Chapter 221

Mixed Models – No Repeated Measures

Introduction

This specialized Mixed Models procedure analyzes data from fixed effects, factorial designs. These designs classify subjects into one or more fixed factors and have only one measurement per subject. This procedure is especially useful when you have covariates and/or unequal variances across a factor.

The **Mixed Models - General** chapter gives a comprehensive overview of this topic. We encourage you to look there for details.

Fixed Effects Models

A fixed effect (or factor) is a variable for which levels in the study represent all levels of interest, or at least all levels that are important for inference (e.g., treatment, dose, etc.). The fixed effects in the model include those factors for which means, standard errors, and confidence intervals will be estimated, and tests of hypotheses will be performed. Fixed factors may be discrete variables or continuous covariates.

The correct model for fixed effects depends on the number of fixed factors, the questions to be answered by the analysis, and the amount of data available for the analysis. When more than one fixed factor may influence the response, it is common to include those factors in the model, along with their interactions (two-way, three-way, etc.). Difficulties arise when there are not sufficient data to model the higher-order interactions. In this case, some interactions must be omitted. It is usually suggested that if you include an interaction in the model, you should also include the main effects (i.e., individual factors) involved in the interaction even if the hypothesis test for the main effects in not significant.

Covariates

Covariates are continuous measurements that are not of primary interest in the study, but potentially have an influence on the response. This procedure permits the user to make comparisons of fixed effect means at specified values of covariates. Commonly, investigators wish to make comparisons of levels of a factor at several values of covariates.

Multiple Comparisons of Fixed Effect Levels

If there is evidence that the means of a fixed factor are, it is usually of interest to perform post-hoc pair-wise comparisons of the least-squares means to further clarify those differences. It is well-known that p-value adjustments need to be made when multiple comparison tests are performed (see Hochberg and Tamhane, 1987, or Hsu, 1996, for general discussion and details of the need for multiplicity adjustment). Such adjustments are usually made to preserve the family-wise error rate (FWER), also called the experiment-wise error rate, of the group of tests. FWER is the probability of incorrectly rejecting at least one of the pair-wise tests.

Family-Wise Error Rate (FWER) Control – Bonferroni Adjustment

The Bonferroni p-value adjustment produces adjusted p-values (probability levels) for which the FWER is controlled strictly (Westfall et al, 1999). The Bonferroni adjustment is applied to all *m* unadjusted (raw) p-values (p_j) as

$$\tilde{p}_j = \min\left(mp_j, 1\right).$$

That is, each p-value is multiplied by the number of tests in the set (family), and if the result is greater than one, it is set to the maximum possible p-value of one. The Bonferroni adjustment is generally considered to be a conservative method for simultaneously comparing levels of fixed effects.

Multiple Comparisons for the Interaction of Two Main Effects

When examining a fixed effect interaction using post-hoc (or planned) multiple comparison tests, a useful method is to compare all levels of one factor at each level of the other factor. This method is termed ‘slicing’. For example, if the interaction of Factor1 and Factor2 is significant, comparing the Factor2 mean at each level of Factor1 could aid in understanding the nature of the interaction.

Multiple Comparisons for Several Covariate Levels

When more than one covariate value is specified for *Compute Means at these Values*, the number of test used in the Bonferroni adjustment can increase dramatically. The number of tests for the Bonferroni adjustment is computed as

$$\textit{Number of Tests} = \textit{Number of Comparisons per Set} \times \textit{Number of Covariate Sets}$$

As an example, suppose that an experiment has two covariates, and a single fixed treatment factor with three levels: Control, T1, and T2. If ‘All Pairs’ were selected as the comparison on the Comparisons tab, then the number of comparisons per set would be three (T1 – Control, T2 – Control, and T2 – T1). Suppose that the researcher desired to compute the hypothesis tests at two values for the first covariate and four values for the second. The number of covariate sets would be 2 × 4 = 8. Therefore, the number of tests used in the Bonferroni adjustment to conserve the overall error-rate would be 3 × 8 = 24. The raw p-value would have to be less than 0.05/24 = 0.00208 in order to declare significance at the 0.05 level.

This example illustrates that care must be taken when specifying the covariate values at which the means and analyses will be computed. As more covariate values are specified, the number of tests in the adjustment increases, which makes it more difficult to find significance.

Example 1 – Two Fixed Factors and Two Covariates

In this example, 24 males and 24 females are randomly allocated to three dose groups: low, medium, and high. The age of each subject is recorded. Their response to a certain stimulus is recorded as a pretest. Next, the assigned dose of a certain compound is administered and their response to the stimuli is measured again.

Researchers wish to investigate how the response to the stimuli is affected by the subject's age, gender, dose, and pretest score. This can be done using a two factor, two covariate mixed model. They want to allow for the possibility of a difference in variance for males versus females.

This example will run all reports and plots so that they may be documented. Usually, only a subset of the reports would be generated. Here is an excerpt of the dataset

2 Factor 2 Covariate Dataset

Test	Dose	Gender	PreTest	Age
64.1	Low	F	52.1	58
69.1	Low	F	56.4	57
69.5	Low	F	53.5	50
81.8	Low	F	72.5	30
82.4	Low	F	65.3	72
86.7	Low	F	73.6	24
87	Low	F	69.2	30
93.8	Low	F	78.5	61
63.8	Low	M	52.4	42
66.1	Low	M	55.9	53
67.1	Low	M	56.5	36
74.7	Low	M	62.6	44
.	.	.	.	.
.	.	.	.	.
.	.	.	.	.
99	High	M	84.6	74
99.5	High	M	76	70
99.5	High	M	77.2	36

Setup

To run this example, complete the following steps:

1 Open the 2 Factor 2 Covariate example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **2 Factor 2 Covariate** and click **OK**.

2 Specify the Mixed Models – No Repeated Measures procedure options

- Find and open the **Mixed Models – No Repeated Measures** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 1** settings template. To load this template, click **Open Example Template** in the Help Center or File menu.

Option	Value
Variables Tab	
Response	**Test**
Number	**2**
Variable 1	**Dose**
Comparison	**All Pairs**
Variable 2	**Gender**
$\neq\sigma^2$	**Checked**
Number	**2**
Variable 1	**Age**
Compute Means at these Values	**30 60**
Variable 2	**PreTest**
Compute Means at these Values	**60 80**
Terms	**1-Way**
Reports Tab	
Run Summary	**Checked**
Variance Estimates	**Checked**
Hypothesis Tests	**Checked**
L Matrices - Terms	**Unchecked**
Comparisons - Sorted by Factors	**Checked**
Comparisons - Sorted by Covariate Values	**Checked**
L Matrices - Comparisons	**Unchecked**
Means - Sorted by Factors	**Checked**
Means - Sorted by Covariate Values	**Checked**
L Matrices - LS Means	**Unchecked**
Fixed Effects Solution	**Checked**
Asymptotic VC Matrix	**Checked**
Hessian Matrix	**Checked**
Show Report Definitions	**Checked**
Plots Tab	
Means Plot(s)	**Checked**
Report Options (*in the Toolbar*)	
Variable Labels	**Column Names**

3 Run the procedure

- Click the **Run** button to perform the calculations and generate the output.

Run Summary

Run Summary

Item	Value
Likelihood Type	Restricted Maximum Likelihood
Fixed Model	AGE+PRETEST+DOSE+GENDER
Number of Subjects	48
Solution Type	Newton-Raphson
Fisher Iterations	5 of a possible 5
Newton Iterations	3 of a possible 40
Max Retries	10
Lambda	1
Log Likelihood	-118.2766
-2 Log Likelihood	236.5532
AIC (Smaller Better)	240.5532
Convergence	Normal
Run Time (Seconds)	1.185

This report provides a summary of the model and the iterations toward the maximum log likelihood.

Likelihood Type

This value indicates that restricted maximum likelihood was used rather than maximum likelihood.

Fixed Model

The model that was fit to the data.

Number of Subjects

The number of rows processed from the database.

Solution Type

The solution type is method used for finding the maximum (restricted) maximum likelihood solution. Newton-Raphson is the recommended method.

Fisher Iterations

Some Fisher-Scoring iterations are used as part of the Newton-Raphson algorithm. The '5 of a possible 5' means five Fisher-Scoring iterations were used, and five was the maximum allowed (as specified on the Maximization tab).

Newton Iterations

The '2 of a possible 40' means two Newton-Raphson iterations were used, while forty was the maximum allowed (as specified on the Maximization tab).

Max Retries

The maximum number of times that lambda was changed and new variance-covariance parameters were found during an iteration was ten. If the values of the parameters result in a negative variance, lambda is divided by two and new parameters are generated. This process continues until a positive variance occurs or until Max Retries is reached.

Lambda

Lambda is a parameter used in the Newton-Raphson process to specify the amount of change in parameter estimates between iterations. One is generally an appropriate selection. When convergence problems occur, a good remedy is to set this to 0.5.

If the values of the parameters result in a negative variance, lambda is divided by two and new parameters are generated. This process continues until a positive variance occurs or until Max Retries is reached.

Log Likelihood

This is the log of the likelihood of the data given the variance-covariance parameter estimates. When a maximum is reached, the algorithm converges.

-2 Log Likelihood

This is minus 2 times the log of the likelihood. When a minimum is reached, the algorithm converges.

AIC

The Akaike Information Criterion is used for comparing covariance structures in models. It gives a penalty for increasing the number of covariance parameters in the model.

Convergence

'Normal' convergence indicates that convergence was reached before the limit.

Run Time (Seconds)

The run time is the amount of time used to solve the problem and generate the output.

Variance Report

Variance Report

Gender	Variance σ^2	Standard Deviation σ
F	7.7742	2.7882
M	8.9230	2.9871

This section gives the variances and standard deviations of the individual groups as designated.

Gender

The name of the variable across which the variances vary.

Variance

The estimated variance of this group.

Standard Deviation

The estimated standard deviation of this group.

Term-by-Term Hypothesis Test Results

Term-by-Term Hypothesis Test Results

Model Term	F-Value	Numerator DF	Denominator DF	Prob Level
Age	2.68	1	41.5	0.1092
PreTest	371.29	1	40.6	0.0000
Dose	27.77	2	41.9	0.0000
Gender	1.14	1	41.8	0.2923

This section contains a F-test for each model term using the methods of Kenward and Roger (1997).

Model Term

This is the name of the term in the model.

F-Value

The F-Value corresponds to the L matrix used for testing this term in the model. The F-Value is based on the F approximation described in Kenward and Roger (1997).

Numerator DF

This is the numerator degrees of freedom for the corresponding term.

Denominator DF

This is the approximate denominator degrees of freedom for this comparison as described in Kenward and Roger (1997).

Prob Level

The Probability Level (or P-value) gives the strength of evidence (smaller Prob Level implies more evidence) that a term in the model has differences among its levels, or a slope different from zero in the case of covariate. It is the probability of obtaining the corresponding F-Value (or greater) if the null hypothesis of equal means (or no slope) is true.

Individual Comparison Tests – Sorted by Factors, Sorted by Covariates

Individual Comparison Tests - Sorted by Factors

Comparison/ Covariate(s)	Comparison Mean Difference	F-Value	Num DF	Denom DF	Raw Prob Level	Bonferroni Prob Level
Dose						
Age=30.0000, PreTest=60.0000		27.77	2	41.9	0.0000	0.0000 [4]
Age=30.0000, PreTest=80.0000		27.77	2	41.9	0.0000	0.0000 [4]
Age=60.0000, PreTest=60.0000		27.77	2	41.9	0.0000	0.0000 [4]
Age=60.0000, PreTest=80.0000		27.77	2	41.9	0.0000	0.0000 [4]
Dose: Low - Medium						
Age=30.0000, PreTest=60.0000	-3.5405	10.70	1	42.0	0.0021	0.0258 [12]
Age=30.0000, PreTest=80.0000	-3.5405	10.70	1	42.0	0.0021	0.0258 [12]
Age=60.0000, PreTest=60.0000	-3.5405	10.70	1	42.0	0.0021	0.0258 [12]
Age=60.0000, PreTest=80.0000	-3.5405	10.70	1	42.0	0.0021	0.0258 [12]

(report continues)

This section shows the F-tests for comparisons of the levels of the fixed terms of the model according to the methods described by Kenward and Roger (1997). The individual comparisons are grouped into subsets of the fixed model terms.

Comparison/Covariate(s)

This is the comparison being made. The first line is F-test for the overall Dose factor. On this line, the levels of dose are compared when the age is 30 and the PreTest is 60.

Comparison Mean Difference

This is the difference in the least squares means for each comparison.

F-Value

The F-Value corresponds to the L matrix used for testing this comparison. The F-Value is based on the F approximation described in Kenward and Roger (1997).

Numerator DF

This is the numerator degrees of freedom for this comparison.

Denominator DF

This is the approximate denominator degrees of freedom for this comparison as described in Kenward and Roger (1997).

Raw Prob Level

The Raw Probability Level (or Raw P-value) gives the strength of evidence for a single comparison, unadjusted for multiple testing. It is the single test probability of obtaining the corresponding difference if the null hypothesis of equal means is true.

Bonferroni Prob Level

The Bonferroni Prob Level is adjusted for multiple tests. The number of tests adjusted for is enclosed in brackets following each Bonferroni Prob Level.

Least Squares (Adjusted) Means – Sorted by Factors

Least Squares (Adjusted) Means - Sorted by Factors

Term/ Covariate(s)	Mean	Standard Error of Mean	95.0% Lower Conf. Limit for Mean	95.0% Upper Conf. Limit for Mean	DF
Intercept					
Age=30.0000, PreTest=60.0000	76.4756	0.6804	75.1015	77.8498	40.9
Age=30.0000, PreTest=80.0000	95.2350	0.8576	93.5041	96.9658	41.8
Age=60.0000, PreTest=60.0000	77.9188	0.6560	76.5949	79.2428	41.9
Age=60.0000, PreTest=80.0000	96.6782	0.9836	94.6906	98.6658	40.2
Dose = [Low]					
Age=30.0000, PreTest=60.0000	72.6032	0.9360	70.7098	74.4967	38.9
Age=30.0000, PreTest=80.0000	91.3626	1.1474	89.0454	93.6797	41.1
Age=60.0000, PreTest=60.0000	74.0464	0.8347	72.3599	75.7330	40.4
Age=60.0000, PreTest=80.0000	92.8058	1.1841	90.4162	95.1953	42.0
Dose = [Medium]					
Age=30.0000, PreTest=60.0000	76.1437	0.8637	74.4006	77.8868	42.0
Age=30.0000, PreTest=80.0000	94.9030	1.0516	92.7802	97.0259	41.6
Age=60.0000, PreTest=60.0000	77.5869	0.9406	75.6886	79.4852	41.9
Age=60.0000, PreTest=80.0000	96.3463	1.2285	93.8648	98.8277	40.8

(Report continues)

This section gives the adjusted means for the levels of each fixed factor at the specified values of the covariates.

Name
This is the level of the fixed term that is estimated on the line.

Mean
The mean is the estimated least squares (adjusted or marginal) mean at the specified value of the covariate.

Standard Error of Mean
This is the standard error of the mean.

95.0% Lower (Upper) Conf. Limit for Mean
These limits give a 95% confidence interval for the mean.

DF
The degrees of freedom used for the confidence limits are calculated using the method of Kenward and Roger (1997).

Means Plots

Means Plots

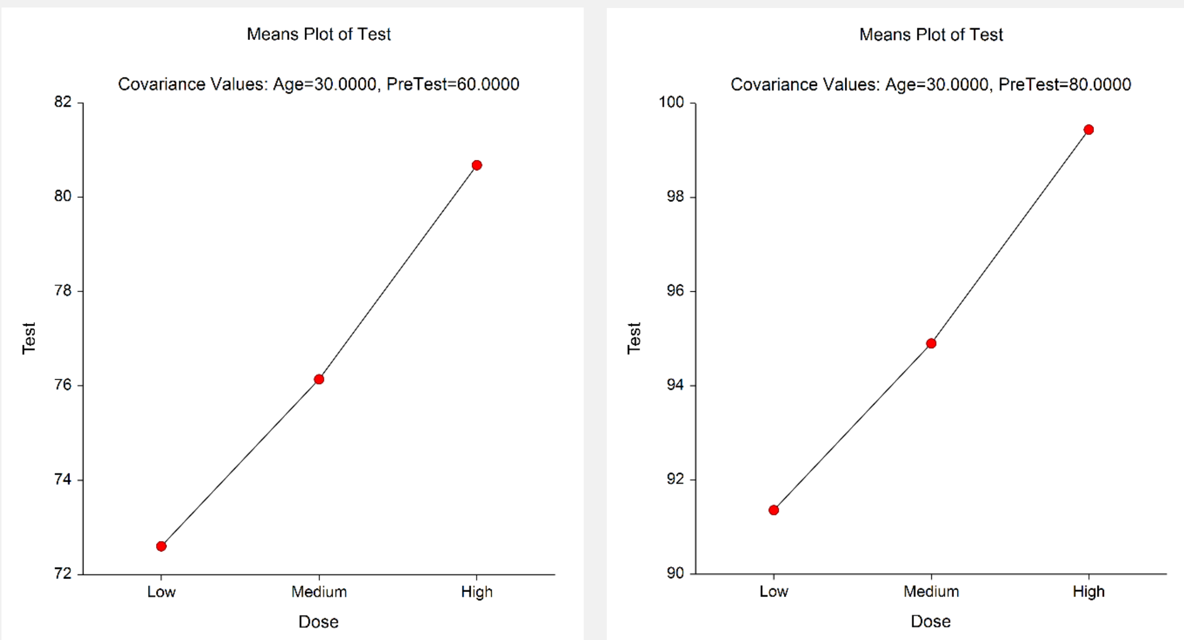

These plots show the means broken up into the categories of the fixed effects of the model.

Solution for Fixed Effects

Solution for Fixed Effects

Effect Name	Effect Estimate (Beta)	Effect Standard Error	Prob Level	95.0% Lower Conf. Limit of Beta	95.0% Upper Conf. Limit of Beta	DF	Effect No.
Intercept	22.4743	3.8707	0.0000	14.6519	30.2968	40.1	1
Age	0.0481	0.0294	0.1092	-0.0112	0.1074	41.5	2
PreTest	0.9380	0.0487	0.0000	0.8396	1.0363	40.6	3
(Dose="Low")	-8.0767	1.0866	0.0000	-10.2697	-5.8838	42.0	4
(Dose="Medium")	-4.5362	1.0817	0.0001	-6.7196	-2.3529	41.8	5
(Dose="High")	0.0000	0.0000					6
(Gender="F")	0.9687	0.9082	0.2923	-0.8643	2.8018	41.8	7
(Gender="M")	0.0000	0.0000					8

This section shows the model estimates for all the model terms (betas).

Effect Name

The Effect Name is the level of the fixed effect that is examine on the line.

Effect Estimate (Beta)

The Effect Estimate is the beta-coefficient for this effect of the model. For main effects terms the number of effects per term is the number of levels minus one. An effect estimate of zero is given for the last effect(s) of each term. There may be several zero estimates for effects of interaction terms.

Effect Standard Error

This is the standard error for the corresponding effect.

Prob Level

The Prob Level tests whether the effect is zero.

95.0% Lower (Upper) Conf. Limit of Beta

These limits give a 95% confidence interval for the effect.

DF

The degrees of freedom used for the confidence limits and hypothesis tests are calculated using the method of Kenward and Roger (1997).

Effect No.

This number identifies the effect of the line.

Asymptotic Variance-Covariance Matrix of Variance Estimates

Asymptotic Variance-Covariance Matrix of Variance Estimates

Parm	R(1,1,1)	R(1,1,2)
R(1,1,1)	5.7409	-0.0009
R(1,1,2)	-0.0009	7.6050

This section gives the asymptotic variance-covariance matrix of the variance components of the model.

Hessian Matrix of Variance Estimates

Hessian Matrix of Variance Estimates

Parm	R(1,1,1)	R(1,1,2)
R(1,1,1)	0.17418787	0.00002123
R(1,1,2)	0.00002123	0.13149167

The Hessian Matrix is directly related to the asymptotic variance-covariance matrix of the variance estimates.