we would like to control by putting it in the model when we examine the variables in which we're interested.

In a repeated measures design, there are two common types of covariates, called *subject-dependent* covariates and *time-dependent* covariates. Subject-dependent covariates are measured once per subject, and are thus constant over the repeated measures. Time-dependent covariates are measured at of the repeated measure occasions. Further, we assume without loss of generality, and for simplicity, that there are no ties; that each subject has a unique value for the covariate.

For example, if we measured the weight of each subject at the beginning of a clinical trial, and wished to use weight as a covariate in a repeated measures analysis, weight would be a subject-dependent covariate. On the other hand, if we measured weight on each occasion for each subject, then weight would be a time-dependent covariate.

We discuss subject-dependent covariates only.

If we attempt to simply add a covariate to the analysis we ran earlier using GLM we have the following program:

```
proc glm data=nontime;
  class group time subj;
  model y=group time group*time
        subj(group) weight;
  test h=group e=subj(group);
  title 'GLM with covariate';
  run;
```

This program produces a problematic result. In the GLM ANOVA table, the covariate WEIGHT is labeled as having 0 *df*, and the sums of squares, mean squares, and *F* statistics are missing for this covariate. This occurs because PROC GLM parameterizes the SUBJ(GROUP) effect with a one column indicator variable for each subject in the study. In other words, there are as many

columns for the SUBJ(GROUP) effect as there are subjects, each with a 1 in the row in X corresponding each unique subject, and a 0 in the other rows. When you include a covariate, the covariate column becomes calculable as an exact linear combination of the columns corresponding to the SUBJ(GROUP) effect and is thus redundant. When GLM finds redundant columns, it assigns them 0 *df* and no variation. The parameter estimates and hypothesis tests are identical to those produced by GLM when there is no covariate.

PROC MIXED behaves differently. The following program produces an analysis for the same data using PROC MIXED.

```
proc mixed data=nontime;
  class group time subj;
  model y=group time group*time
        weight;
  random subj(group);
  title 'MIXED with covariate (in
        model statement)';
  run;
```

Note that in this program, the covariate, WEIGHT, is used in the MODEL statement thus, considered a fixed effect, while the SUBJ(GROUP) effect is used in the RANDOM statement, thus considered a random effect. The columns corresponding to the SUBJ(GROUP) effect are thus placed in the X matrix, while the columns corresponding to the covariate, WEIGHT, are placed in the Z matrix. The redundancy is avoided.

If we decided to consider WEIGHT to be a random effect, and thus moved it to the RANDOM statement in MIXED, we obtain the same results as we did using GLM: the covariate has 0 *df*, no SS, and the parameter estimates MIXED produces are identical to those produced by GLM.

Discussion:

The MIXED procedure is extraordinarily flexible, and a useful addition to the toolbox for the design and analysis of linear models studies (useful for design because it is silly to design a study for which the tools to analyze the resulting data properly are not available). However, the model that it fits is a more complex model than the univariate general linear model that PROC GLM fits, and the procedure is correspondingly more complex. To use this procedure properly requires training and caution.

Contact Information

Robert M. Hamer, Ph.D.
 Department of Psychiatry
 UMDNJ Robert Wood Johnson Medical School
 675 Hoes Lane
 Piscataway NJ 08854
 732 235 4218
 hamer@umdnj.edu

Pippa Simpson, Ph.D.
Arkansas Children's Hospital
PEDS/CARE, So. Campus Room 301
Little Rock, Arkansas 72202
(501) 320-6631
Psimpson@care.ach.uams.edu

Acknowledgment

The authors wish to thank Leonard Feldt, Ph.D., for helpful comments.

References

Littell, R.C., Milliken, G.A., Stroup, W.W., and Wolfinger, R.D. *The SAS System for Mixed Models*. Cary, NC: SAS Institute, 1996.

Milliken, G.A., and Johnson, D.E. *Analysis of Messy Data Volume I: Designed Experiments*. London: Chapman & Hall, 1992.

Milliken, G. A., and Johnson, D.E. *Analysis of Messy Data Volume II: Nonreplicated Experiments*. New York: Chapman & Hall, 1989.

Milliken, G. A., and Johnson, D.E. *Analysis of Messy Data Volume III: Analysis of Covariance*. In press, 2000.