Technical report

SPSS

# Linear Mixed-Effects Modeling in SPSS: An Introduction to the MIXED Procedure

## Table of contents

## Introduction

The linear mixed-effects models (MIXED) procedure in SPSS enables you to fit linear mixed-effects models to data sampled from normal distributions. Recent texts, such as those by McCulloch and Searle (2000) and Verbeke and Molenberghs (2000), comprehensively review mixed-effects models. The MIXED procedure fits models more general than those of the general linear model (GLM) procedure and it encompasses all models in the variance components (VARCOMP) procedure. This report illustrates the types of models that MIXED handles. We begin with an explanation of simple models that can be fitted using GLM and VARCOMP, to show how they are translated into MIXED. We then proceed to fit models that are unique to MIXED.

The major capabilities that differentiate MIXED from GLM are that MIXED handles correlated data and unequal variances. Correlated data are very common in such situations as repeated measurements of survey respondents or experimental subjects. MIXED extends repeated measures models in GLM to allow an unequal number of repetitions. It also handles more complex situations in which experimental units are nested in a hierarchy. MIXED can, for example, process data obtained from a sample of students selected from a sample of schools in a district.

In a linear mixed-effects model, responses from a subject are thought to be the sum (linear) of so-called fixed and random effects. If an effect, such as a medical treatment, affects the population mean, it is fixed. If an effect is associated with a sampling procedure (e.g., subject effect), it is random. In a mixed-effects model, random effects contribute only to the covariance structure of the data. The presence of random effects, however, often introduces correlations between cases as well. Though the fixed effect is the primary interest in most studies or experiments, it is necessary to adjust for the covariance structure of the data. The adjustment made in procedures like GLM-Univariate is often not appropriate because it assumes independence of the data.

The MIXED procedure solves these problems by providing the tools necessary to estimate fixed and random effects in one model. MIXED is based, furthermore, on maximum likelihood (ML) and restricted maximum likelihood (REML) methods, versus the analysis of variance (ANOVA) methods in GLM. ANOVA methods produce an optimum estimator (minimum variance) for balanced designs, whereas ML and REML yield asymptotically efficient estimators for balanced and unbalanced designs. ML and REML thus present a clear advantage over ANOVA methods in modeling real data, since data are often unbalanced. The asymptotic normality of ML and REML estimators, furthermore, conveniently allows us to make inferences on the covariance parameters of the model, which is difficult to do in GLM.

## Data preparation for MIXED

Many datasets store repeated observations on a sample of subjects in "one subject per row" format. MIXED, however, expects that observations from a subject are encoded in separate rows. To illustrate, we select a subset of cases from the data that appear in Potthoff and Roy (1964). The data shown in Figure 1 encode, in one row, three repeated measurements of a dependent variable ("dist1" to "dist3") from a subject observed at different ages ("age1" to "age3").

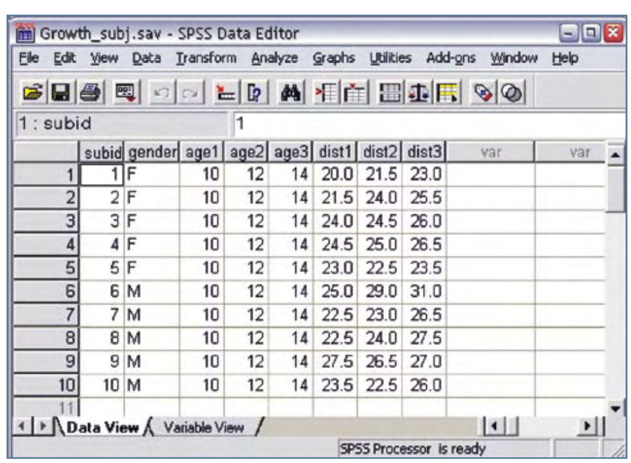


| | subid | gender | age1 | age2 | age3 | dist1 | dist2 | dist3 | var | var |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 1 | F | 10 | 12 | 14 | 20.0 | 21.5 | 23.0 | | |
| 2 | 2 | F | 10 | 12 | 14 | 21.5 | 24.0 | 25.5 | | |
| 3 | 3 | F | 10 | 12 | 14 | 24.0 | 24.5 | 26.0 | | |
| 4 | 4 | F | 10 | 12 | 14 | 24.5 | 25.0 | 26.5 | | |
| 5 | 5 | F | 10 | 12 | 14 | 23.0 | 22.5 | 23.5 | | |
| 6 | 6 | M | 10 | 12 | 14 | 25.0 | 29.0 | 31.0 | | |
| 7 | 7 | M | 10 | 12 | 14 | 22.5 | 23.0 | 26.5 | | |
| 8 | 8 | M | 10 | 12 | 14 | 22.5 | 24.0 | 27.5 | | |
| 9 | 9 | M | 10 | 12 | 14 | 27.5 | 26.5 | 27.0 | | |
| 10 | 10 | M | 10 | 12 | 14 | 23.5 | 22.5 | 26.0 | | |
| 11 | | | | | | | | | | |

**Figure 1. MIXED, however, requires that measurements at different ages be collapsed into one variable, so that each subject has three cases. The Data Restructure Wizard in SPSS simplifies the tedious data conversion process. We choose "Data->Restructure" from the pull-down menu, and select the option "Restructure selected variables into cases." We then click the "Next" button to reach the dialog shown in Figure 2.**

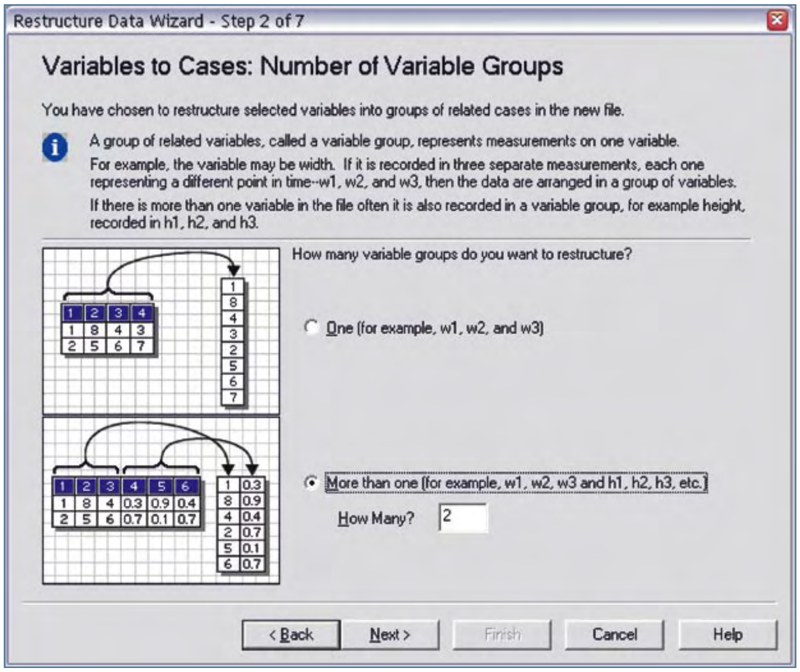


Figure 2. We need to convert two groups of variables ("age" and "dist") into cases. We therefore enter "2" and click "Next." This brings us to the "Select Variables" dialog box.

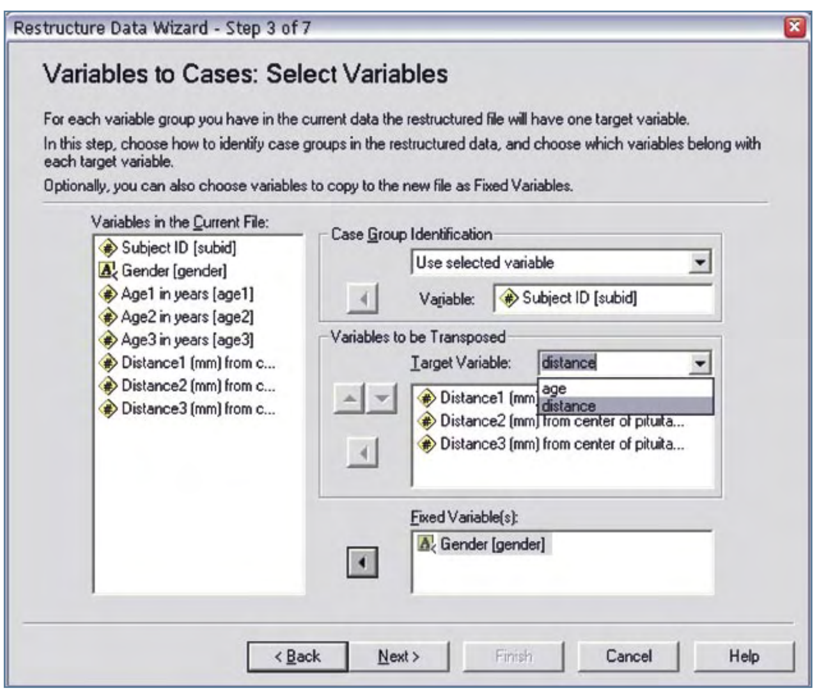


Figure 3. In the "Select Variables" dialog box, we first specify "Subject ID [subid]" as the case group identification. We then enter the names of new variables in the target variable drop-down list. For the target variable "age," we drag "age1," "age2," and "age3" to the list box in the "Variables to be Transposed" group. We similarly associate variables "dist1," "dist2," and "dist3" with the target variable "distance." We then drag variables that do not vary within a subject to the "Fixed Variable(s)" box. Clicking "Next" brings us to the "Create Index Variables" dialog box. We accept the default of one index variable, then click "Next" to arrive at the final dialog box.

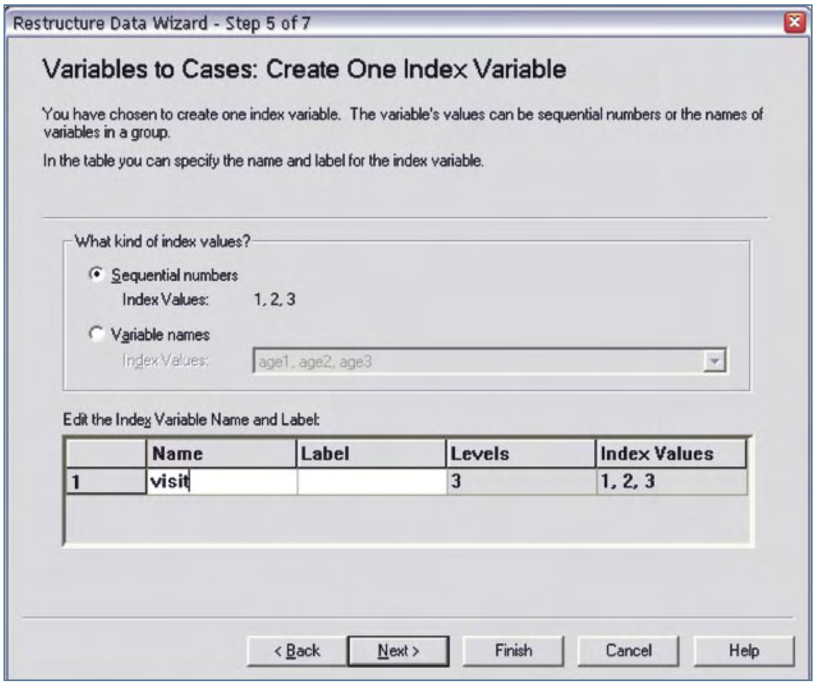


Figure 4. In the "Create One Index Variable" dialog box, we enter "visit" as the name of the indexing variable and click "Finish."

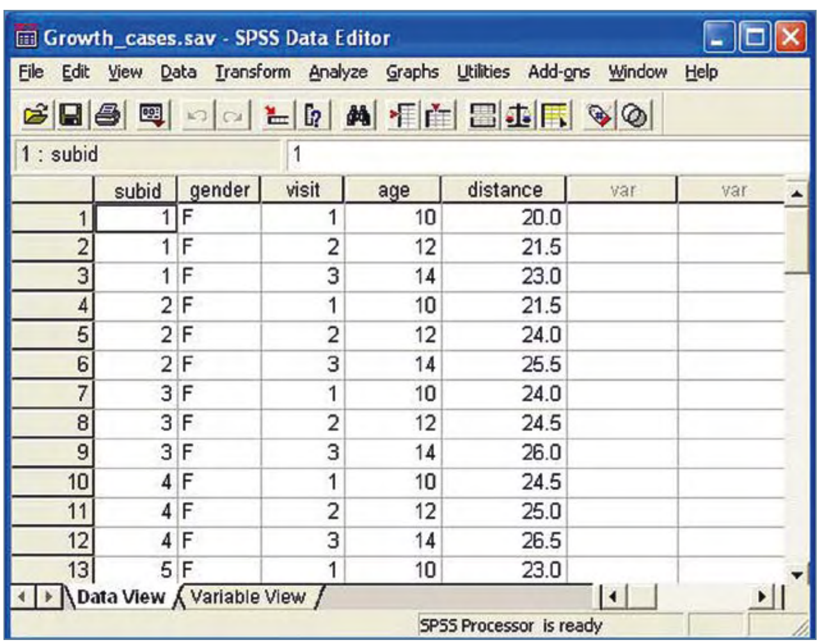


| | subid | gender | visit | age | distance | var | var |
|---|---|---|---|---|---|---|---|
| 1 | 1 | F | 1 | 10 | 20.0 | | |
| 2 | 1 | F | 2 | 12 | 21.5 | | |
| 3 | 1 | F | 3 | 14 | 23.0 | | |
| 4 | 2 | F | 1 | 10 | 21.5 | | |
| 5 | 2 | F | 2 | 12 | 24.0 | | |
| 6 | 2 | F | 3 | 14 | 25.5 | | |
| 7 | 3 | F | 1 | 10 | 24.0 | | |
| 8 | 3 | F | 2 | 12 | 24.5 | | |
| 9 | 3 | F | 3 | 14 | 26.0 | | |
| 10 | 4 | F | 1 | 10 | 24.5 | | |
| 11 | 4 | F | 2 | 12 | 25.0 | | |
| 12 | 4 | F | 3 | 14 | 26.5 | | |
| 13 | 5 | F | 1 | 10 | 23.0 | | |

Figure 5. We now have three cases for each subject.

We can also perform the conversion using the following command syntax:

```
VARSTOCASES
/MAKE age FROM age1 age2 age3
/MAKE distance FROM dist1 dist2 dist3
/INDEX = visit(3)
/KEEP = subid gender.
```

The command syntax is easy to interpret—it collapses the three age variables into “age” and the three response variables into “distance.” At the same time, a new variable, “visit,” is created to index the three new cases within each subject. The last subcommand means that the two variables that are constant within a subject should be kept.

## Fitting fixed-effects models

### With iid residual errors

A fitted model has the form $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$, where $\mathbf{y}$ is a vector of responses, $\mathbf{X}$ is the fixed-effects design matrix, $\boldsymbol{\beta}$ is a vector of fixed-effects parameters and $\boldsymbol{\varepsilon}$ is a vector of residual errors. In this model, we assume that $\boldsymbol{\varepsilon}$ is distributed as $\mathbf{R} = \boldsymbol{\sigma}^2\mathbf{I}$ where $\mathbf{R}$ is an unknown covariance matrix. A common belief is that $\mathbf{R} = \boldsymbol{\sigma}^2\mathbf{I}$. We can use GLM or MIXED to fit a model with this assumption. Using a subset of the growth study dataset, we illustrate how to use MIXED to fit a fixed-effects model. The following command (Example 1) fits a fixed-effects model that investigates the effect of the variables “gender” and “age” on “distance,” which is a measure of the growth rate.

*Example 1: Fixed-effects model using MIXED*

Command syntax:

```
MIXED DISTANCE BY GENDER WITH AGE
/FIXED = GENDER AGE | SSTYPE(3)
/PRINT = SOLUTION TESTCOV.
```

Output:

**Type III Tests of Fixed Effects [a]**

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 27.000 | 38.356 | .000 |
| gender | 1 | 27 | 7.621 | .010 |
| age | 1 | 27.000 | 11.040 | .003 |

a. Dependent Variable: Distance (mm) from center of pituitary t pteryo-maxillary fissure.

Figure 6

**Estimates of Fixed Effects [b]**

| Parameter | Estimate | Std. Error | df | t | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Intercept | 17.050 | 2.620 | 27.000 | 6.507 | .000 | 11.673 | 22.427 |
| [gender=F] | -1.933 | .700 | 27.000 | -2.761 | .010 | -3.370 | -.496 |
| [gender=M] | .000[a] | .000 | . | . | . | . | . |
| age | .713 | .214 | 27.000 | 3.323 | .003 | .273 | 1.152 |

a. This parameter is set to zero because it is redundant.

b. Dependent Variable: Distance (mm) from center of pituitary to pteryo-maxillary fissure.

Figure 7

The command in Example 1 produces a "Type III Tests of Fixed Effects" table (Figure 6). Both "gender" and "age" are significant at the .05 level. This means that "gender" and "age" are potentially important predictors of the dependent variable. More detailed information on fixed-effects parameters may be obtained by using the subcommand /PRINT SOLUTION. The "Estimates of Fixed Effects" table (Figure 7) gives estimates of individual parameters, as well as their standard errors and confidence intervals. We can see that the mean distance for males is larger than that for females. Distance, moreover, increases with age. MIXED also produces an estimate of the residual error variance and its standard error. The /PRINT TESTCOV option gives us the Wald statistic and the confidence interval for the residual error variance estimate.

**Estimates of Covariance Parameters** [a]

| Parameter | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|
| Residual | 3.679 | 1.001 | 3.674 | .000 | 2.158 | 6.271 |

a. Dependent Variable: Distance (mm) from center of pituitary to pteryo-maxillary fissure.

Figure 8

Example 1 is simple—users familiar with the GLM procedure can fit the same model using GLM.

*Example 2: Fixed-effects model using GLM*

Command syntax:

```
GLM DISTANCE BY GENDER WITH AGE
/METHOD = SSTYPE(3)
/PRINT = PARAMETER
/DESIGN = GENDER AGE.
```

Output:

**Tests of Between-Subjects Effects**

Dependent Variable: Distance (mm) from center of pituitary to pteryo-maxillary fissur

| Source | Type III Sum of Squares | df | Mean Square | F | Sig. |
|---|---|---|---|---|---|
| Corrected Model | 68.646[a] | 2 | 34.323 | 9.331 | .001 |
| Intercept | 141.095 | 1 | 141.095 | 38.356 | .000 |
| gender | 28.033 | 1 | 28.033 | 7.621 | .010 |
| age | 40.613 | 1 | 40.613 | 11.040 | .003 |
| Error | 99.321 | 27 | 3.679 | | |
| Total | 18372.000 | 30 | | | |
| Corrected Total | 167.967 | 29 | | | |

a. R Squared = .409 (Adjusted R Squared = .365)

Figure 9

We see in Figure 9 that GLM and MIXED produced the same Type III tests and parameter estimates. Note, however, that in the MIXED "Type III Tests of Fixed Effects" table (Figure 6), there is no column for the sum of squares. This is because, for some complex models, the test statistics in MIXED may not be expressed as a ratio of two sums of squares. They are thus omitted from the ANOVA table.

**Parameter Estimates**

Dependent Variable: Distance (mm) from center of pituitary to pteryo-maxillary fissure

| | | | | | 95% Confidence Interval | |
|---|---|---|---|---|---|---|
| Parameter | B | Std. Error | t | Sig. | Lower Bound | Upper Bound |
| Intercept | 17.050 | 2.620 | 6.507 | .000 | 11.673 | 22.427 |
| [gender=F] | -1.933 | .700 | -2.761 | .010 | -3.370 | -.496 |
| [gender=M] | 0[a] | . | . | . | . | . |
| age | .713 | .214 | 3.323 | .003 | .273 | 1.152 |

a. This parameter is set to zero because it is redundant.

Figure 10

**With non-iid residual errors**

The assumption may be violated in some situations. This often happens when repeated measurements are made on each subject. In the growth study dataset, for example, the response variable of each subject is measured at various ages. We may suspect that error terms within a subject are correlated. A reasonable choice of the residual error covariance will therefore be a block diagonal matrix, where each block is a first-order autoregressive (AR1) covariance matrix.

*Example 3: Fixed-effects model with correlated residual errors*

Command syntax:

```
MIXED DISTANCE BY GENDER WITH AGE
/FIXED GENDER AGE
/REPEATED VISIT | SUBJECT(SUBID) COVTYPE(AR1)
/PRINT SOLUTION TESTCOV R.
```

Output:

**Type III Tests of Fixed Effects [a]**

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 25.723 | 75.036 | .000 |
| gender | 1 | 8.701 | 3.702 | .088 |
| age | 1 | 23.687 | 22.772 | .000 |

a. Dependent Variable: Distance (mm) from center of pituitary t pteryo-maxillary fissure.

Figure 11

**Estimates of Fixed Effects [b]**

| | | | | | | 95% Confidence Interval | |
|---|---|---|---|---|---|---|---|
| Parameter | Estimate | Std. Error | df | t | Sig. | Lower Bound | Upper Bound |
| Intercept | 17.243 | 1.947 | 26.760 | 8.857 | .000 | 13.246 | 21.239 |
| [gender=F] | -2.072 | 1.077 | 8.701 | -1.924 | .088 | -4.522 | .377 |
| [gender=M] | .000[a] | .000 | . | . | . | . | . |
| age | .713 | .149 | 23.687 | 4.772 | .000 | .404 | 1.021 |

a. This parameter is set to zero because it is redundant.

b. Dependent Variable: Distance (mm) from center of pituitary to pteryo-maxillary fissure.

Figure 12

**Estimates of Covariance Parameters** [a]

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Repeated Measures | AR1 diagonal | 3.809 | 1.467 | 2.597 | .009 | 1.791 | 8.101 |
| | AR1 rho | .729 | .120 | 6.072 | .000 | .401 | .892 |

a. Dependent Variable: Distance (mm) from center of pituitary to pteryo-maxillary fissure.

Figure 13

**Residual Covariance (R) Matrix** [a]

| | [visit = 1] | [visit = 2] | [visit = 3] |
|---|---|---|---|
| [visit = 1] | 3.809 | 2.778 | 2.026 |
| [visit = 2] | 2.778 | 3.809 | 2.778 |
| [visit = 3] | 2.026 | 2.778 | 3.809 |

First-Order Autoregressive

a. Dependent Variable: Distance (mm) from center of pituitary to pteryo-maxillary fissure

Figure 14

Example 3 uses the /REPEATED subcommand to specify a more general covariance structure for the residual errors. Since there are three observations per subject, we assume that the set of three residual errors for each subject is a sample from a three-dimensional normal distribution with a first-order autoregressive (AR1) covariance matrix. Residual errors within each subject are therefore correlated, but are independent across subjects. The MIXED procedure, by default, uses the REML method to estimate the covariance matrix. An alternative is to request ML estimates by using the /METHOD=ML subcommand.

The command syntax in Example 3 also produces the “Residual Covariance (R) Matrix” (Figure 14), which shows the estimated covariance matrix of the residual error for one subject. We see from the “Estimates of Covariance Parameters” table (Figure 13) that the correlation parameter has a relatively large value (.729) and that the p-value of the Wald test is less than .05. The autoregressive structure may fit the data better than the model in Example 1.

We also see that, for the tests of fixed effects, the denominator degrees of freedom are not integers. This is because these statistics do not have exact F distributions. The values for denominator degrees of freedom are obtained by a Satterthwaite approximation. We see in the new model that gender is not significant at the .05 level. This demonstrates that ignoring the possible correlations in your data may lead to incorrect conclusions. MIXED is therefore usually a better alternative to GLM and VARCOMP when data are correlated.

## Fitting simple mixed-effects models

### Balanced design

MIXED, as its name implies, handles complicated models that involve fixed and random effects. Levels of an effect are, in some situations, only a sample of all possible levels. If we want to study the efficiency of workers in different environments, for example, we don’t need to include all workers in the study—a sample of workers is usually enough. The worker effect should be considered random, due to the sampling process. A mixed-effects model has, in general, the form $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\boldsymbol{\gamma} + \boldsymbol{\varepsilon}$ where the extra term $\mathbf{Z}\boldsymbol{\gamma}$ models the random effects. $\mathbf{Z}$ is the design matrix of random effects and $\boldsymbol{\gamma}$ is a vector of random-effects parameters. We can use GLM and MIXED to fit mixed-effects models. MIXED, however, fits a much wider class of models. To understand the functionality of MIXED, we first look at several simpler models that can be created in MIXED and GLM. We also look at the similarity between MIXED and VARCOMP in these models.

In examples 4 through 6, we use a semiconductor dataset that appeared in Pinheiro and Bates (2000) to illustrate the similarity between GLM, MIXED, and VARCOMP. The dependent variable in this dataset is "current" and the predictor is "voltage." The data are collected from a sample of ten silicon wafers. There are eight sites on each wafer and five measurements are taken at each site. We have, therefore, a total of 400 observations and a balanced design.

*Example 4: Simple mixed-effects model with balanced design using MIXED*

Command syntax:

```
MIXED CURRENT BY WAFER WITH VOLTAGE
/FIXED VOLTAGE | SSTYPE(3)
/RANDOM WAFER
/PRINT SOLUTION TESTCOV.
```

Output:

**Type III Tests of Fixed Effects [a]**

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 16.559 | 3774.499 | .000 |
| voltage | 1 | 389.000 | 7958.177 | .000 |

a. Dependent Variable: current.

Figure 15

**Estimates of Fixed Effects [a]**

| Parameter | Estimate | Std. Error | df | t | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Intercept | 7.082868 | .115287 | 16.559 | -61.437 | .000 | -7.326596 | 6.839139 |
| voltage | 9.648660 | .037012 | 389.000 | 260.688 | .000 | 9.575890 | 9.721429 |

a. Dependent Variable: current.

Figure 16

**Estimates of Covariance Parameters [a]**

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Residual | | .175 | .013 | 13.946 | .000 | .152 | .202 |
| wafer | Variance | .093 | .046 | 2.026 | .043 | .036 | .246 |

a. Dependent Variable: current.

Figure 17

*Example 5: Simple mixed-effects model with balanced design using GLM*

Command syntax:

```
GLM CURRENT BY WAFER WITH VOLTAGE
/RANDOM = WAFER
/METHOD = SSTYPE(3)
/PRINT = PARAMETER
/DESIGN = WAFER VOLTAGE.
```

Output:

**Tests of Between-Subjects Effects**

Dependent Variable: current

| Source | | Type III Sum of Squares | df | Mean Square | F | Sig. |
|---|---|---|---|---|---|---|
| Intercept | Hypothesis | 2229.645 | 1 | 2229.645 | 3774.499 | .000 |
| | Error | 9.782 | 16.56 | .591[a] | | |
| wafer | Hypothesis | 35.223 | 9 | 3.914 | 22.319 | .000 |
| | Error | 68.211 | 389 | .175[b] | | |
| voltage | Hypothesis | 11916.369 | 1 | 11916.369 | 7958.177 | .000 |
| | Error | 68.211 | 389 | .175[b] | | |

a. .111 MS(wafer) + .889 MS(Error)

b. MS(Error)

Figure 18

**Expected Mean Squares [a,b]**

| Source | Variance Component | | |
|---|---|---|---|
| | Var(wafer) | Var(Error) | Quadratic Term |
| Intercept | 4.444 | 1.000 | Intercept |
| wafer | 40.000 | 1.000 | |
| voltage | .000 | 1.000 | voltage |
| Error | .000 | 1.000 | |

a. For each source, the expected mean square equals the sum of the coefficients in the cells times the variance components, plus a quadratic term involving effects in the Quadratic Term cell.

b. Expected Mean Squares are based on the Type III

Figure 19

*Example 6: Variance components model with balanced design*

Command syntax:

```
VARCOMP CURRENT BY WAFER WITH VOLTAGE
/RANDOM = WAFER
/METHOD = REML.
```

Output:

**Variance Estimates**

| Component | Estimate |
|---|---|
| Var(wafer) | .093 |
| Var(Error) | .175 |

Dependent Variable: current
Method: Restricted Maximum Likelihood Estimatior

Figure 20

In Example 4, "voltage" is entered as a fixed effect and "wafer" is entered as a random effect. This example tries to model the relationship between "current" and "voltage" using a straight line, but the intercept of the regression line will vary from wafer to wafer according to a normal distribution. In the Type III tests for "voltage," we see a significant relationship between "current" and "voltage." If we delve deeper into the parameter estimates table, the regression coefficient of "voltage" is 9.65. This indicates a positive relationship between "current" and "voltage." In the "Estimates of Covariance Parameters" table (Figure 17), we have estimates for the residual error variance and the variance due to the sampling of wafers.

We repeat the same model in Example 5 using GLM. Note that MIXED produces Type III tests for fixed effects only, but GLM includes fixed and random effects. GLM treats all effects as fixed during computation and constructs F statistics by taking the ratio of the appropriate sums of squares. Mean squares of random effects in GLM are estimates of functions of the variance parameters of random and residual effects. These functions can be recovered from “Expected Mean Squares” (Figure 19). In MIXED, the outputs are much simpler because the variance parameters are estimated directly using ML or REML. As a result, there are no random-effect sums of squares.

When we have a balanced design, as in examples 4 through 6, the tests of fixed effects are the same for GLM and MIXED. We can also recover the variance parameter estimates of MIXED by using the sum of squares in GLM. In MIXED, for example, the estimate of the residual variance is 0.175, which is the same as the MS(Error) in GLM. The variance estimate of random effect “wafer” is 0.093, which can be recovered in GLM using the “Expected Mean Squares” table (Figure 19) in Example 5:

Var(WAFER) = [MS(WAFER)-MS(Error)]/40 = 0.093

This is equal to MIXED’s estimate. One drawback of GLM, however, is that you cannot compute the standard error of the variance estimates.

VARCOMP is, in fact, a subset of MIXED. These two procedures therefore always provide the same variance estimates, as seen in examples 4 and 6. VARCOMP only fits relatively simple models. It can only handle random effects that are iid. No statistics on fixed effects are produced. If your primary objective is to make inferences about fixed effects and your data are correlated, MIXED is a better choice.

An important note: Due to the different estimation methods that are used, GLM and MIXED often do not produce the same results. The next section gives an example of situations in which they produce different results.

**Unbalanced design**

One situation in which MIXED and GLM disagree is with an unbalanced design. To illustrate this, we removed some cases in the semiconductor dataset, so that the design is no longer balanced.

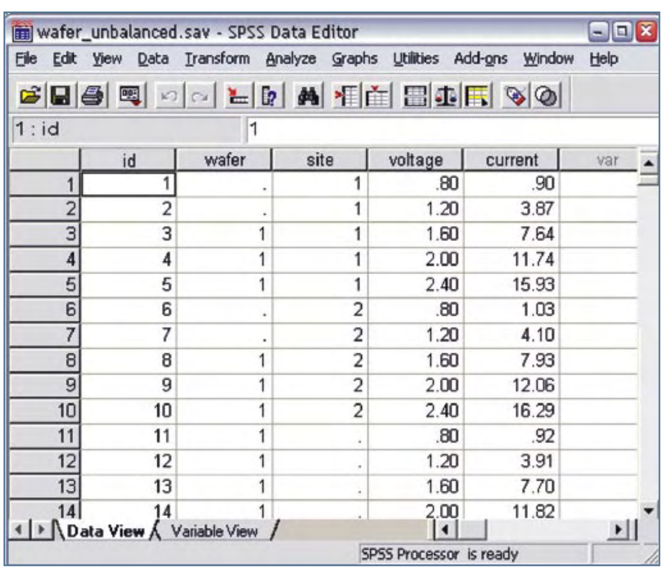


| | id | wafer | site | voltage | current | var |
|---|---|---|---|---|---|---|
| 1 | 1 | . | 1 | .80 | .90 | |
| 2 | 2 | . | 1 | 1.20 | 3.87 | |
| 3 | 3 | 1 | 1 | 1.60 | 7.64 | |
| 4 | 4 | 1 | 1 | 2.00 | 11.74 | |
| 5 | 5 | 1 | 1 | 2.40 | 15.93 | |
| 6 | 6 | . | 2 | .80 | 1.03 | |
| 7 | 7 | . | 2 | 1.20 | 4.10 | |
| 8 | 8 | 1 | 2 | 1.60 | 7.93 | |
| 9 | 9 | 1 | 2 | 2.00 | 12.06 | |
| 10 | 10 | 1 | 2 | 2.40 | 16.29 | |
| 11 | 11 | 1 | . | .80 | .92 | |
| 12 | 12 | 1 | . | 1.20 | 3.91 | |
| 13 | 13 | 1 | . | 1.60 | 7.70 | |
| 14 | 14 | 1 | . | 2.00 | 11.82 | |

Figure 21

We then rerun examples 4 through 6 with this unbalanced dataset. The output is shown in examples 4a through 6a. We want to see whether the three methods—GLM, MIXED and VARCOMP—still agree with each other.

*Example 4a: Mixed-effects model with unbalanced design using MIXED*

Command syntax:

```
MIXED CURRENT BY WAFER WITH VOLTAGE
/FIXED VOLTAGE | SSTYPE(3)
/RANDOM WAFER
/PRINT SOLUTION TESTCOV.
```

Output:

**Type III Tests of Fixed Effects [a]**

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 16.495 | 3709.960 | .000 |
| voltage | 1 | 385.037 | 7481.118 | .000 |

a. Dependent Variable: current.

Figure 22

**Estimates of Fixed Effects [a]**

| Parameter | Estimate | Std. Error | df | t | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Intercept | -7.098 | .117 | 16.495 | -60.909 | .000 | -7.344 | -7.098 |
| voltage | 9.656 | .037 | 385.037 | 259.771 | .000 | 9.583 | 9.730 |

a. Dependent Variable: current.

Figure 23

**Estimates of Covariance Parameters [a]**

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Residual | | .174451 | .013 | 13.874 | .000 | .151 | .201 |
| wafer | Variance | .095725 | .047 | 2.027 | .043 | .036 | .252 |

a. Dependent Variable: current.

Figure 24

*Example 5a: Mixed-effects model with unbalanced design using GLM*

Command syntax:

```
GLM CURRENT BY WAFER WITH VOLTAGE
/RANDOM = WAFER
/METHOD = SSTYPE(3)
/PRINT = PARAMETER
/DESIGN = WAFER VOLTAGE.
```

Output:

**Tests of Between-Subjects Effects**

Dependent Variable: current

| Source | | Type III Sum of Squares | df | Mean Square | F | Sig. |
|---|---|---|---|---|---|---|
| Intercept | Hypothesis | 2193.281 | 1 | 2193.281 | 3724.816 | .000 |
| | Error | 9.746 | 16.551 | .589[a] | | |
| wafer | Hypothesis | 35.495 | 9 | 3.944 | 22.607 | .000 |
| | Error | 67.163 | 385 | .174[b] | | |
| voltage | Hypothesis | 11772.307 | 1 | 11772.307 | 67482.629 | .000 |
| | Error | 67.163 | 385 | .174[b] | | |

a. .110 MS(wafer) + .890 MS(Error)

b. MS(Error)

Figure 25

**Expected Mean Squares a,b**

| Source | Variance Component | | |
|---|---|---|---|
| | Var(wafer) | Var(Error) | Quadratic Term |
| Intercept | 4.352 | 1.000 | Intercept |
| wafer | 39.591 | 1.000 | |
| voltage | .000 | 1.000 | voltage |
| Error | .000 | 1.000 | |

a. For each source, the expected mean square equals the sum of the coefficients in the cells times the variance components, plus a quadratic term involving effects in the Quadratic Term cell.

b. Expected Mean Squares are based on the Type III Sums of Squares.

Figure 26

*Example 6a: Variance components model with unbalanced design*

Command syntax:

```
VARCOMP CURRENT BY WAFER WITH VOLTAGE
/RANDOM = WAFER
/METHOD = REML.
```

Output:

**Variance Estimates**

| Component | Estimate |
|---|---|
| Var(wafer) | .0957247 |
| Var(Error) | .1744505 |

Dependent Variable: current
Method: Restricted Maximum Likelihood Estimatior

Figure 27

Since the data have changed, we expect examples 4a through 6a to differ from examples 4 through 6. We will focus instead on whether examples 4a, 5a, and 6a agree with each other.

In Example 4a, the F statistic for the “voltage” effect is 67481.118, but Example 5a gives an F statistic value of 67482.629. Apart from the test of fixed effects, we also see a difference in covariance parameter estimates.

Examples 4a and 6a, however, show that VARCOMP and MIXED can produce the same variance estimates, even in an unbalanced design. This is because MIXED and VARCOMP offer maximum likelihood or restricted maximum likelihood methods in estimation, while GLM estimates are based on the method-of-moments approach.

MIXED is generally preferred because it is asymptotically efficient (minimum variance), whether or not the data are balanced. GLM, however, only achieves its optimum behavior when the data are balanced.

## Fitting mixed-effects models

**With subjects**

In the semiconductor dataset, "current" is a dependent variable measured on a batch of wafers. These wafers are therefore considered subjects in a study. An effect of interest (such as "site") may often vary with subjects ("wafer"). One scenario is that the (population) means of "current" at separate sites are different. When we look at the current measured at these sites on individual wafers, however, they hover below or above the population mean according to some normal distribution. It is therefore common to enter an "effect by subject" interaction term in a GLM or MIXED model to account for the subject variations.

In the dataset there are eight sites and ten wafers. The site*wafer effect, therefore, has 80 parameters, which can be denoted by $\gamma_{ij}$, i=1...10 and j=1...8. A common assumption is that $\gamma_{ij}$'s are assumed to be iid normal with zero mean and an unknown variance. The mean is zero because $\gamma_{ij}$'s are used to model only the population variation. The mean of the population is modeled by entering "site" as a fixed effect in GLM and MIXED. The results of this model for MIXED and GLM are shown in examples 7 and 8.

*Example 7: Fitting random effect*subject interaction using MIXED*

Command syntax:

```
MIXED CURRENT BY WAFER SITE WITH VOLTAGE
/FIXED SITE VOLTAGE |SSTYPE(3)
/RANDOM SITE*WAFER | COVTYPE(ID).
```

Output:

**Type III Tests of Fixed Effects** [a]

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 329.796 | 0467.974 | .000 |
| site | 7 | 72.000 | 1.140 | .348 |
| voltage | 1 | 319.000 | 6639.444 | .000 |

a. Dependent Variable: current.

Figure 28

**Estimates of Covariance Parameters** [a]

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Residual | | .155 | .012 | 12.629 | .000 | .133 | .182 |
| site * wafer | Variance | .104 | .023 | 4.586 | .000 | .068 | .159 |

a. Dependent Variable: current.

Figure 29

*Example 8: Fitting random effect*subject interaction using GLM*

Command syntax:

```
GLM CURRENT BY WAFER SITE WITH VOLTAGE
/RANDOM = WAFER
/METHOD = SSTYPE(3)
/DESIGN = SITE SITE*WAFER VOLTAGE.
```

Output:

**Tests of Between-Subjects Effects**

Dependent Variable: current

| Source | | Type III Sum of Squares | df | Mean Square | F | Sig. |
|---|---|---|---|---|---|---|
| Intercept | Hypothesis | 2229.645 | 1 | 2229.645 | 0467.974 | .000 |
| | Error | 70.246 | 329.796 | .213[a] | | |
| site | Hypothesis | 5.371 | 7 | .767 | 1.140 | .348 |
| | Error | 48.462 | 72 | .673[b] | | |
| wafer * site | Hypothesis | 48.462 | 72 | .673 | 4.329 | .000 |
| | Error | 49.600 | 319 | .155[c] | | |
| voltage | Hypothesis | 11916.369 | 1 | 1916.369 | 6639.444 | .000 |
| | Error | 49.600 | 319 | .155[c] | | |

a. .111 MS(wafer * site) + .889 MS(Error)

b. MS(wafer * site)

c. MS(Error)

Figure 30

**Expected Mean Squares [a,b]**

| Source | Variance Component | | Quadratic Term |
|---|---|---|---|
| | Var(wafer * site) | Var(Error) | |
| Intercept | .556 | 1.000 | Intercept, site |
| site | 5.000 | 1.000 | site |
| wafer * site | 5.000 | 1.000 | |
| voltage | .000 | 1.000 | voltage |
| Error | .000 | 1.000 | |

a. For each source, the expected mean square equals the sum of the coefficients in the cells times the variance components, plus a quadratic term involving effects in the Quadratic Term cell.

b. Expected Mean Squares are based on the Type III Sums of Squares.

Figure 31

Since the design is balanced, the results of GLM and MIXED in examples 7 and 8 match. This is similar to examples 4 and 5. We see from the results of Type III tests that "voltage" is still an important predictor of "current," while "site" is not. The mean currents at different sites are thus not significantly different from each other, so we can use a simpler model without the fixed effect "site." We should still, however, consider a random-effects model, because ignoring the subject variation may lead to incorrect standard error estimates of fixed effects or false significant tests.

Up to this point, we examined primarily the similarities between GLM and MIXED. MIXED, in fact, has a much more flexible way of modeling random effects. Using the SUBJECT and COVTYPE options, Example 9 presents an equivalent form of Example 7.

*Example 9: Fitting random effect*subject interaction using SUBJECT specification*

Command syntax:

```
MIXED CURRENT BY SITE WITH VOLTAGE
/FIXED SITE VOLTAGE |SSTYPE(3)
/RANDOM SITE | SUBJECT(WAFER) COVTYPE(ID).
```

The SUBJECT option tells MIXED that each subject will have its own set of random parameters for the random effect "site." The COVTYPE option will specify the form of the variance covariance matrix of the random parameters within one subject. The command syntax attempts to specify the distributional assumption in a multivariate form, which can be written as:

$$\begin{pmatrix}\gamma_{1,1}\\ \vdots \\ \gamma_{1,8}\end{pmatrix}\begin{pmatrix}\gamma_{2,1}\\ \vdots \\ \gamma_{2,8}\end{pmatrix}\ldots,\begin{pmatrix}\gamma_{10,1}\\ \vdots \\ \gamma_{10,8}\end{pmatrix} \text{ are iid } N\left[\begin{pmatrix}0\\0\\ \vdots \\0\end{pmatrix},\begin{pmatrix}\sigma_g^2 & 0 & \cdots & 0\\ 0 & \sigma_g^2 & \cdots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \cdots & \sigma_g^2\end{pmatrix}\right]$$

Figure 32

Under normality, this assumption is equivalent to that in Example 7. One advantage of the multivariate form is that you can easily specify other covariance structures by using the COVTYPE option. The flexibility in specifying covariance structures helps us to fit a model that better describes the data. If, for example, we believe that the variances of different sites are different, we can specify a diagonal matrix as covariance type and the assumption becomes:

$$\begin{pmatrix}\gamma_{1,1}\\ \vdots \\ \gamma_{1,8}\end{pmatrix}\begin{pmatrix}\gamma_{2,1}\\ \vdots \\ \gamma_{2,8}\end{pmatrix}\ldots,\begin{pmatrix}\gamma_{10,1}\\ \vdots \\ \gamma_{10,8}\end{pmatrix} \text{ are iid } N\left[\begin{pmatrix}0\\0\\ \vdots \\0\end{pmatrix},\begin{pmatrix}\sigma_1^2 & 0 & \cdots & 0\\ 0 & \sigma_2^2 & \cdots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \cdots & \sigma_8^2\end{pmatrix}\right]$$

Figure 33

The result of fitting the same model using this assumption is given in Example 10.

*Example 10: Using COVTYPE in a random-effects model*

Command syntax:

```
MIXED CURRENT BY SITE WITH VOLTAGE
/FIXED SITE VOLTAGE |SSTYPE(3)
/RANDOM SITE | SUBJECT(WAFER) COVTYPE(DIAG)
/PRINT G TESTCOV.
```

Output:

**Type III Tests of Fixed Effects a**

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 311.313 | 0467.974 | .000 |
| site | 7 | 16.310 | 1.267 | .325 |
| voltage | 1 | 319.000 | 6639.444 | .000 |

a. Dependent Variable: current.

Figure 34

**Estimates of Covariance Parameters** [a]

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Residual | | .155 | .012 | 12.629 | .000 | .133 | .182 |
| site [subject = wafer] | Var: [site=1] | .136 | .079 | 1.726 | .084 | .044 | .424 |
| | Var: [site=2] | .096 | .060 | 1.599 | .110 | .028 | .326 |
| | Var: [site=3] | .183 | .101 | 1.812 | .070 | .062 | .539 |
| | Var: [site=4] | .119 | .071 | 1.681 | .093 | .037 | .382 |
| | Var: [site=5] | .071 | .048 | 1.475 | .140 | .019 | .269 |
| | Var: [site=6] | .073 | .049 | 1.484 | .138 | .019 | .272 |
| | Var: [site=7] | .030 | .029 | 1.046 | .296 | .005 | .198 |
| | Var: [site=8] | .120 | .071 | 1.685 | .092 | .038 | .385 |

a. Dependent Variable: current.

Figure 35

**Random Effect Covariance Structure (G)** [a]

| | [site=1] \| wafer | [site=2] \| wafer | [site=3] \| wafer | [site=4] \| wafer | [site=5] \| wafer | [site=6] \| wafer | [site=7] \| wafer | [site=8] \| wafer |
|---|---|---|---|---|---|---|---|---|
| [site=1] \| wafer | .136 | .000 | .000 | .000 | .000 | .000 | .000 | .000 |
| [site=2] \| wafer | .000 | .096 | .000 | .000 | .000 | .000 | .000 | .000 |
| [site=3] \| wafer | .000 | .000 | .183 | .000 | .000 | .000 | .000 | .000 |
| [site=4] \| wafer | .000 | .000 | .000 | .119 | .000 | .000 | .000 | .000 |
| [site=5] \| wafer | .000 | .000 | .000 | .000 | .071 | .000 | .000 | .000 |
| [site=6] \| wafer | .000 | .000 | .000 | .000 | .000 | .073 | .000 | .000 |
| [site=7] \| wafer | .000 | .000 | .000 | .000 | .000 | .000 | .030 | .000 |
| [site=8] \| wafer | .000 | .000 | .000 | .000 | .000 | .000 | .000 | .120 |

Diagonal

a. Dependent Variable: current.

Figure 36

In Example 10, we request one extra table, the estimated covariance matrix of the random effect "site." It is an eight-by-eight diagonal matrix in this case. Note that changing the covariance structure of a random effect also changes the estimates and tests of fixed effects. We want, in practice, an objective method to select suitable covariance structures for our random effects. In the section "Covariance Structure Selection," we revisit examples 9 and 10 to show how to select covariance structures for random effects.

### Multilevel analysis

The use of the SUBJECT and COVTYPE options in /RANDOM and /REPEATED brings many options for modeling the covariance structures of random effects and residual errors. It is particularly useful when modeling data obtained from a hierarchy. Example 11 illustrates the simultaneous use of these options in a multilevel model. We selected data from six schools from the Junior School Project of Mortimore, et al. (1988). We investigate below how the socioeconomic status (SES) of a student affects his or her math scores over a three-year period.

*Example 11: Multilevel mixed-effects model*

Command syntax:

```
MIXED MATHTEST BY SCHOOL CLASS STUDENT GENDER SES SCHLYEAR
/FIXED GENDER SES SCHLYEAR SCHOOL
/RANDOM SES |SUBJECT(SCHOOL*CLASS) COVTYPE(ID)
/RANDOM SES |SUBJECT(SCHOOL*CLASS*STUDENT) COVTYPE(ID)
/REPEATED SCHLYEAR | SUBJECT(SCHOOL*CLASS*STUDENT) COVTYPE(AR1)
/PRINT SOLUTION TESTCOV.
```

Output:

**Type III Tests of Fixed Effects** [a]

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 15.332 | 1076.489 | .000 |
| gender | 1 | 134.839 | .979 | .324 |
| ses | 2 | 16.815 | 3.888 | .041 |
| schlyear | 2 | 202.538 | 55.376 | .000 |
| school | 5 | 13.120 | .872 | .525 |

a. Dependent Variable: Math test.

Figure 37

**Estimates of Fixed Effects** [b]

| Parameter | Estimate | Std. Error | df | t | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Intercept | 29.097 | 2.184 | 20.036 | 13.324 | .000 | 24.542 | 33.652 |
| [gender=0] | -1.026 | 1.037 | 134.839 | -.989 | .324 | -3.077 | 1.025 |
| [gender=1] | .000[a] | .000 | . | . | . | . | . |
| [ses=1.00] | 5.803 | 2.331 | 21.478 | 2.490 | .021 | .963 | 10.644 |
| [ses=2.00] | .304 | 1.782 | 13.877 | .170 | .867 | -3.522 | 4.129 |
| [ses=3.00] | .000[a] | .000 | . | . | . | . | . |
| [schlyear=0] | -4.377 | .457 | 118.116 | -9.575 | .000 | -5.282 | -3.471 |
| [schlyear=1] | -4.126 | .468 | 219.557 | -8.825 | .000 | -5.047 | -3.204 |
| [schlyear=2] | .000[a] | .000 | . | . | . | . | . |
| [school=1] | -2.751 | 2.405 | 12.873 | -1.144 | .274 | -7.952 | 2.450 |
| [school=2] | -.784 | 2.865 | 18.557 | -.274 | .787 | -6.792 | 5.223 |
| [school=3] | 2.269 | 2.645 | 14.518 | .858 | .405 | -3.385 | 7.923 |
| [school=4] | -1.911 | 2.811 | 9.329 | -.680 | .513 | -8.236 | 4.415 |
| [school=5] | -.686 | 2.545 | 15.575 | -.270 | .791 | -6.092 | 4.720 |
| [school=6] | .000[a] | .000 | . | . | . | . | . |

a. This parameter is set to zero because it is redundant.

b. Dependent Variable: Math test.

Figure 38

**Estimates of Covariance Parameters** [a]

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Repeated Measures | AR1 diagonal | 12.686 | 1.667 | 7.609 | .000 | 9.805 | 16.413 |
| | AR1 rho | -.027 | .142 | -.190 | .850 | -.296 | .246 |
| ses [subject = school | Variance | 6.450 | 4.991 | 1.292 | .196 | 1.415 | 29.391 |
| ses [subject = school | Variance | 30.409 | 4.782 | 6.358 | .000 | 22.342 | 41.387 |

a. Dependent Variable: Math test.

Figure 39

In Example 11, the goal is to discover whether socioeconomic status (“ses”) is an important predictor for mathematics achievement (“mathtest”). To do so, we use the factor “ses” as a fixed effect. We also want to adjust for the possible sampling variation due to different classes and students. “Ses” is therefore also used twice as a random effect. The first random effect tries to adjust for the variation of the “ses” effect owing to class variation. In order to identify all classes in the dataset, school*class is specified in the SUBJECT option. The second random effect also tries to adjust for the variation of the “ses” effect owing to student variation. The subject specification is thus school*class*student. All of the students are followed for three years; the school year (“schlyear”) is therefore used as a fixed effect to adjust for possible trends in this period. The /REPEATED subcommand is also used to model the possible correlation of the residual errors within each student.

We have a relatively small dataset. Since there are only six schools, we can only use school as a fixed effect while adjusting for possible differences between schools. In this example, there is only one random effect at each level. With SPSS 11.5 or later, you can specify more than one random effect in MIXED. If multiple random effects are specified on the same RANDOM subcommand, you can model their correlation by using a suitable COVTYPE specification. If the random effects are specified on separate RANDOM subcommands, they are assumed to be independent.

In the Type III tests of fixed effects, in Example 11, we see that socioeconomic status does impact student performance. The parameter estimates of “ses” for students with “ses=1” (fathers have managerial or professional occupations) indicate that these students perform better than students at other socioeconomic levels. The effect “schlyear” is also significant in the model and the students’ performances increase with “schlyear.”

From “Estimates of Covariance Parameters” (Figure 39), we notice that the estimate of the “AR1 rho” parameter is not significant, which means that a simple, scaled-identity structure may be used. For the variation of “ses” due to school* class, the estimate is very small compared to other sources of variance and the Wald test indicates that it is not significant. We can therefore consider removing the random effect from the model.

We see from this example that the major advantages of MIXED are that it is able to look at different aspects of a dataset simultaneously and that all of the statistics are already adjusted for all effects in the model. Without MIXED, we must use different tools to study different aspects of the models. An example of this is using GLM to study the fixed effects and using VARCOMP to study the covariance structure. This is not only time consuming, but the assumptions behind the statistics are usually violated.

### Custom hypothesis tests

Apart from predefined statistics, MIXED allows users to construct custom hypotheses on fixed- and random-effects parameters through the use of the /TEST subcommand. To illustrate, we use a dataset from Pinheiro and Bates (2000). The data consist of a CT scan on a sample of ten dogs. The dogs’ left and right lymph nodes were scanned and the intensity of each scan was recorded in the variable pixel. The following mixed-model command syntax tests whether there is a difference between the left and right lymph nodes.

*Example 12: Custom hypothesis testing in mixed-effects model*

Command syntax:

```
MIXED PIXEL BY SIDE
/FIXED SIDE
/RANDOM SIDE | SUBJECT(DOG) COVTYPE(UN)
/TEST(0) ‘Side (fixed)’ SIDE 1 -1
/TEST(0) ‘Side (random)’ SIDE 1 -1 | SIDE 1 -1
/PRINT LMATRIX.
```

Output:

**Contrast Coefficients** [a]

| | | L1 |
|---|---|---|
| Fixed Effects | Intercept | 0 |
| | [side=L] | 1 |
| | [side=R] | -1 |
| Random Effects | [side=L] \| dog | 0 |
| | [side=R] \| dog | 0 |

a. Side (fixed)

Figure 40

Contrast Estimates [a,b]

| Contrast | Estimate | Std. Error | df | Test Value | t | Sig. | 95% Confidence Interval | |
|---|---|---|---|---|---|---|---|---|
| | | | | | | | Lower Bound | Upper Bound |
| L1 | 8.502 | 7.337 | 7.898 | .000 | 1.159 | .280 | -8.454 | 25.458 |

a. Side (fixed)

b. Dependent Variable: pixel.

Figure 41

Contrast Estimates [a,b]

| Contrast | Estimate | Std. Error | df | Test Value | t | Sig. | 95% Confidence Interval | |
|---|---|---|---|---|---|---|---|---|
| | | | | | | | Lower Bound | Upper Bound |
| L1 | 8.502 | 3.205 | 86.681 | .000 | 2.653 | .009 | 2.131 | 14.873 |

a. Side (random)

b. Dependent Variable: pixel.

Figure 42

The output of the two /TEST subcommands is shown above. The first test looks at differences in the left and right sides in the general population (broad inference space). We should use the second test to test the differences between the left and right sides for the sample of dogs used in this particular study (narrow inference space). In the second test, the average differences of the random effects over the ten dogs are added to the statistics. MIXED automatically calculates the average over subjects. Note that the contrast coefficients for random effects are scaled by one/(number of subjects). Though the average difference for the random effect is zero, it affects the standard error of the statistic. We see that statistics of the two tests are the same, but the second has a smaller standard error. This means that if we make an inference on a larger population, there will be more uncertainty. This is reflected in the larger standard error of the test. The hypothesis in this example is not significant in the general population, but it is significant for the narrow inference. A larger sample size is therefore often needed to test a hypothesis about the general population.

### Covariance structure selection

In examples 3 and 11, we see the use of Wald statistics in covariance structure selection. Another approach to testing hypotheses on covariance parameters uses likelihood ratio tests. The statistics are constructed by taking the differences of the -2 Log likelihoods of two nested models. Under the null hypothesis that the covariance parameters are 0 in the population, this difference follows a chi-squared distribution with degrees of freedom equal to the difference in the number of parameters of the models.

To illustrate the use of the likelihood ratio test, we again look at the model in examples 9 and 10. In Example 9, we use a scaled identity as the covariance matrix of the random effect “site.” In Example 10, however, we use a diagonal matrix with unequal diagonal elements. Our goal is to discover which model better fits the data. We obtain the -2 Log likelihood values and other criteria about the two models from the information criteria tables shown on the next page.

*Information criteria for Example 9*

| Information Criteria [a] | |
|---|---|
| -2 Restricted Log Likelihood | 523.532 |
| Akaike's Information Criterion (AIC) | 527.532 |
| Hurvich and Tsai's Criterion (AICC) | 527.563 |
| Bozdogan's Criterion (CAIC) | 537.469 |
| Schwarz's Bayesian Criterion (BIC) | 535.469 |

The information criteria are displayed in smaller-is-better forms.

a. Dependent Variable: current.

Figure 43

| Information Criteria [a] | |
|---|---|
| -2 Restricted Log Likelihood | 519.290 |
| Akaike's Information Criterion (AIC) | 537.290 |
| Hurvich and Tsai's Criterion (AICC) | 537.763 |
| Bozdogan's Criterion (CAIC) | 582.009 |
| Schwarz's Bayesian Criterion (BIC) | 573.009 |

The information criteria are displayed in smaller-is-better forms.

a. Dependent Variable: current.

Figure 44

The likelihood ratio test statistic for testing Example 9 (null hypothesis) versus Example 10 is 523.532 - 519.290 = 4.242. This statistic has a chi-squared distribution and the degrees of freedom are determined by the difference (seven) in the number of parameters in the two models. The p-value of this statistic is 0.752, which is not significant at the 0.05 level. The likelihood ratio test indicates, therefore, that we may use the simpler model in Example 9.

Apart from Wald statistics and likelihood ratio tests, we can also use such information criteria as Akaike's Information Criterion (AIC) and Schwarz's Bayesian Criterion (BIC) to search for the best model.

## Random coefficient models

In many situations, it is impossible to use a single regression line to describe the behavior of every individual. To account for possible variations between individuals, we can treat the regression coefficients as random variables. This type of model is therefore called the random coefficient model. We typically assume that the regression coefficients have normal distributions. Here we have a dataset that was used by Willet (1988) and Singer (1998) as illustration. The data are the performances of 35 individuals in an opposite-naming task on four consecutive occasions. The performance profiles of the 35 individuals are shown in the following graph.

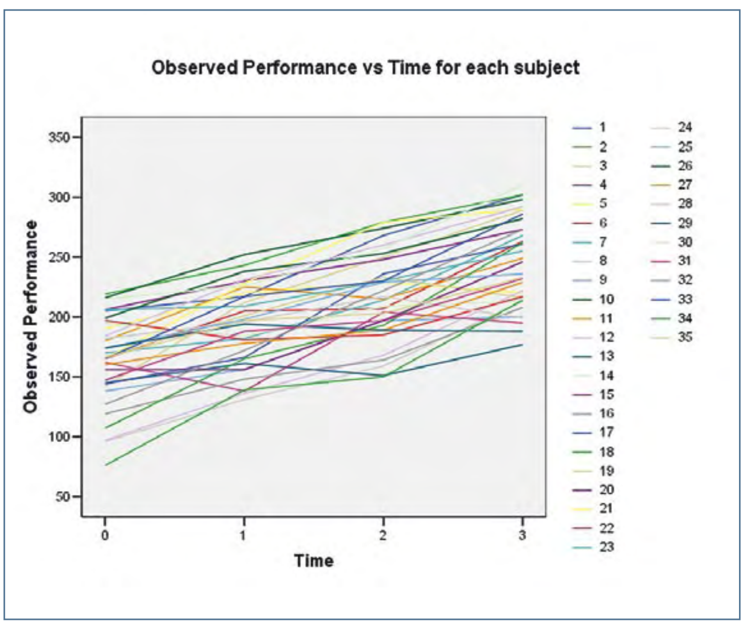


Figure 45

We can see that most individuals exhibit an increasing trend over time. Since a single regression line will not fit all of them, it makes sense to use a random coefficient model. If we restrict ourselves to linear models, there are three possible model types:

- Random intercept
- Random slopes
- Random intercept and slopes

**Random intercept models**

As the name suggests, random intercept models assume that each individual has a different intercept. In this model, we assume that the intercepts have an iid normal distribution with a mean of zero and some unknown variance.

*Example 13: Random intercept models*

Command syntax:

```
MIXED Y WITH TIME
/FIXED INTERCEPT TIME
/RANDOM INTERCEPT | SUBJECT(ID) COVTYPE(ID)
/PRINT SOLUTION TESTCOV.
```

Output:

Estimates of Fixed Effects [a]

| Parameter | Estimate | Std. Error | df | t | Sig. | 95% Confidence Interval | |
|---|---|---|---|---|---|---|---|
| | | | | | | Lower Bound | Upper Bound |
| Intercept | 164.374 | 5.777 | 46.060 | 28.455 | .000 | 152.747 | 176.002 |
| time | 26.960 | 1.466 | 104.000 | 18.395 | .000 | 24.054 | 29.866 |

a. Dependent Variable: Performance.

Figure 46

Estimates of Covariance Parameters [a]

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval | |
|---|---|---|---|---|---|---|---|
| | | | | | | Lower Bound | Upper Bound |
| Residual | | 375.901 | 52.128 | 7.211 | .000 | 286.440 | 493.303 |
| Intercept [subject = id | Variance | 904.805 | 242.590 | 3.730 | .000 | 534.979 | 530.289 |

a. Dependent Variable: Performance.

Figure 47

The coefficients you see in the Estimates of Fixed Effects table are the estimated population regression line. Since we are using a random intercept model, MIXED automatically estimates the variance of the random intercepts. The estimates are found in the Estimates of Covariance Parameters table. The estimated variance of the intercept is about 904.805, which suggests that different individuals have different intercepts.

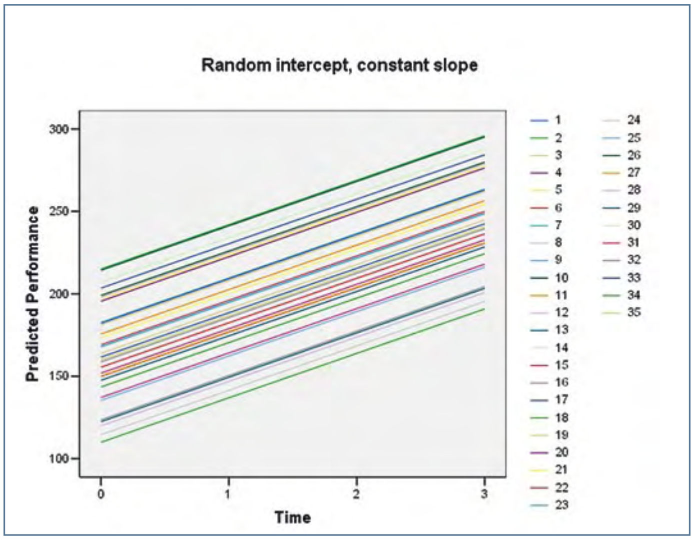


Figure 48

**Random slopes models**

Analogous to a random intercept model, a random slopes model assumes that each individual has a different slope. In this model, we assume that the slopes have an iid normal distribution with a mean of zero and an unknown variance.

*Example 14: Random slopes models*

Command syntax:

```
MIXED Y WITH TIME
/FIXED INTERCEPT TIME
/RANDOM TIME | SUBJECT(ID) COVTYPE(ID)
/PRINT SOLUTION TESTCOV.
```

Output:

Estimates of Fixed Effects [a]

| Parameter | Estimate | Std. Error | df | t | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Intercept | 164.374 | 3.793 | 104.000 | 43.336 | .000 | 156.853 | 171.896 |
| time | 26.960 | 2.941 | 66.286 | 9.166 | .000 | 21.088 | 32.832 |

a. Dependent Variable: Performance.

Figure 49

Estimates of Covariance Parameters[a]

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Residual | | 719.345 | 99.755 | 7.211 | .000 | 548.147 | 944.013 |
| time [subject = id] | Variance | 158.946 | 51.507 | 3.086 | .002 | 84.220 | 299.975 |

a. Dependent Variable: Performance.

Figure 50

As with the random intercepts model, MIXED provides the estimated population regression line in the Estimates of Fixed Effects table, and the variance of the random slopes in the Estimates of Covariance Parameters table. The estimated variance of the random slopes is 158.946, which is highly significant.

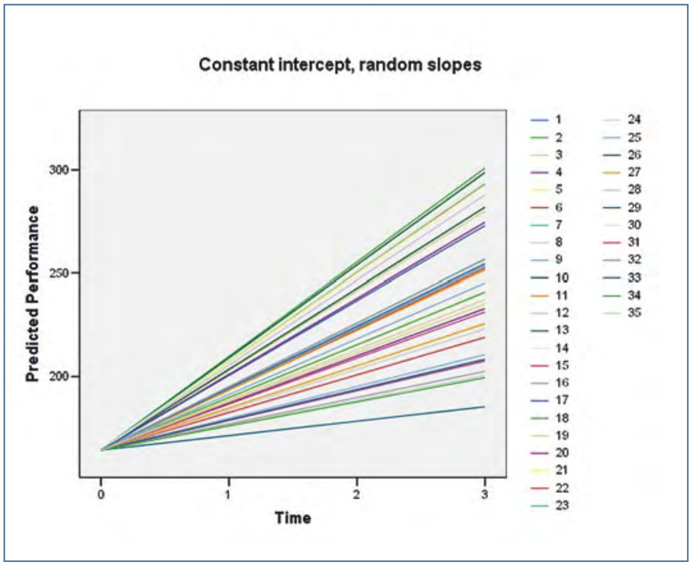


Figure 51

In comparing the predicted profile plots in examples 13 and 14 to the observed profile plots, we notice that neither the random intercepts nor the random slopes model can completely explain the variations in the data. We therefore need to consider a more complicated model that has both random intercepts and random slopes.

**Random intercepts and slopes models**

When both intercepts and slopes are random, MIXED has more flexibility in modeling the data. In this model, pairs of intercepts and slopes are assumed to have iid bivariate normal distribution with a mean of zero and some unknown covariance matrix.

*Example 15: Random intercepts and slopes models*

Command syntax:

```
MIXED Y WITH TIME
/FIXED INTERCEPT TIME
/RANDOM INTERCEPT TIME | SUBJECT(ID) COVTYPE(UN)
/PRINT SOLUTION TESTCOV.
```

Output:

**Estimates of Covariance Parameters** [a]

| Parameter | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|
| Residual | 719.345 | 99.755 | 7.211 | .000 | 548.147 | 944.013 |
| time [subject = id] Variance | 158.946 | 51.507 | 3.086 | .002 | 84.220 | 299.975 |

a. Dependent Variable: Performance.

Figure 52

**Estimates of Covariance Parameters** [a]

| Parameter | | Estimate | Std. Error | Wald Z | Sig. | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|---|---|
| Residual | | 159.477 | 26.957 | 5.916 | .000 | 114.504 | 222.115 |
| Intercept + time [subject = id] | UN (1,1) | 1198.777 | 318.381 | 3.765 | .000 | 712.310 | 2017.472 |
| | UN (2,1) | -179.256 | 88.963 | -2.015 | .044 | 353.621 | -4.890 |
| | UN (2,2) | 132.401 | 40.211 | 3.293 | .001 | 73.009 | 240.106 |

a. Dependent Variable: Performance.

Figure 53

In addition to estimating the population regression line, MIXED also estimates the variance of the intercepts, the variance of the slopes, and the covariance between the intercepts and the slopes. All of the variance and covariance parameters in this model are significant at the 0.05 level. We can see that the predicted profiles of the 35 individuals as shown below match the observed profile much better than the profiles produced by the previous two models.

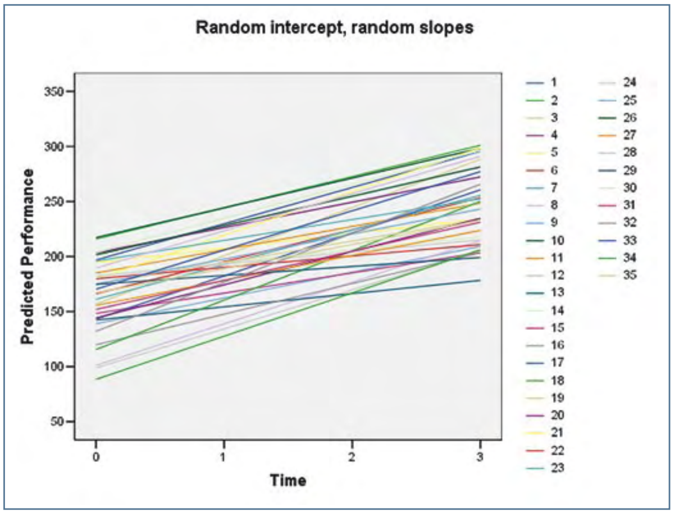


Figure 54

If we compare the AIC of the three random coefficient models, we see that the random intercepts and slopes model has the smallest AIC. It is therefore the best model of the three.

| Model | AIC |
|---|---|
| Random intercept | 1304.340 |
| Random slope | 1361.464 |
| Random intercept and slope | 1274.823 |

Figure 55

### Estimated marginal means

Estimated marginal means (EMMEANS) are also known as modified population marginal means or predicted means. In most cases, they are also the same as least squares means, which are group means that are estimated from the fitted model. In general, they are preferred over observed means, which do not account for the underlying model of your data. In SPSS for Windows, there are two ways to compute EMMEANS. The first method is to spell out the contrast matrix directly and use MIXED's /TEST subcommand to compute them. This is a laborious task, however, and prone to errors. The /EMMEANS subcommand is therefore introduced to simplify the calculations. To illustrate, we apply the method to an example dataset containing salary and demographic information for 474 individuals.

In the following example, we fit a fixed-effects model that predicts employee salary by using gender, minority group membership, job classification, and education as predictors. Based on the model, we would like to find the predicted salary for each job category.

*Example 16: EMMEANS*

Command syntax:

```
MIXED SALARY BY GENDER MINORITY JOBCAT WITH EDUC
/FIXED GENDER MINORITY JOBCAT EDUC
/PRINT LMATRIX SOLUTION
/EMMEANS = TABLES(JOBCAT) COMPARE ADJ(SIDAK).
```

Output:

**Type III Tests of Fixed Effects [a]**

| Source | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Intercept | 1 | 468.000 | 50.475 | .000 |
| gender | 1 | 468.000 | 33.730 | .000 |
| minority | 1 | 468.000 | 4.053 | .045 |
| jobcat | 2 | 468.000 | 189.799 | .000 |
| educ | 1 | 468.000 | 56.511 | .000 |

a. Dependent Variable: Current Salary.

Figure 56

**Estimates [b]**

| Employment Category | Mean | Std. Error | df | 95% Confidence Interval Lower Bound | 95% Confidence Interval Upper Bound |
|---|---|---|---|---|---|
| Clerical | 8599.033[a] | 553.351 | 468.000 | 27511.673 | 29686.393 |
| Custodial | 2998.593[a] | 1961.688 | 468.000 | 29143.787 | 36853.399 |
| Manager | 5338.089[a] | 1309.287 | 468.000 | 52765.280 | 57910.899 |

a. Covariates appearing in the model are evaluated at the following values: Educational Level (years) = 13.49.

b. Dependent Variable: Current Salary.

Figure 57

All the effects are significant at the 0.05 level, therefore it's logical to try to discover the mean salary of a particular demographic group and compare it to that of other groups. The EMMEANS subcommand can help to answer these types of questions. If you specify the option TABLES(JOBCAT) on an EMMEANS subcommand, it computes the predicted mean of each job category using the fitted model. In general, these predicted means are different from the observed cell means. The output is shown in the Estimates table (Figure 57). It shows that managers have the highest average salary ($55,338) and clerks have the lowest average salary ($28,599).

In order to discover whether salaries in different job categories are significantly different from each other, you can use the COMPARE option to instruct MIXED to perform all pairwise comparisons among all job categories. If you only want to compare categories to a reference category, you can use the optional keyword REFCAT to specify the reference category. The ADJ(SIDAK) option will instruct MIXED to use the Sidak multiple tests adjustment when calculating p-values. The results are shown in the Pairwise Comparisons table (Figure 58). The p-values suggest that all pairs are significant at the 0.05 level, except the comparison between the clerical group and the custodial group. The COMPARE option also performs a univariate test to discover whether the means of all job categories are equal. In this example, the univariate test's p-value is less than 0.05, so we reject the null hypothesis of equal category means.

The previous example is relatively simple. Next, we will illustrate the use of EMMEANS in a more sophisticated model that is similar to Example 11. The model we are going to use is essentially the same as the one used in Example 11, but with the addition of the GENDER*SES interaction.

**Pairwise Comparisons** [b]

| (I) Employment Category | (J) Employment Category | Mean Difference (I-J) | Std. Error | df | Sig.[a] | 95% Confidence Interval for Difference[a] Lower Bound | Upper Bound |
|---|---|---|---|---|---|---|---|
| Clerical | Custodial | -4399.561 | 1999.853 | 468.000 | .083 | -9191.838 | 392.717 |
| | Manager | -26739.057* | 1389.206 | 468.000 | .000 | -30068.032 | -23410.081 |
| Custodial | Clerical | 4399.561 | 1999.853 | 468.000 | .083 | -392.717 | 9191.838 |
| | Manager | -22339.496* | 2483.022 | 468.000 | .000 | -28289.597 | -16389.395 |
| Manager | Clerical | 26739.057* | 1389.206 | 468.000 | .000 | 23410.081 | 30068.032 |
| | Custodial | 22339.496* | 2483.022 | 468.000 | .000 | 16389.395 | 28289.597 |

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Sidak.

b. Dependent Variable: Current Salary.

Figure 58

**Univariate Tests** [a]

| Numerator df | Denominator df | F | Sig. |
|---|---|---|---|
| 2 | 468.000 | 189.799 | .000 |

The F tests the effect of Employment Category. This te
is based on the linearly independent pairwise
comparisons among the estimated marginal means.

a. Dependent Variable: Current Salary.

Figure 59

*Example 17: EMMEANS*

Command syntax:

```
MIXED MATHTEST BY SCHOOL CLASS STUDENT
    GENDER SES SCHLYEAR
/FIXED GENDER SES GENDER*SES SCHLYEAR SCHOOL
/RANDOM SES |SUBJECT(SCHOOL*CLASS) COVTYPE(ID)
/RANDOM SES |SUBJECT(SCHOOL*CLASS*STUDENT) COVTYPE(ID)
/REPEATED SCHLYEAR | SUBJECT(SCHOOL*CLASS*STUDENT)
    COVTYPE(AR1)
/PRINT SOLUTION TESTCOV
/EMMEAN TABLE(SES*GENDER) COMPARE(SES) ADJ(SIDAK).
```

**Estimates** [a]

| Socioeconomic Status | Gender | Mean | Std. Error | df | 95% Confidence Interval Lower Bound | Upper Bound |
|---|---|---|---|---|---|---|
| I & II | Boy | 31.287 | 2.271 | 44.353 | 26.710 | 35.864 |
| | Girl | 30.800 | 2.088 | 34.606 | 26.559 | 35.040 |
| III to VI | Boy | 25.211 | 1.283 | 19.171 | 22.528 | 27.895 |
| | Girl | 25.558 | 1.260 | 17.387 | 22.904 | 28.212 |
| Other | Boy | 23.500 | 1.676 | 31.892 | 20.085 | 26.915 |
| | Girl | 26.898 | 1.765 | 37.100 | 23.322 | 30.473 |

a. Dependent Variable: Math test.

Figure 60

Output:

**Pairwise Comparisons** [b]

| Gender | (I) Socioeconomic Status | (J) Socioeconomic Status | Mean Difference (I-J) | Std. Error | df | Sig.[a] | 95% Confidence Interval for Difference[a] Lower Bound | Upper Bound |
|---|---|---|---|---|---|---|---|---|
| Boy | I & II | III to VI | 6.076 | 2.614 | 37.195 | .075 | -.458 | 12.610 |
| | | Other | 7.788* | 2.867 | 40.812 | .029 | .650 | 14.925 |
| | III to VI | I & II | -6.076 | 2.614 | 37.195 | .075 | -12.610 | .458 |
| | | Other | 1.712 | 2.099 | 26.131 | .807 | -3.642 | 7.065 |
| | Other | I & II | -7.788* | 2.867 | 40.812 | .029 | -14.925 | -.650 |
| | | III to VI | -1.712 | 2.099 | 26.131 | .807 | -7.065 | 3.642 |
| Girl | I & II | III to VI | 5.242 | 2.440 | 29.808 | .115 | -.930 | 11.414 |
| | | Other | 3.902 | 2.783 | 37.687 | .426 | -3.051 | 10.854 |
| | III to VI | I & II | -5.242 | 2.440 | 29.808 | .115 | -11.414 | .930 |
| | | Other | -1.340 | 2.172 | 29.159 | .904 | -6.842 | 4.162 |
| | Other | I & II | -3.902 | 2.783 | 37.687 | .426 | -10.854 | 3.051 |
| | | III to VI | 1.340 | 2.172 | 29.159 | .904 | -4.162 | 6.842 |

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Sidak.

b. Dependent Variable: Math test.

Figure 61

**Univariate Tests [a]**

| Gender | Numerator df | Denominator df | F | Sig. |
|---|---|---|---|---|
| Boy | 2 | 31.715 | 3.848 | .032 |
| Girl | 2 | 30.831 | 2.308 | .116 |

Each F tests the simple effects of Socioeconomic Status within ea level combination of the other effects shown. These tests are bas on the linearly independent pairwise comparisons among the estimated marginal means.

a. Dependent Variable: Math test.

Figure 62

The command syntax in Example 17 requests the predicted means of all gender-socioeconomic status combinations. Since this model involves random effects, the predicted means are computed by averaging the random effects over subjects. The predicted means of the six gender-socioeconomic status combinations are shown in the Estimates table (Figure 60). Among these simple effects comparisons, only Socioeconomic Status I & II and Socioeconomic Status Other/Boy are significantly different from each other, with p-value 0.029.

The COMPARE(SES) option in Example 17 indicates that we want to perform a univariate test to determine whether the means of socioeconomic status are the same within each gender. The results (see Figure 62) indicate that the means of socioeconomic status among boys are significant at 0.05 level but not among girls. This agrees with the Pairwise Comparisons table (Figure 61).

## References


McCulloch, C.E., and Searle, S.R. (2000). Generalized, Linear, and Mixed Models. John Wiley and Sons.

Mortimore, P., Sammons, P., Stoll, L., Lewis, D. and Ecob, R. (1988). School Matters: the Junior Years. Wells, Open Books.

Pinheiro J.C., and Bates, D.M. (2000). Mixed-Effects Models in S and S-PLUS. Springer.

Potthoff, R.F., and Roy, S.N. (1964). “A generalized multivariate analysis of variance model useful especially for growth curve problems.” Biometrika, 51:313-326.

Singer J.D. (1998). “Using SAS PROC MIXED to fix multilevel models, hierarchical models and individual growth models.” Journal of Educational and Behavioral Statistics, 24:323-355.

Verbeke, G., and Molenberghs, G. (2000). Linear Mixed Models for Longitudinal Data. Springer.

Willett, J.B. (1989). “Questions and answers in the measurement of change.” In E.Z. Rothkopf (Ed.) Review of Research in Education, 15:345-422. Washington, DC: American Education Research Association.


## About SPSS Inc.

SPSS Inc. (NASDAQ: SPSS) is the world’s leading provider of predictive analytics software and solutions. The company’s predictive analytics technology connects data to effective action by drawing reliable conclusions about current conditions and future events. More than 250,000 public sector, academic, and commercial customers, including more than 95 percent of the Fortune 1000, rely on SPSS technology to help increase revenue, reduce costs, improve processes, and detect and prevent fraud. Founded in 1968, SPSS is headquartered in Chicago, Illinois. For additional information, please visit www.spss.com.

**To learn more, please visit www.spss.com. For SPSS office locations and telephone numbers, go to www.spss.com/worldwide.**