

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 so obtained are analyzed using an analysis of variance table that produces an F-test.

General Linear Models (GLM)

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 $i=1,2,\dots,I$; $j=1,2,\dots,J$; and $k=1,2,\dots,K$. 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 i^{th} level of a factor A.

b_j the contribution of the j^{th} level of a factor B.

$(ab)_{ij}$ the combined contribution of the i^{th} level of a factor A and the j^{th} level of a factor B.

e_{ijk} the contribution of the k^{th} 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.

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

Procedure Options

This section describes the options available in this procedure.

Variables Tab

These panels specify the variables used in the analysis and the model.

Response Variables

Response Variable(s)

Specifies the response (dependent) variable to be analyzed. If you specify more than one variable here, a separate analysis is run for each variable.

Covariate Specification

Covariate(s)

One or more covariates may be specified, causing an analysis of covariance (ANCOVA) to be run. Note that you cannot specify covariates if any of your factors are of the nested type.

Factor Specification

Factor Variable

At least one factor variable must be specified. This variable's values indicate how the values of the response variable should be categorized. Examples of factor variables are gender, age groups, "yes" or "no" responses, etc. Note that the values in the variable may be either numeric or text. The treatment of text variables is specified for each variable by the *Data Type* option on the database.

General Linear Models (GLM)

Type

This option specifies whether the factor is fixed, random, or nested.

A **fixed** factor includes all possible levels, like male and female for gender, includes representative values across the possible range of values, like low, medium, and high temperatures, or includes a set of values to which inferences will be limited, like New York, California, and Maryland.

A **random** factor is one in which the chosen levels represent a random sample from the population of values. For example, you might select four classes from the hundreds in your state or you might select ten batches from an industrial process. The key is that a random sample is chosen. In *NCSS*, a random factor is “crossed” with other random and fixed factors. Two factors are crossed when each level of one includes all levels of the other.

A **nested** factor is a special type of random factor whose levels (values) are not repeated for all combinations of the factors before it. That is, if factor B is nested in factor A, each level of factor A has its own set of values for factor B.

For example, suppose that factor A represents three fourth-grade classrooms of twenty students in a particular state. Further suppose that factor B represents the sixty children in these classrooms. If factors A and B were crossed, then all sixty children would somehow simultaneously be attending all three classrooms. However, if each classroom has a mutually exclusive set of twenty children, we say that children are nested within classrooms or B is nested within A. Notice that nesting occurs when each level of the first factor (the classrooms) contains separate levels of the second factor (the children).

Note that nested factors should be numbered consecutively, just like random and fixed factors. In the preceding example, you would number the children from one to sixty. You cannot have two individuals with the same identification number.

Comparisons

Comparisons are only valid for fixed factors. This option lets you specify comparisons that you want to run on this factor. 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.

Comparisons may be specified by simply listing the weights. For example, suppose a factor has three levels (unique values). Further suppose that the first level represents a control group, the second a treatment at one dose, and the third a treatment at a higher dose. Three comparisons come to mind: compare each of the treatment groups to the control group and compare the two treatment groups to each other. These three comparisons would be

Control vs. Treatment 1	-1,1,0
Control vs. Treatment 2	-1,0,1
Treatment 1 vs. Treatment 2	0,-1,1

You might also be interested in comparing the control group with the average of both treatment groups. The weights for this comparison would be -2,1,1.

When a factor is quantitative, it might be of interest to divide the response pattern into linear, quadratic, cubic, or other components. If the sample sizes are equal and the factor levels are equally spaced, these so-called components of trend may be studied by the use of simple contrasts. For example, suppose a quantitative factor has three levels: 5, 10, and 15. Contrasts to test the linear and quadratic trend components would be

Linear trend	-1,0,1
Quadratic trend	1,-2,1

General Linear Models (GLM)

If the sample sizes for the groups are unequal (the design is unbalanced), adjustments must be made for the differing sample sizes.

NCSS will automatically generate some of the more common sets of contrasts, or it will let you specify up to three custom contrasts yourself. The following common sets are designated by this option.

- **None**

No comparisons are generated.

- **Standard Set**

This option generates a standard set of contrasts in which the mean of the first level is compared to the average of the rest, the mean of the second group is compared to the average of those remaining, and so on.

The following example displays the type of contrast generated by this option. Suppose there are four levels (groups) in the factor. The contrasts generated by this option are:

-3,1,1,1	Compare the first-level mean with the average of the rest.
0,-2,1,1	Compare the second-level mean with the average of the rest.
0,0,-1,1	Compare the third-level mean with the fourth-level mean.

- **Polynomial**

This option generates a set of orthogonal contrasts that allow you to test various trend components from linear up to sixth order. These contrasts are appropriate even if the levels are unequally spaced or the group sample sizes are unequal. Of course, these contrasts are only appropriate for data that are at least ordinal. Usually, you would augment the analysis of this type of data with a multiple regression analysis.

The following example displays the type of contrasts generated by this option. Suppose there are four equally spaced levels in the factor and each group has two observations. The contrasts generated by this option are (scaled to whole numbers):

-3,-1,1,3	Linear component.
1,-1,-1,1	Quadratic component.
-1,3,-3,1	Cubic component.

- **Linear Trend**

This option generates a set of orthogonal contrasts and retains only the linear component. This contrast is appropriate even if the levels are unequally spaced and the group sample sizes are unequal. See Orthogonal Polynomials above for more detail.

- **Linear-Quadratic Trend**

This option generates the complete set of orthogonal polynomials, but only the results for the first two (the linear and quadratic) are reported.

- **Linear-Cubic Trend**

This option generates the complete set of orthogonal polynomials, but only the results for the first three are reported.

- **Linear-Quartic Trend**

This option generates the complete set of orthogonal polynomials, but only the results for the first four are reported.

General Linear Models (GLM)

- **Each with First**

This option generates a set of nonorthogonal contrasts appropriate for comparing each of the remaining levels with the first level. The following example displays the type of contrast generated by this option. Suppose there are four levels (groups) in the factor. The contrasts generated by this option are:

-1,1,0,0	Compare the first- and second-level means.
-1,0,1,0	Compare the first- and third-level means.
-1,0,0,1	Compare the first- and fourth-level means.

- **Each with Last**

This option generates a set of nonorthogonal contrasts appropriate for comparing each of the remaining levels with the last level. The following example displays the type of contrast generated by this option. Suppose there are four levels (groups) in the factor. The contrasts generated by this option are:

-1,0,0,1	Compare the first- and fourth-level means.
0,-1,0,1	Compare the second- and fourth-level means.
0,0,-1,1	Compare the third- and fourth-level means.

- **Custom**

This option indicates that the contrasts listed in the corresponding three boxes of the Comparison panel should be used.

Model Specification

This section specifies the experimental design model.

Which Model Terms

A design in which main effect and interaction terms are included is called a saturated model. Often, it is useful to omit various interaction terms from the model. This option lets you specify which interactions to keep very easily. If the selection provided here is not flexible enough for your needs, you can specify *custom* here and enter the model directly.

The options included here are as follows.

- **Full Model**

The complete, saturated model is analyzed. This option requires that you have no missing cells, although you can have an unbalanced design. Hence, you cannot use this option with Latin square or fractional factorial designs.

- **Up to 1-Way**

A main-effects only model is run. All interactions are omitted.

- **Up to 2-Way**

All main-effects and two-way interactions are included in the model.

- **Up to 3-Way**

All main-effects, two-way, and three-way interactions are included in the model.

- **Up to 4-Way**

All main-effects, two-way, three-way, and four-way interactions are included in the model.

General Linear Models (GLM)

- **Custom**

This option indicates that you want the Custom Model (given in the next box) to be used.

- **Write Model in ‘Custom Model’ Field**

When this option is checked, no analysis is performed when the procedure is run. Instead, a copy of the full model is stored in the Custom Model box. You can then delete selected terms from the model without having to enter all the terms you want to keep.

Custom Model

When Custom Model (see Which Model Terms above) is selected, the model itself is entered here. If all main effects and interactions are desired, you can enter the word “ALL” here. For complicated designs, it is usually easier to check the next option, Write Model in ‘Custom Model’ Field, and run the procedure. The appropriate model will be generated and placed in this box. You can then edit it as you desire.

The model is entered using letters separated by the plus sign. For example, a three-factor factorial in which only two-way interactions are needed would be entered as follows:

$A+B+AB+C+AC+BC$.

A simple repeated-measures design would look like this:

$A+B(A)+C+AC+BC(A)$.

Custom Comparisons Tab

This panel is used when the Comparison option of one or more factors is set to Custom. The Custom option means that the contrast coefficients are to be entered by the user. The boxes on this panel contain the user-supplied contrast coefficients. The first row is for factor one, the second row for factor two, and so on.

Custom Comparisons

The following options are only used if Comparisons is set to 'Custom' on the Variables tab.

Custom (1-3)

This option lets you write a user-specified comparison by specifying the weights of that comparison. Note that there are no numerical restrictions on these coefficients. They do not even have to sum to zero. However, this is recommended. If the coefficients do sum to zero, the comparison is called a contrast. The significance tests anticipate that only one or two of these comparisons are to be run. If you run several, you should make some type of Bonferroni adjustment to your alpha value.

When you put in your own contrasts, you must be careful that you specify the appropriate number of weights. For example, if the factor has four levels, four weights must be specified, separated by commas. Extra weights are ignored. If too few weights are specified, the missing weights are set to zero.

These comparison coefficients designate weighted averages of the level-means that are to be statistically tested. The null hypothesis is that the weighted average is zero. The alternative hypothesis is that the weighted average is nonzero. The weights (comparison coefficients) are specified here.

As an example, suppose you want to compare the average of the first two levels with the average of the last two levels in a six-level factor. You would enter “-1,-1,0,0,1,1.”

As a second example, suppose you want to compare the average of the first two levels with the average of the last three levels in a six-level factor. The contrast would be

$-3,-3,0,2,2,2$.

Note that in each case, we have used weights that sum to zero. This is why we could not use ones in the second example.

Reports Tab

The following options control which reports are displayed.

Select Reports

EMS Report ... Means Report

Specify whether to display the indicated reports.

Test Alpha

The value of alpha for the statistical tests and power analysis. Usually, this number will range from 0.10 to 0.001. A common choice for alpha is 0.05, but this value is a legacy from the age before computers when only printed tables were available. You should determine a value appropriate for your particular study.

Multiple Comparison Tests

Note that if covariates are used, no multiple comparison tests are displayed.

Bonferroni Test (All-Pairs) ... Tukey-Kramer Confidence Intervals

These options specify which MC tests and confidence intervals to display.

Tests for Two-Factor Interactions

This option specifies whether multiple comparison tests are generated for two-factor interaction terms. When checked, the means of two-factor interactions will be tested by each active multiple comparison test. The multiple comparison test will treat the means as if they came from a single

factor. For example, suppose factor A has two levels and factor B has three levels. The AB interaction would then have six levels. The active multiple comparison tests would be run on these six means.

Care must be used when interpreting multiple comparison tests on interaction means. Remember that these means contain not only the effects of the interaction, but also the main effects of the two factors. Hence these means contain the combined effects of factor A, factor B, and the AB interaction. You cannot interpret the results as representing only the AB interaction.

Multiple Comparison Tests – Options

MC Alpha

Specifies the alpha value used by the multiple-comparison tests.

MC Decimals

Specify how many decimals to display in the multiple comparison sections.

Report Options

Precision

Specify the precision of numbers in the report. Single precision will display seven-place accuracy, while the double precision will display thirteen-place accuracy.

Variable Names

Indicate whether to display the variable names or the variable labels.

Value Labels

Indicate whether to display the data values or their labels.

Plots Tab

These options specify the plots of group means.

Select Plots

Means Plot(s)

Specify whether to display the indicated plots. Click the plot format button to change the plot settings.

Y-Axis Scaling

Specify the method for calculating the minimum and maximum along the vertical axis. *Separate* means that each plot is scaled independently. *Uniform* means that all plots use the overall minimum and maximum of the data. This option is ignored if a minimum or maximum is specified.

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.

You may follow along here by making the appropriate entries or load the completed template **Example 1** by clicking on Open Example Template from the File menu of the General Linear Models (GLM) window.

1 Open the Randomized Block dataset.

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Click on the file **Randomized Block.NCSS**.
- Click **Open**.

2 Open the General Linear Models (GLM) window.

- Using the Analysis menu or the Procedure Navigator, find and select the **General Linear Models (GLM)** procedure.
- On the menus, select **File**, then **New Template**. This will fill the procedure with the default template.

3 Specify the variables.

- Select the **Variables tab**.
- Double-click in the **Response Variables** box. This will bring up the variable selection window.
- Select **Response** from the list of variables and then click **Ok**.
- Double-click in the **Factor 1 Variable** box. This will bring up the variable selection window.
- Select **Block** from the list of variables and then click **Ok**.
- Select **Random** in the **Type** box for Factor 1.
- Double-click in the **Factor 2 Variable** box. This will bring up the variable selection window.
- Select **Treatment** from the list of variables and then click **Ok**.
- Select **Fixed** in the **Type** box for Factor 2.
- Select **Linear** in the **Comparisons** box for Factor 2.

4 Specify the reports.

- Select the **Reports tab**.
- Check the **Tukey-Kramer Test** option of the Multiple Comparison Tests.

5 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the green Run button.

General Linear Models (GLM)

We will now document this output, one section at a time.

Expected Mean Squares Section

Expected Mean Squares Section

Source 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.

Source 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 Section

Analysis of Variance Table

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
A (Blocks)	2	10648.67	5324.333			
B (Treatment)	3	4650.917	1550.306	1.09	0.421359	0.177941
AB	6	8507.333	1417.889			
S	0	0				
Total (Adjusted)	11	23806.92				
Total	12					

* Term significant at alpha = 0.05

General Linear Models (GLM)

Source 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.

Prob Level

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 (Alpha=0.05)

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.

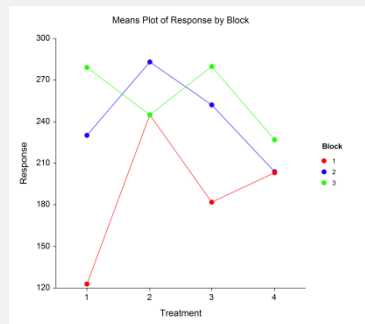


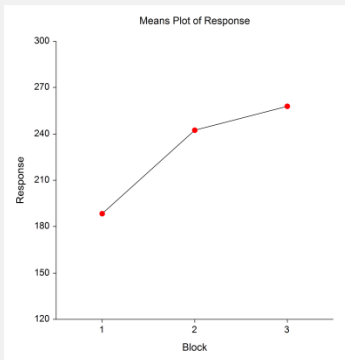
Generally, you would consider the power of the test when you accept the null hypothesis. The power will give you some idea of what actions you might take to make your results significant. If you accept the null hypothesis with high power, there is not much left to do. At least you know that the means 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 Section, Plots Sections

Means and Standard Errors , and Plots Sections

Term	Count	Mean	Standard Error
All	12	229.4167	
A: Blocks			
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
.	.	.	.
.	.	.	.



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. These standard errors are not the same as the simple standard errors calculated separately for each group. 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.

General Linear Models (GLM)

Plot of Means

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

Multiple-Comparison Sections**Tukey-Kramer Multiple-Comparison Test**

Response: Response
Term B: Treatment

Alpha=0.050 Error Term=AB DF=6 MSE=1417.889 Critical Value=4.8956

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

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.

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 for 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.

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.

Planned-Comparison Section

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

Planned Comparison: B Linear Trend

Response: Response
Term B: Treatment

Alpha=0.050 Error Term=AB DF=6 MSE=1417.889

Comparison Value=-3.950387 T-Value=0.1817 Prob>|T|=0.861794 Decision(0.05)=Do Not Reject
Comparison Standard Error=21.74005 Comparison Confidence Interval = -57.14637 to 49.24559

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

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.

Comparison Value

The value of the comparison. This is formed by multiplying the Comparison Coefficient times the Mean for each group and summing.

T-Value

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

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

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.

Prob>|T|

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

Decision(0.05)

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

General Linear Models (GLM)

Comparison Standard Error

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

Group

The label for this group.

Comparison Coefficient

The coefficient (weight) used for this group. Note that for our example, the weights are appropriate for the linear-trend component of a set of orthogonal polynomials.

Count

The number of observations in the mean.

Mean

The value of the sample mean.

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

General Linear Models (GLM)

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.

General Linear Models (GLM)

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

General Linear Models (GLM)

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

Row	Column				
	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

General Linear Models (GLM)

Latin Square dataset (continued)

Rows	Columns	Letters	Yield
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

General Linear Models (GLM)

Heart dataset (continued)

Exercise	Subject	Time	Heart Rate
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

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
A: Exercise	2	331.1667	165.5833	0.61	0.555761	0.133134
B(A): Subject	15	4064.833	270.9889			
C: Time	1	277.7778	277.7778	50.51	0.000004*	0.999998
AC	2	234.7222	117.3611	21.34	0.000041*	0.999793
BC(A)	15	82.5	5.5			
S	0	0				
Total (Adjusted)	35	4991				
Total	36					

* Term significant at alpha = 0.05

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	.5527708
10	18	67.05556	.5527708
AC: Exercise,Time			
Daily,0	6	71.83334	.9574271
Daily,10	6	60.16667	.9574271
None,0	6	73	.9574271
None,10	6	73.83334	.9574271
Weekly,0	6	73	.9574271
Weekly,10	6	67.16666	.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_{ij} 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

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

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

General Linear Models (GLM)

ANCOVA dataset (continued)

State	Age	IQ	X1	X2	X3	X4
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

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
X(Age)	1	5.239314	5.239314	0.47	0.497230	0.101741
A: State	2	28.38448	14.19224	1.28	0.293886	0.253298
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.

General Linear Models (GLM)

Second Analysis of Covariance Report

Analysis of Variance Table

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
X(Age)	1	9.740934	9.740934	0.94	0.341886	0.153727
X(X3)	1	22.27164	22.27164	2.15	0.155572	0.290793
X(X4)	1	21.07455	21.07455	2.03	0.166672	0.277862
A: State	2	46.57466	23.28733	2.25	0.127408	0.412350
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

General Linear Models (GLM)

Plant dataset (continued)

Row	Plant	Leaf	Calcium
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 Section

Source 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

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
A: Plant	3	7.560346	2.520115	7.67	0.009725*	
B(A): Leaf	8	2.6302	0.328775	49.41	0.000000*	
S	12	0.07985	0.006654167			
Total (Adjusted)	23	10.2704				
Total	24					

* Term significant at alpha = 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.

You may follow along here by making the appropriate entries or load the completed template **Example 6** by clicking on Open Example Template from the File menu of the General Linear Models (GLM) window.

1 Open the Grain Yield dataset.

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Click on the file **Grain Yield.NCSS**.
- Click **Open**.

2 Open the General Linear Models (GLM) window.

- Using the Analysis menu or the Procedure Navigator, find and select the **General Linear Models (GLM)** procedure.
- On the menus, select **File**, then **New Template**. This will fill the procedure with the default template.

3 Specify the variables and model.

- Select the **Variables** tab.
- Set the **Response Variable** to **Yield**.

- Set **Factor Variable 1** to **Rep**.
- Set its **Type** to **Random**.

- Set **Factor Variable 2** to **Nitrogen**.
- Set its **Type** to **Fixed**.

- Set **Factor Variable 3** to **Plot**.
- Set its **Type** to **Nested**.

- Set **Factor Variable 4** to **Management**.
- Set its **Type** to **Fixed**.

- Set **Factor Variable 5** to **Subplot**.
- Set its **Type** to **Nested**.

- Set **Factor Variable 6** to **Variety**.
- Set its **Type** to **Fixed**.

- Set **Which Model Terms** to **Custom Model**.
- Set the **Custom Model** to **A+B+C(AB)+D+BD+CD(AB)+F+BF+DF+BDF**.

4 Specify the reports.

- Select the **Reports** tab.
- Check the **EMS Report** and the **ANOVA Report**, and the **Means Report**.

General Linear Models (GLM)

5 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the green Run button.

The following reports will appear.

Expected Mean Squares Section

Expected Mean Squares Section				
Source 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	S

Analysis of Variance Table Section

Analysis of Variance Table						
Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
A: Rep	2	0.7319945	0.3659973	0.66	0.543910	
B: Nitrogen	4	61.64082	15.41021	27.70	0.000097*	0.999999
C(AB): Plot	8	4.451351	0.5564188	1.12	0.361290	
D: Management	2	42.93611	21.46805	82.00	0.000000*	1.000000
BD	8	1.102973	0.1378717	0.53	0.822648	0.182126
CD(AB)	20	5.236335	0.2618167	0.53	0.942667	
F: Variety	2	206.0132	103.0066	207.87	0.000000*	1.000000
BF	8	14.14451	1.768063	3.57	0.001916*	0.969208
DF	4	3.851769	0.9629423	1.94	0.114899	0.552586
BDF	16	3.699232	0.231202	0.47	0.953759	0.264199
S	60	29.73249	0.4955415			
Total (Adjusted)	134	373.5407				
Total	135					

* Term significant at alpha = 0.05

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