Chapter 212

General Linear Models (GLM)

Introduction

This procedure performs an analysis of variance or analysis of covariance on up to ten factors using the general linear models approach. The experimental design may include up to two nested terms, making possible various repeated measures and split-plot analyses.

Because the program allows you to control which interactions are included and which are omitted, it can analyze designs with confounding such as Latin squares and fractional factorials.

Kinds of Research Questions

A large amount of research consists of studying the influence of a set of independent variables on a response (dependent) variable. Many experiments are designed to look at the influence of a single independent variable (factor) while holding other factors constant. These experiments are called single-factor experiments and are analyzed with the one-way analysis of variance (ANOVA). A second type of design considers the impact of one factor across several values of other factors. This experimental design is called the factorial design.

The factorial design is popular among researchers because it not only lets you study the individual effects of several factors in a single experiment, but it also lets you study their interaction. Interaction is present when the response variable fails to behave the same at values of one factor when a second factor is varied. Since factors seldom work independently, the study of their interaction becomes very important.

This procedure will also analyze repeated-measures and split-plot designs. These designs are popular in many disciplines in which experiments are needed that take several measurements on an individual through time. Examples are pre-post type tests administered to various groups of individuals.

Analysis of covariance (ANCOVA) is another design that may be analyzed using this procedure. ANCOVA is useful when you want to improve precision by removing various extraneous sources of variation from your study.

The Linear Model

We begin with an infinite population of individuals with many measurable characteristics. These individuals are (mentally) separated into two or more treatment populations based on one or more of these characteristics. A random sample of the individuals in each population is drawn. A treatment is applied to each individual in the sample and an outcome is measured. The data obtained are analyzed using an analysis of variance table that produces an F-test.

A mathematical model may be formulated that underlies each analysis of variance. This model expresses the response variable as the sum of parameters of the population. For example, a linear mathematical model for a two-factor experiment is

$$Y_{ijk} = m + a_i + b_j + (ab)_{ij} + e_{ijk}$$

where $i = 1, 2, …, I$; $j = 1, 2, …, J$; and $k = 1, 2, …, K$. This model expresses the value of the response variable, Y, as the sum of five components:

m the mean.

a_i the contribution of the ith level of a factor A.

b_j the contribution of the jth level of a factor B.

$(ab)_{ij}$ the combined contribution of the ith level of a factor A and the jth level of a factor B.

e_{ijk} the contribution of the kth individual. This is often called the "error."

Note that this model is the sum of various constants. This type of model is called a linear model. It becomes the mathematical basis for our discussion of the analysis of variance. Also note that this serves only as an example. Many linear models could be formulated for the two-factor experiment.

Assumptions

The following assumptions are made when using the F-test.

1. The response variable is continuous.

2. The e_{ijk} follow the normal probability distribution with mean equal to zero.

3. The variances of the e_{ijk} are equal for all values of i, j, and k.

4. The individuals are independent.

Limitations

There are few limitations when using these tests. Sample sizes may range from a few to several hundred. If your data are discrete with at least five unique values, you can assume that you have met the continuous variable assumption. Perhaps the greatest restriction is that your data comes from a random sample of the population. If you do not have a random sample, the F-test will not work.

When missing cells occur in your design, you must take special care to be sure that appropriate interaction terms are removed from the ANOVA model.

Special restrictions apply when you are running an analysis with nested terms, as in repeated measures designs. First of all, you cannot have covariates with nested terms. Second, although the sample sizes of groups (the "between" factor) may be unequal, all data must be present for each nested factor. For example, if you are running a pre-post design, you must have both pre- and post- scores for each individual. You cannot include individuals that have only one or the other.

Comparisons

Comparisons are only valid for fixed factors. A comparison is formulated in terms of the means as follows:

$$C_i = \sum_{j=1}^{J} w_{ij}\, m_j$$

In this equation, there are J levels in the factor, the means for each level of the factor are denoted m_i, and w_{ij} represents a set of J weight values for the i^{th} comparison. The comparison value, C_i, is tested using a t-test. Note that if the w_{ij} sum to zero across j, the comparison is called a "contrast" of the means.

Multiple Comparison Procedures

The multiple comparison procedures are discussed in the One-Way Analysis of Variance chapter. Note that if covariates are used, no multiple comparisons are shown.

Data Structure

The data must be entered in a format that puts the response in one variable and the values of each of the factors in other variables. An example of the data for a randomized-block design is shown next.

Randomized Block Dataset

Block	Treatment	Response
1	1	123
2	1	230
3	1	279
1	2	245
2	2	283
3	2	245
1	3	182
2	3	252
3	3	280
1	4	203
2	4	204
3	4	227

Example 1 – Running a GLM ANOVA

This section presents an example of how to run an analysis of the data presented above. These data are contained in the Randomized Block dataset.

Setup

To run this example, complete the following steps:

1   Open the Randomized Block example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select Randomized Block and click OK.

2   Specify the General Linear Models (GLM) procedure options

- Find and open the **General Linear Models (GLM)** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 1** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

Response Variable(s) ....**Response**
Factor Variable 1 ....**Block**
  Type..**Random**
Factor Variable 2 ....**Treatment**
  Type..**Fixed**
  Comparisons...**Linear**

Reports Tab

Compute Power..**Checked**
Tukey-Kramer Test...**Checked**
Show Notes ....**Checked**

3   Run the procedure

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

Expected Mean Squares

Expected Mean Squares

Model Term	DF	Term Fixed?	Denominator Term	Expected Mean Square
A: Block	2	No	S(AB)	S+bsA
B: Treatment	3	Yes	AB	S+sAB+asB
AB	6	No	S(AB)	S+sAB
S(AB)	0	No		S

Note: Expected Mean Squares are for the balanced cell-frequency case.

The expected mean square expressions are provided to show the appropriate error term for each factor. The correct error term for a factor is that term that is identical except for the factor being tested.

Model Term

The source of variation or term in the model.

DF

The degrees of freedom, which is the number of observations used by this term.

Term Fixed?

Indicates whether the term is fixed or random.

Denominator Term

Indicates the term used as the denominator in the F-ratio.

Expected Mean Square

This expression represents the expected value of the corresponding mean square if the design was completely balanced. *S* represents the expected value of the mean square error (sigma). The uppercase letters represent either the adjusted sum of squared treatment means if the factor is fixed, or the variance component if the factor is random. The lowercase letter represents the number of levels for that factor, and *s* represents the number of replications of the experimental layout.

These EMS expressions are provided to determine the appropriate error term for each factor. The correct error term for a factor is that term whose EMS is identical except for the factor being tested.

In this example, the appropriate error term for factor B is the AB interaction. The appropriate error term for AB is S (mean square error). Since there are zero degrees of freedom for S, the terms A and AB cannot be tested.

Analysis of Variance Table

Analysis of Variance Table

Model Term	DF	Sum of Squares	Mean Square	F-Ratio	P-Value	Power†
A: Block	2	10648.67	5324.333			
B: Treatment	3	4650.917	1550.306	1.0934	0.42136	0.17794
AB	6	8507.333	1417.889			
S	0	5.880629E-12				
Total (Adjusted)	11	23806.92				
Total	12					

† Power was calculated using the observed F-Ratio as the population effect size with a significance level of α = 0.05.

Model Term

The source of variation, which is the term in the model.

DF

The degrees of freedom, which is the number of observations used by the corresponding model term.

Sum of Squares

This is the sum of squares for this term. It is usually included in the ANOVA table for completeness, not for direct interpretation.

Mean Square

An estimate of the variation accounted for by this term; it is the sum of squares divided by the degrees of freedom.

F-Ratio

The ratio of the mean square for this term and the mean square of its corresponding error term. This is also called the F-test value.

P-Value

The significance level of the above F-ratio, or the probability of an F-ratio larger than that obtained by this analysis. For example, to test at an alpha of 0.05, this probability would have to be less than 0.05 to make the F-ratio significant. Note that if the value is significant at the specified value of alpha, a star is placed to the right of the F-Ratio.

Power

Power is the probability of rejecting the hypothesis that the means are equal when they are in fact not equal. Power is one minus the probability of type II error (β). The power of the test depends on the sample size, the magnitudes of the variances, the alpha level, and the actual differences among the population means.

The power value calculated here assumes that the population standard deviation is equal to the observed standard deviation and that the differences among the population means are exactly equal to the differences among the sample means.

High power is desirable. High power means that there is a high probability of rejecting the null hypothesis when the null hypothesis is false. This is a critical measure of precision in hypothesis testing.

Generally, you would consider the power of the test when you accept the null hypothesis. The power will give you some idea of what actions you might take to make your results significant. If you accept the null hypothesis with high power, there is not much left to do. At least you know that the means are not different. However, if you accept the null hypothesis with low power, you can take one or more of the following actions:

1.   Increase your alpha level. Perhaps you should be testing at alpha = 0.05 instead of alpha = 0.01. Increasing the alpha level will increase the power.

2.   Increasing your sample size will increase the power of your test if you have low power. If you have high power, an increase in sample size will have little effect.

3.   Decrease the magnitude of the variance. Perhaps you can redesign your study so that measurements are more precise and extraneous sources of variation are removed.

Means and Standard Errors and Means Plots

Means and Standard Errors

Term	Count	Mean	Standard Error
All	12	229.4167	
A: Block			
1	4	188.25	0
2	4	242.25	0
3	4	257.75	0
B: Treatment			
1	3	210.6667	21.74005
2	3	257.6667	21.74005
3	3	238	21.74005
4	3	211.3333	21.74005
AB: Block, Treatment			
1, 1	1	123	0
1, 2	1	245	0
1, 3	1	182	0
1, 4	1	203	0
2, 1	1	230	0
2, 2	1	283	0
2, 3	1	252	0
2, 4	1	204	0
3, 1	1	279	0
3, 2	1	245	0
3, 3	1	280	0
3, 4	1	227	0

Means Plots
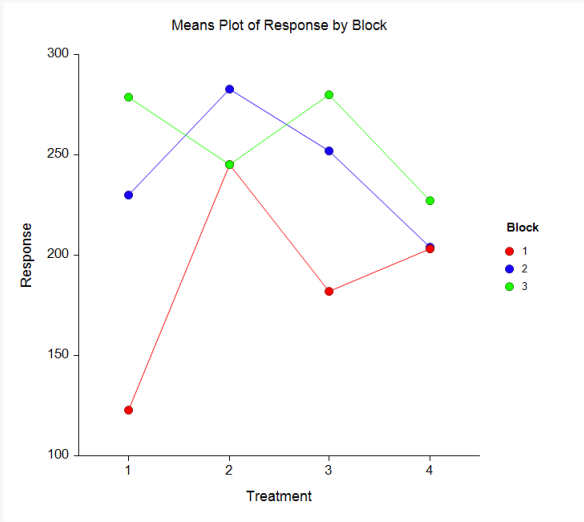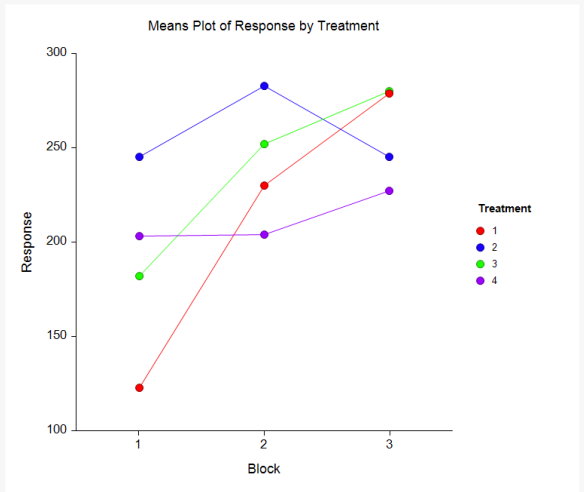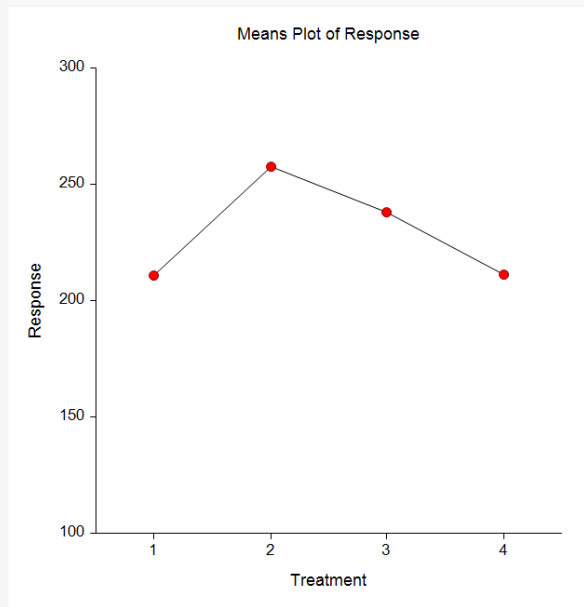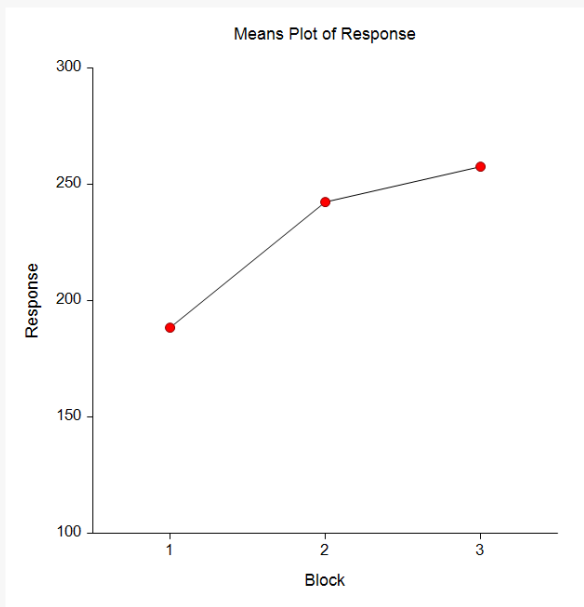
Term

The label for this line of the report.

Count

The number of observations in the mean.

Mean

The value of the sample mean.

Standard Error

The standard error of the mean. Note that these standard errors are the square root of the mean square of the error term for this term divided by the count. <u>These standard errors are not the same as the simple standard errors calculated separately for each group</u>. The standard errors reported here are those appropriate for testing multiple comparisons.

Note that the standard errors for the means of Block are zero since there is no error term for this factor. This may be seen by looking at the Expected Mean Squares Report above.

Means Plots

These plots display the means for each factor and two-way interactions. Note how easily you can see patterns in the plots.

Multiple Comparison Reports

Tukey-Kramer Multiple Comparison Test
───

Comparison Term: B: Treatment
───

Group	Count	Mean	Different From Groups
1	3	210.6667	
2	3	257.6667	
3	3	238	
4	3	211.3333	

───
α = 0.05, Error Term = AB, DF = 6, MSE = 1417.889, Critical Value = 4.8952

Notes:
This report provides multiple comparison tests for all pairwise differences between the means.

These sections present the results of the multiple-comparison procedures selected. These reports all use a uniform format that will be described by considering Tukey-Kramer Multiple-Comparison Test. The reports for the other procedures are similar. For more information on the interpretation of the various multiple-comparison procedures, turn to the section by that name in the One-Way ANOVA chapter.

Group

The label for this group.

Count

The number of observations in the mean.

Mean

The value of the sample mean.

Different From Groups

A list of those groups that are significantly different from this group according to this multiple-comparison procedure. All groups not listed are not significantly different from this group.

Alpha

The level of significance that you selected.

Error Term

The term in the ANOVA model that is used as the error term.

DF

The degrees of freedom of the error term.

MSE

The value of the mean square error.

Critical Value

The value of the test statistic that is "just significant" at the given value of alpha. This value depends on which multiple-comparison procedure you are using. It is based on the t-distribution or the studentized range distribution. It is the value of t, F, or q in the corresponding formulas.

Planned Comparison Reports

This section presents the results of any planned comparisons that were selected.

Planned Comparison: Linear Trend

Comparison Term: B: Treatment

Test Results

Comparison Value	Standard Error	T-Statistic	Two-Sided P-Value	Reject H0 at α = 0.05?	95% Confidence Interval Limits for the Comparison Value	
					Lower	Upper
-3.950387	21.74005	0.1817	0.86179	No	-57.14637	49.24559

α = 0.05, Error Term = AB, DF = 6, MSE = 1417.889

Group Details

Group	Comparison Coefficient	Count	Mean
1	-0.6708204	3	210.6667
2	-0.2236068	3	257.6667
3	0.2236068	3	238
4	0.6708204	3	211.3333

Notes:
This section presents the results of the planned comparison testing whether a first-order polynomial (a linear trend) is significant.

Comparison Value

The value of the comparison. This is formed by multiplying the comparison coefficient by the mean for each group and summing.

Standard Error

This is the standard error of the estimated comparison value. It is the denominator of the T-Value (above).

T-Statistic

The t-test used to test whether the above Comparison Value is significantly different from zero is

$$t_f = \frac{\sum_{i=1}^{k} c_i \, M_i}{\sqrt{MSE \sum_{i=1}^{k} \frac{c_i^2}{n_i}}}$$

where **MSE** is the mean square error, **f** is the degrees of freedom associated with **MSE**, **k** is the number of groups, c_i is the comparison coefficient for the i^{th} group, M_i is the mean of the i^{th} group, and n_i is the sample size of the i^{th} group.

Two-Sided P-Value

The significance level of the above T-Value. The Comparison is statistically significant if this value is less than the specified alpha.

Reject H0 at α = 0.05?

The decision based on the specified value of the multiple-comparison alpha.

Group

The label for this group.

Comparison Coefficient

The coefficient (weight) used for this group.

Count

The number of observations in the mean.

Mean

The value of the sample mean.

Alpha

The level of significance that you selected.

Error Term

The term in the ANOVA model that is used as the error term.

DF

The degrees of freedom of the error term.

MSE

The value of the mean square error.

GLM ANOVA Checklist

This checklist, prepared by a professional statistician, is a flowchart of the steps you should complete to conduct a valid analysis. Since this topic is vast, this flowchart will give only a brief summary. You should consult appropriate statistical books in your field for further details. We recommend Winer (1990) and Keppel (1991) as good books to use, but there are many others available that are equally useful.

Step 1 – Data Preparation

Introduction

This step involves scanning your data for anomalies, keypunch errors, typos, and so on. You would be surprised how often we hear of people completing an analysis, only to find that they had mistakenly selected the wrong variables.

Sample Size

The sample size (number of nonmissing rows) has a lot of ramifications. The analysis of variance was originally developed under the assumption that the sample sizes of each treatment combination are equal. In practice this seldom happens, but the closer you can get to equal sample sizes the better.

Missing Values

The number and pattern of missing values are always issues to consider. Usually, we assume that missing values occur at random throughout your data. If this is not true, your results will be biased since a particular segment of the population is underrepresented.

If you have missing values, it will be important to identify the degree of unbalance in your design. You should also check to see if there are any missing cells. If there are, you cannot run a full model. You will have to assume some interactions are zero and remove them from the ANOVA model.

Type of Data

The mathematical basis of the F-test assumes that the data are continuous. Because of the rounding that occurs when data are recorded, all data are technically discrete. The validity of assuming the continuity of the data then comes down to determining when we have too much rounding. For example, most statisticians would not worry about human-age data that was rounded to the nearest year. However, if these data were rounded to the nearest ten years or further to only three groups (young, adolescent, and adult), most statisticians question the validity of the probability statements. Some studies have shown that

the F-test is reasonably accurate when the data have only five possible values (most would call this discrete data). If your data contain less than five unique values, any probability statements made are tenuous.

Also, you should double-check to ensure that you are going to use the appropriate design. Our experience is that many researchers use a factorial design when they should be using a repeated measures design. Consider again the examples of each type of design and make sure you are using the correct one.

Outliers

Generally, outliers cause distortion in most popular statistical tests. You must scan your data for outliers (the box plot is an excellent tool for doing this). If you have outliers, you have to decide if they are one-time occurrences or if they would occur in another sample. If they are one-time occurrences, you can remove them and proceed. If you know they represent a certain segment of the population, you have to decide between biasing your results (by removing them) or leaving them in and invalidating the normality assumption.

Step 2 – Setup and Run the GLM ANOVA Panel

Introduction

Now comes the fun part: running the program. **NCSS** is designed to be simple to operate, but it can still seem complicated. When you go to run a procedure such as this for the first time, take a few minutes to read through the chapter again and familiarize yourself with the issues involved.

Enter Variables

The templates are set with ready-to-run defaults. About all you have to do is select the appropriate variables (columns of data).

Select All Plots

As a rule, you should select the means plots. They add a great deal to your ability to interpret the data.

Specify Alpha

Most beginners at statistics forget this important step and let the alpha value default to the standard 0.05. You should consciously decide what value of alpha is appropriate for your study. The 0.05 default came about when people had to rely on printed probability tables and there were only two values available: 0.05 or 0.01. Now you can set the value to whatever is appropriate.

A special note on setting the Multiple Comparison alpha. You will often want to reset this value to 0.10 so that the individual tests are made at a more reasonable significance level.

Step 3 – Check Assumptions

Introduction

Testing the assumptions of normality and equal variance is often difficult in a multi-way analysis of variance. We suggest that you make several passes through your data using our one-way ANOVA program, studying each factor separately. We suggest this because the one-way ANOVA program displays extensive diagnostic information for checking equal variance and normality. Although this method does not account for the interactions among the factors, it is often the best you can do to assess the validity of your assumptions.

Sometimes, the ANOVA model can be recoded so that you can run it through our regression program. When this is possible, you can analyze the residuals to assess normality and equal variance.

Random Sample

These statistical procedures were designed with the assumption that the sample population was selected randomly. The validity of this assumption depends on the method used to select the sample. If you have not used valid sampling techniques, the F-test will not work.

Check Descriptive Statistics

You should check the Means and Standard Errors Section first to determine if the Counts and the Means are reasonable. If you have selected the wrong variable, these values will alert you.

Step 4 – Interpret Findings

Introduction

You are now ready to conduct your tests. The basic plan of attack for analyzing your output is as follows:

1. Glance through the reports, checking the means, the F-tests, and so forth for obvious problems.

2. Look at the power of the nonsignificant tests. Could the lack of significance be the result of a small sample size?

3. Determine which main effects and interactions are significant.

4. Use care in interpreting a main effect when its interaction with another term is significant.

5. Use planned comparisons, paired comparisons, and plots of means to view the experimental results and discuss what they reveal.

Examples of Various Experimental Designs

We will now present examples of how to run various popular types of experimental designs.

Randomized-Block Design

The randomized-block design is a very popular experimental design. The focus of the analysis is on a set of two or more treatments. A blocking variable is used to account for extraneous factors. Each block receives all treatments. These treatments are randomly assigned within the block.

The data in the Randomized Block dataset show how to enter the data for this type of design. You should designate the block term as *random* and the treatment term as *fixed*. Set the Which Model Terms option to Up to 1-Way (removing the interaction term). In a typical randomized-block design, the interaction term becomes the error term, so it does not have to be fit separately. Doing this will reduce the amount of time needed to complete the calculations.

Single-Factor Repeated-Measures Design

The single-factor repeated-measures design is similar to the randomized-block design. In this design the individuals (analogous to the blocks) are measured over time. Unlike the randomized-block design, however, the treatments are not applied in random order. Instead, the treatments are always applied in the same order. For example, you might conduct a pre-test, apply some treatment to the individuals, and conduct a post-test. You cannot apply the post-test first.

The data in the Randomized Block dataset show how to enter the data for this type of design if you think of blocks as the individuals and treatments as time of measurement.

It turns out that even though the randomization method is different, the analysis of this design is identical to that described above for the randomized-block design. The individuals become the blocks. This variable is designated *random*. The repeated-measures variable (the variable representing time) becomes the treatment. This variable is designated as *fixed*. Set the Which Model Terms option (Model Tab) to Up to 1-Way to omit the interaction term.

Example 2 – Latin-Square Design

Fractional-rep designs are known for their ability to provide insight about several factors with a minimum number of observations. This efficiency comes from an experimental setup that ignores many interaction terms. The Latin square is one such design. It may be analyzed with **NCSS**. The following table shows a set of Latin-square data from page 313 of Snedecor and Cochran (1972).

Latin-Square Data from Snedecor and Cochran

	Column				
Row	**1**	**2**	**3**	**4**	**5**
1	B:257	E:230	A:279	C:287	D:202
2	D:245	A:283	E:245	B:280	C:260
3	E:182	B:252	C:280	D:246	A:250
4	A:203	C:204	D:227	E:193	B:259
5	C:231	D:271	B:266	A:334	E:338

The following table shows the data as it would be entered for analysis in **NCSS**. The *Custom Model* statement "A+B+C" would be used since many of the interactions cannot be estimated. The factors would be designated as fixed or random depending on the experimental situation.

Latin Square Dataset

Rows	Columns	Letters	Yield
1	1	B	257
1	2	E	230
1	3	A	279
1	4	C	287
1	5	D	202
2	1	D	245
2	2	A	283
2	3	E	245
2	4	B	280
2	5	C	260
3	1	E	182
3	2	B	252
3	3	C	280
3	4	D	246
3	5	A	250
4	1	A	203
4	2	C	204
4	3	D	227
4	4	E	193
4	5	B	259
5	1	C	231
5	2	D	271
5	3	B	266
5	4	A	334
5	5	E	338

Example 3 – Repeated-Measures Design

A Repeated Measures ANOVA is a particular type of three-factor design that uses two error terms. In this design, treatments are applied to experimental units of different sizes. For example, in an educational study, one treatment might be applied to whole classrooms. A second treatment might consist of the students' responses to a pre-test and a post-test. Such a design employs two error terms. One error term is the between-classes error for testing the first factor. The other error term is the within-student error for testing the second factor.

This procedure analyzes data from an experimental design represented by the following mathematical model:

$$Y_{ijkl} = \mu + A_i + S_{ij} + B_k + AB_{ik} + e_{ijkl}$$

In this model, A is the between-group treatment, S_{ij} is the between-group error, B is the within-subject treatment, and e_{ijkl} is the within-subject error.

This is a specialized technique with strict assumptions. An advanced statistical text dealing with the topic should be consulted before the technique is employed.

The data below illustrate how the data should be set up. An experiment was conducted to study the effects of exercise on heart rate. The subjects were randomly divided into three groups of six. The first group did not have a regular exercise plan. The second group exercised once a week. The third group exercised daily. Each subject's heart rate was recorded when the experiment began and again at the end of ten weeks. These data are stored in a dataset called Heart. You might want to open this dataset and run the analysis yourself.

Heart Dataset

Exercise	Subject	Time	Heart Rate
None	1	0	87
None	1	10	89
None	2	0	67
None	2	10	65
None	3	0	55
None	3	10	58
None	4	0	66
None	4	10	68
None	5	0	88
None	5	10	90
None	6	0	75
None	6	10	73
Weekly	7	0	84
Weekly	7	10	78
Weekly	8	0	78
Weekly	8	10	72
Weekly	9	0	64
Weekly	9	10	53
Weekly	10	0	73
Weekly	10	10	65
Weekly	11	0	84
Weekly	11	10	82
Weekly	12	0	55
Weekly	12	10	53
Daily	13	0	72
Daily	13	10	55
Daily	14	0	83
Daily	14	10	72
Daily	15	0	75
Daily	15	10	63
Daily	16	0	55
Daily	16	10	49
Daily	17	0	83
Daily	17	10	68
Daily	18	0	63
Daily	18	10	54

To run this analysis, you specify *Heart Rate* as the Response Variable, *Exercise* as Factor 1 (designate it as fixed), *Subject* as Factor 2 (designate it as nested), and *Time* as Factor 3 (designate it as fixed). Select the full model. When the analysis is complete, the following output is displayed.

Repeated-Measure ANOVA Report for Heart Rate Data

Analysis of Variance Table

Model Term	DF	Sum of Squares	Mean Square	F-Ratio	P-Value
A: Exercise	2	331.1667	165.5833	0.6110	0.55576
B(A): Subject	15	4064.833	270.9889		
C: Time	1	277.7778	277.7778	50.5051	0.00000*
AC	2	234.7222	117.3611	21.3384	0.00004*
BC(A)	15	82.5	5.5		
S	0	0			
Total (Adjusted)	35	4991			
Total	36				

* Term significant at α = 0.05

Means and Standard Errors

Term	Count	Mean	Standard Error
All	36	69.83334	
A: Exercise			
Daily	12	66	4.752095
None	12	73.41666	4.752095
Weekly	12	70.08334	4.752095
C: Time			
0	18	72.61111	0.5527708
10	18	67.05556	0.5527708
AC: Exercise, Time			
Daily, 0	6	71.83334	0.9574271
Daily, 10	6	60.16667	0.9574271
None, 0	6	73	0.9574271
None, 10	6	73.83334	0.9574271
Weekly, 0	6	73	0.9574271
Weekly, 10	6	67.16666	0.9574271

Example 4 – Analysis of Covariance

The analysis of covariance uses features from both analysis of variance and multiple regression. The usual one-way classification model in analysis of variance is

$$Y_{ij} = \mu_i + e_{1ij}$$

where Y_{ij} is the j^{th} observation in the i^{th} group, μ_i represents the true mean of the i^{th} group, and e_{1ij} are the residuals or errors in the above model (usually assumed to be normally distributed). Suppose you have measured a second variable with values X_{ij} that is linearly related to Y. Further suppose that the slope of the relationship between Y and X is constant from group to group. You could then write the analysis of covariance model

$$Y_{ij} = \mu_i + \beta\left(X_{ij} - X_{..}\right) + e_{2ij}$$

where $X..$ represents the overall mean of X. If X and Y are closely related, you would expect that the errors, e_{2ij}, would be much smaller than the errors, e_{1ij}, giving you more precise results.

The analysis of covariance is useful for many reasons, but it does have the (highly) restrictive assumption that the slope is constant over all the groups. This assumption is often violated, which limits the technique's usefulness. You will want to study more about this technique in statistical texts before you use it.

Running an analysis of covariance is easy in **NCSS**. You fill out the procedure template as usual for an ANOVA. To change your ANOVA into an ANCOVA, you simply specify one or more covariates. We will now take you through an extended example showing how to run an Ancova as well as how to test the assumption of equal slopes. The following data give the home state, age, and IQ of thirty teenagers. The variables X1-X4 are for use in testing the Ancova assumption of equal slopes and they will be explained later.

Suppose we wish to test for differences in IQ among the three states while controlling for age (the covariate). These data are contained in the ANCOVA database. You should open this database now if you want to follow along.

ANCOVA Dataset

State	Age	IQ	X1	X2	X3	X4
Iowa	12	100	-1	-1	-12	-12
Iowa	13	102	-1	-1	-13	-13
Iowa	12	97	-1	-1	-12	-12
Iowa	14	96	-1	-1	-14	-14
Iowa	15	105	-1	-1	-15	-15
Iowa	18	106	-1	-1	-18	-18
Iowa	12	105	-1	-1	-12	-12
Iowa	14	103	-1	-1	-14	-14
Iowa	12	99	-1	-1	-12	-12
Iowa	10	98	-1	-1	-10	-10
Utah	14	104	0	2	0	28
Utah	11	105	0	2	0	22
Utah	12	106	0	2	0	24
Utah	15	103	0	2	0	30
Utah	17	102	0	2	0	34
Utah	18	99	0	2	0	36
Utah	19	107	0	2	0	38
Utah	16	105	0	2	0	32
Utah	15	103	0	2	0	30
Utah	14	103	0	2	0	28
Texas	15	105	1	-1	15	-15
Texas	16	106	1	-1	16	-16
Texas	12	103	1	-1	12	-12
Texas	13	99	1	-1	13	-13
Texas	14	93	1	-1	14	-14
Texas	11	104	1	-1	11	-11
Texas	18	103	1	-1	18	-18
Texas	19	100	1	-1	19	-19
Texas	18	101	1	-1	18	-18
Texas	16	104	1	-1	16	-16

We begin by loading the ANCOVA dataset and the *GLM ANOVA* options panel. We specify *IQ* as the *Response Variable, Age* as the *Covariate,* and *State* as *Factor 1*. We run the procedure and the analysis of covariance table is displayed.

Analysis of Covariance Report

Analysis of Variance Table

Model Term	DF	Sum of Squares	Mean Square	F-Ratio	P-Value
X(Age)	1	5.239314	5.239314	0.4740	0.49723
A: State	2	28.38448	14.19224	1.2841	0.29389
S	26	287.3607	11.05233		
Total (Adjusted)	29	328.8			
Total	30				

Notice that now, in addition to the test for factor A, we also have a test for the covariate. This test, the one along the line labeled "X(Age)," tests the significance of the covariate. If it is not significant (as is the case in this example), analysis of covariance should not be used. However, if it is significant, you may proceed to the next F-test, the one dealing with factor A (State). This is the test that is usually desired in the analysis of covariance. It tests whether the adjusted means of the three states are different. The means are adjusted as if all three states had the same age. That is, the means for each state are adjusted to the average value of age. These adjusted

means are shown in the *Means and Effects* report. If you run the analysis without the covariate, you'll notice that these means are different.

Since the covariate (Age) is not significant, you should stop here. However, for the sake of instruction, we will assume that the covariate is significant and proceed to test whether the slopes between IQ and Age are the same in the three states. The following steps will lead you through this test:

1. Construct a new contrast variable for each degree of freedom of the factor. In our current example, the three levels (states) of factor A yield two degrees of freedom, so we must create two contrast variables. These are shown as X1 and X2.

2. Multiply each of these new variables by the covariate variable. In our example, X3 = (X1)(Age) and X4 = (X2)(Age).

3. Run another ANCOVA, using the same setup as before except now you fit the three covariates Age, X3, and X4. Call these the Model 2 results and call the previous results with just the single covariate the Model 1 results.

Second Analysis of Covariance Report

Analysis of Variance Table

Model Term	DF	Sum of Squares	Mean Square	F-Ratio	P-Value
X(Age)	1	9.740934	9.740934	0.9402	0.34189
X(X3)	1	22.27164	22.27164	2.1498	0.15557
X(X4)	1	21.07455	21.07455	2.0342	0.16667
A: State	2	46.57466	23.28733	2.2478	0.12741
S	24	248.6402	10.36001		
Total (Adjusted)	29	328.8			
Total	30				

4. Finally, create the F-test for equality of slopes as follows. The formula is

$$F_{k,m} = \frac{(SSE_1 - SSE_2) \, / \, k}{MSE_2}$$

where k is the degrees of freedom of the factor (in our example, this is 2), m is the degrees of freedom of the mean square for error in model2, SSE_1 and SSE_2 are the sums of squares error for model1 and model2, and MSE_2 is the mean square for error in model2.

The calculations for this example proceed as follows:

$F_{2,24}$ = [(287.3607-248.6402)/2]/10.36001 = 1.86875.

This F-ratio would then be compared against a tabulated 0.05 F-value, 3.403, which you could find in the probability calculator or in a statistics book. Since 1.86875 < 3.403, we would not reject the equality of slopes assumption in this case.

One final note, you should generate a scatter plot, which shows the response variable on the vertical axis, the covariate on the horizontal axis, and uses different symbols for each group. The least squares trend line can also be displayed. This plot will let you visually assess the validity of the assumption of equal slope.

Example 5 – Hierarchical-Classification Design

Snedecor and Cochran (1967), page 286, present an example of a hierarchical-classification design. In this example, four plants were selected at random, and three leaves were randomly selected from each plant. Two samples were taken from each leaf, and the amount of calcium in the sample was recorded. The data are displayed below. The data are stored in a database called Plant.

Plant Dataset

Row	Plant	Leaf	Calcium
1	1	1	3.28
2	1	1	3.09
3	1	2	3.52
4	1	2	3.48
5	1	3	2.88
6	1	3	2.80
7	2	4	2.46
8	2	4	2.44
9	2	5	1.87
10	2	5	1.92
11	2	6	2.19
12	2	6	2.19
13	3	7	2.77
14	3	7	2.66
15	3	8	3.74
16	3	8	3.44
17	3	9	2.55
18	3	9	2.55
19	4	10	3.78
20	4	10	3.87
21	4	11	4.07
22	4	11	4.12
23	4	12	3.31
24	4	12	3.31

To run this analysis, you specify *Calcium* as the Response Variable, *Plant* as Factor 1 (designate it as random), and *Leaf* as Factor 2 (designate it as nested). Select the *Full Model*. When the analysis is complete, the following output is displayed.

Plant Data Example Report

Expected Mean Squares

Model Term	DF	Term Fixed?	Denominator Term	Expected Mean Square
A: Plant	3	No	B(A)	S+sB+bsA
B(A): Leaf	8	No	S(AB)	S+sB
S(AB)	12	No		S

Note: Expected Mean Squares are for the balanced cell-frequency case.

Analysis of Variance Table

Model Term	DF	Sum of Squares	Mean Square	F-Ratio	P-Value
A: Plant	3	7.560346	2.520115	7.6652	0.00973*
B(A): Leaf	8	2.6302	0.328775	49.4089	0.00000*
S	12	0.07985	0.006654167		
Total (Adjusted)	23	10.2704			
Total	24				

* Term significant at α = 0.05

Example 6 – Split-Split-Plot Design

Gomez and Gomez (1984), page 142-153, present an example of a split-split-plot design. This design is an extension of the split-plot design in which the plots are divided into subplots. Thus, there are two nesting factors and three levels of precision. There are three plot sizes: the main plot, the subplot, and the sub-subplot.

Their example is of a 5 x 3 x 3 factorial experiment with three replications (blocks). The treatments are five levels of nitrogen on a main plot, three types of management applied to the subplot, and three varieties of rice applied to sub-subplots.

The data are too lengthy to be presented here, but they are contained in the Grain Yield dataset.

Setup

To run this example, complete the following steps:

1   Open the Grain Yield example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **Grain Yield** and click **OK**.

2   Specify the General Linear Models (GLM) procedure options

- Find and open the **General Linear Models (GLM)** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 6** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

Response Variable(s) ....**Yield**
Factor Variable 1 ....**Rep**
  Type...**Random**
Factor Variable 2 ....**Nitrogen**
  Type...**Fixed**
Factor Variable 3 ....**Plot**
  Type...**Nested**
Factor Variable 4 ....**Management**
  Type...**Fixed**
Factor Variable 5 ....**Subplot**
  Type...**Nested**
Factor Variable 6 ....**Variety**
  Type...**Fixed**
Which Model Terms.......................................**Custom Model**
Custom Model...**A+B+C(AB)+D+BD+CD(AB)+F+BF+DF+BDF**

3   Run the procedure

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

Expected Mean Squares

Expected Mean Squares

Model Term	DF	Term Fixed?	Denominator Term	Expected Mean Square
A: **Rep**	2	No	C(AB)	S+defsC+bcdefsA
B: **Nitrogen**	4	Yes	C(AB)	S+defsC+acdefsB
C(AB): **Plot**	8	No	S(ABCDEF)	S+defsC
D: **Management**	2	Yes	CD(AB)	S+efsCD+abcefsD
BD	8	Yes	CD(AB)	S+efsCD+acefsBD
CD(AB)	20	No	S(ABCDEF)	S+efsCD
F: **Variety**	2	Yes	S(ABCDEF)	S+abcdesF
BF	8	Yes	S(ABCDEF)	S+acdesBF
DF	4	Yes	S(ABCDEF)	S+abcesDF
BDF	16	Yes	S(ABCDEF)	S+acesBDF
S(ABCDEF)	60	No		S

Note: Expected Mean Squares are for the balanced cell-frequency case.

Analysis of Variance Table

Analysis of Variance Table

Model Term	DF	Sum of Squares	Mean Square	F-Ratio	P-Value
A: Rep	2	0.7319945	0.3659973	0.6578	0.54391
B: Nitrogen	4	61.64082	15.41021	**27.6953**	**0.00010***
C(AB): Plot	8	4.451351	0.5564188	1.1229	0.36129
D: Management	2	42.93611	21.46805	**81.9965**	**0.00000***
BD	8	1.102973	0.1378717	0.5266	0.82265
CD(AB)	20	5.236335	0.2618167	0.5283	0.94267
F: Variety	2	206.0132	103.0066	**207.8667**	**0.00000***
BF	8	14.14451	1.768063	**3.5679**	**0.00192***
DF	4	3.851769	0.9629423	**1.9432**	**0.11490**
BDF	16	3.699232	0.231202	**0.4666**	**0.95376**
S	60	29.73249	0.4955415		
Total (Adjusted)	134	373.5407			
Total	135				

* Term significant at α = 0.05

A look at Gomez and Gomez (1984) page 153 will show that these are the exact same answers that were given there.