

Chapter 225

One-Way Analysis of Covariance (ANCOVA)

Introduction

This procedure performs analysis of covariance (ANCOVA) with one group variable and one covariate. This procedure uses multiple regression techniques to estimate model parameters and compute least squares means. This procedure also provides standard error estimates for least squares means and their differences and calculates many multiple comparison tests (Tukey-Kramer, Dunnett, Bonferroni, Scheffe, Sidak, and Fisher's LSD) with corresponding adjusted P-values and simultaneous confidence intervals. The procedure also provides response vs covariate by group scatter plots and residuals for checking model assumptions.

This procedure will analyze One-Way ANOVA models if no covariate is entered and simple linear regression models if no group variable is entered. This allows you to complete the ANCOVA analysis if either the group variable or covariate is determined to be non-significant. For additional options related to one-way ANOVA and simple linear regression analyses, we suggest you use the corresponding procedures in **NCSS**.

This procedure cannot be used to analyze models that include more than one covariate variable or more than one group variable. If the model you want to analyze includes more than one covariate variable and/or more than one group variable, use the General Linear Models (GLM) for Fixed Factors procedure instead.

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 or group) while holding other factors constant. These experiments are called single-factor or single-group experiments and are analyzed with the one-way analysis of variance (ANOVA). Analysis of covariance (ANCOVA) is useful when you want to improve precision by removing extraneous sources of variation from your study by including a covariate.

The ANCOVA Model

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 assuming equal slopes as

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

One-Way Analysis of Covariance (ANCOVA)

where $\bar{X}_{..}$ 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 classical 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.

If it is not reasonable to conclude that the slopes are equal, then a covariate-by-group interaction term should be included in the model.

Assumptions

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

1. The response variable is continuous.
2. The treatments do not affect the value of the covariate, X_{ij} .
3. The e_{2ij} follow the normal probability distribution with mean equal to zero.
4. The variances of the e_{2ij} are equal for all values of i and j .
5. The individuals are independent.

Normality of Residuals

The residuals are assumed to follow the normal probability distribution with zero mean and constant variance. This can be evaluated using a normal probability plot of the residuals. Also, normality tests are used to evaluate this assumption. The most popular of the five normality tests provided is the Shapiro-Wilk test.

Unfortunately, a breakdown in any of the other assumptions results in a departure from this assumption as well. Hence, you should investigate the other assumptions first, leaving this assumption until last.

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.

Representing Group Variables

Categorical group variables take on only a few unique values. For example, suppose a therapy variable has three possible values: A, B, and C. One question is how to include this variable in the regression model. At first glance, we can convert the letters to numbers by recoding A to 1, B to 2, and C to 3. Now we have numbers. Unfortunately, we will obtain completely different results if we recode A to 2, B to 3, and C to 1. Thus, a direct recode of letters to numbers will not work.

To convert a categorical variable to a form usable in regression analysis, we must create a new set of numeric variables. If a categorical variable has k values, $k - 1$ new binary variables must be generated.

One-Way Analysis of Covariance (ANCOVA)

Indicator (Binary) Variables

Indicator (dummy or binary) variables are created as follows. A *reference value* is selected. Usually, the most common value or the control is selected as the reference value. Next, a variable is generated for each of the values other than the reference value. For example, suppose that C is selected as the reference value. An indicator variable is generated for each of the remaining values: A and B. The value of the indicator variable is one if the value of the original variable is equal to the value of interest, or zero otherwise. Here is how the original variable T and the two new indicator variables TA and TB look in a short example.

<u>T</u>	<u>TA</u>	<u>TB</u>
A	1	0
A	1	0
B	0	1
B	0	1
C	0	0
C	0	0

The generated variables, TA and TB, would be used as columns in the design matrix, X , in the model.

Representing Interactions of Numeric and Categorical Variables

When the interaction between a group variable and a covariate is to be included in the model, all proceeds as above, except that an interaction variable must be generated for each categorical variable. This can be accomplished automatically in **NCSS** based on the slopes assumption. When assuming that the slopes are unequal all applicable covariate-by-group interaction variables are automatically created.

In the following example, the interaction between the group variable T and the covariate variable X is created.

<u>T</u>	<u>TA</u>	<u>TB</u>	<u>X</u>	<u>XTA</u>	<u>XTB</u>
A	1	0	1.2	1.2	0
A	1	0	1.4	1.4	0
B	0	1	2.3	0	2.3
B	0	1	4.7	0	4.7
C	0	0	3.5	0	0
C	0	0	1.8	0	0

When the variables XTA and XTB are added to the model, they will account for the interaction between T and X .

Technical Details

This section presents the technical details of the analysis method (multiple regression) using a mixture of summation and matrix notation.

The Linear Model

The linear model can be written as

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$$

where \mathbf{Y} is a vector of N responses, \mathbf{X} is an $N \times p$ design matrix, $\boldsymbol{\beta}$ is a vector of p fixed and unknown parameters, and \mathbf{e} is a vector of N unknown, random error values. Define the following vectors and matrices:

$$\mathbf{Y} = \begin{bmatrix} y_1 \\ \vdots \\ y_j \\ \vdots \\ y_N \end{bmatrix}, \mathbf{X} = \begin{bmatrix} 1 & x_{11} & \cdots & x_{1p} \\ \vdots & \vdots & \cdots & \vdots \\ 1 & x_{1j} & \cdots & x_{pj} \\ \vdots & \vdots & \cdots & \vdots \\ 1 & x_{1N} & \cdots & x_{pN} \end{bmatrix}, \mathbf{e} = \begin{bmatrix} e_1 \\ \vdots \\ e_j \\ \vdots \\ e_N \end{bmatrix}, \mathbf{1} = \begin{bmatrix} 1 \\ \vdots \\ 1 \\ \vdots \\ 1 \end{bmatrix}, \boldsymbol{\beta} = \begin{bmatrix} \beta_0 \\ \beta_1 \\ \vdots \\ \beta_p \end{bmatrix}$$

\mathbf{X} is the design matrix that includes the covariate, binary variables formed from the group variable, and variables resulting from the covariate-by-group interaction (if included).

Least Squares

Using this notation, the least squares estimates of the model coefficients, \mathbf{b} , are found using the equation.

$$\mathbf{b} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$$

The vector of predicted values of the response variable is given by

$$\hat{\mathbf{Y}} = \mathbf{X}\mathbf{b}$$

The residuals are calculated using

$$\mathbf{e} = \mathbf{Y} - \hat{\mathbf{Y}}$$

One-Way Analysis of Covariance (ANCOVA)

Estimated Variances

An estimate of the variance of the residuals is computed using

$$s^2 = \frac{\mathbf{e}'\mathbf{e}}{N - p - 1}$$

An estimate of the variance of the model coefficients is calculated using

$$\mathbf{v} \begin{pmatrix} b_0 \\ b_1 \\ \vdots \\ b_p \end{pmatrix} = s^2 (\mathbf{X}'\mathbf{X})^{-1}$$

An estimate of the variance of the predicted mean of Y at a specific value of X , say X_0 , is given by

$$s_{Y_m|X_0}^2 = s^2 (1, X_0) (\mathbf{X}'\mathbf{X})^{-1} \begin{pmatrix} 1 \\ X_0 \end{pmatrix}$$

An estimate of the variance of the predicted value of Y for an individual for a specific value of X , say X_0 , is given by

$$s_{Y_I|X_0}^2 = s^2 + s_{Y_m|X_0}^2$$

Hypothesis Tests of the Intercept and Coefficients

Using these variance estimates and assuming the residuals are normally distributed, hypothesis tests may be constructed using the Student's t distribution with $N - p - 1$ degrees of freedom using

$$t_{b_i} = \frac{b_i - B_i}{s_{b_i}}$$

Usually, the hypothesized value of B_i is zero, but this does not have to be the case.

Confidence Intervals of the Intercept and Coefficients

A $100(1 - \alpha)\%$ confidence interval for the true model coefficient, β_i , is given by

$$b_i \pm (t_{1-\alpha/2, N-p-1}) s_{b_i}$$

Confidence Interval of Y for Given X

A $100(1 - \alpha)\%$ confidence interval for the mean of Y at a specific value of X , say X_0 , is given by

$$b'X_0 \pm (t_{1-\alpha/2, N-p-1}) s_{Y_m|X_0}$$

A $100(1 - \alpha)\%$ prediction interval for the value of Y for an individual at a specific value of X , say X_0 , is given by

$$b'X_0 \pm (t_{1-\alpha/2, N-p-1}) s_{Y_I|X_0}$$

One-Way Analysis of Covariance (ANCOVA)

R² (Percent of Variation Explained)

Several measures of the goodness-of-fit of the model to the data have been proposed, but by far the most popular is R^2 . R^2 is the square of the correlation coefficient between Y and \hat{Y} . It is the proportion of the variation in Y that is accounted by the variation in the independent variables. R^2 varies between zero (no linear relationship) and one (perfect linear relationship).

R^2 , officially known as the *coefficient of determination*, is defined as the sum of squares due to the linear regression model divided by the adjusted total sum of squares of Y . The formula for R^2 is

$$R^2 = 1 - \left(\frac{\mathbf{e}'\mathbf{e}}{\mathbf{Y}'\mathbf{Y} - \frac{(\mathbf{1}'\mathbf{Y})^2}{\mathbf{1}'\mathbf{1}}} \right)$$

$$= \frac{SS_{Model}}{SS_{Total}}$$

R^2 is probably the most popular measure of how well a model fits the data. R^2 may be defined either as a ratio or a percentage. Since we use the ratio form, its values range from zero to one. A value of R^2 near zero indicates no linear relationship, while a value near one indicates a perfect linear fit. Although popular, R^2 should not be used indiscriminately or interpreted without scatter plot support. Following are some qualifications on its interpretation:

1. *Additional independent variables.* It is possible to increase R^2 by adding more independent variables, but the additional independent variables may cause an increase in the mean square error, an unfavorable situation. This usually happens when the sample size is small.
2. *Range of the independent variables.* R^2 is influenced by the range of the independent variables. R^2 increases as the range of the X 's increases and decreases as the range of the X 's decreases.
3. *Slope magnitudes.* R^2 does not measure the magnitude of the slopes.
4. *Linearity.* R^2 does not measure the appropriateness of a linear model. It measures the strength of the linear component of the model. Suppose the relationship between X and Y was a perfect sphere. Although there is a perfect relationship between the variables, the R^2 value would be zero.
5. *Predictability.* A large R^2 does not necessarily mean high predictability, nor does a low R^2 necessarily mean poor predictability.
6. *Sample size.* R^2 is highly sensitive to the number of observations. The smaller the sample size, the larger its value.

 \bar{R}^2 (Adjusted R^2)

R^2 varies directly with N , the sample size. In fact, when $N = p$, $R^2 = 1$. Because R^2 is so closely tied to the sample size, an adjusted R^2 value, called \bar{R}^2 , has been developed. \bar{R}^2 was developed to minimize the impact of sample size. The formula for \bar{R}^2 is

$$\bar{R}^2 = 1 - \frac{(N - 1)(1 - R^2)}{N - p - 1}$$

Least Squares Means

As opposed to raw or arithmetic means which are simply averages of the grouped raw data values, least squares means are adjusted for the other terms in the model, such as covariates. In balanced designs with no covariates, the least squares group means will be equal to the raw group means. In unbalanced designs or when covariates are present, the least squares means usually are different from the raw means.

The least squares means and associated comparisons (i.e., differences) can be calculated using a linear contrast vector, \mathbf{c}_i . Means and differences are estimated as

$$\mathbf{c}_i' \mathbf{b},$$

with estimated standard error,

$$SE(\mathbf{c}_i' \mathbf{b}) = s \sqrt{\mathbf{c}_i' (\mathbf{X}' \mathbf{X})^{-1} \mathbf{c}_i}$$

where s is the square root of the estimated mean square error (MSE) from the model based on v degrees of freedom.

For an ANCOVA model with a group variable with 4 levels and a covariate, and if level 4 were the reference value, the components of the contrast vector would take the form

$$\mathbf{c}_i = (I, \mu_1, \mu_2, \mu_3, X)$$

where I represents the indicator for the intercept and X is the value of the covariate which the mean or difference is evaluated. The contrast vector used to estimate μ_2 would be

$$\mathbf{c}_i = (1, 0, 1, 0, X)$$

The contrast vector used to estimate $\mu_1 - \mu_2$ would be

$$\mathbf{c}_i = (0, 1, -1, 0, 0)$$

Confidence intervals for the estimable functions are of the form

$$\mathbf{c}_i' \mathbf{b} \pm c_\alpha SE(\mathbf{c}_i' \mathbf{b})$$

where c_α is the critical value, usually selected to keep the experimentwise error rate equal to α for the collection of all comparisons (see Multiple Comparisons).

Two-sided significance tests for the mean and the difference (against a null value of zero) use the test statistic

$$T_i = \frac{|\mathbf{c}_i' \mathbf{b}|}{SE(\mathbf{c}_i' \mathbf{b})} \geq c_\alpha$$

Multiple Comparisons

Given that the analysis of variance (ANOVA) test finds a significant difference among treatment means, the next task is to determine which treatments are different. Multiple comparison procedures (MCPs) are methods that pinpoint which treatments are different.

The discussion to follow considers the following experiment. Suppose an experiment studies how two gasoline additives influence the miles per gallon obtained. Three types of gasoline were studied. The first sample received additive W, the second received additive V, and the third did not receive an additive (the control group).

If the F-test from an ANOVA for this experiment is significant, we do not know which of the three possible pairs of groups are different. MCPs can help solve this dilemma.

Multiple Comparison Considerations

Whenever MCPs are to be used, the researcher needs to contemplate the following issues.

Exploration Versus Decision-Making

When conducting exploration (or data snooping), you make several comparisons to discover the underlying factors that influence the response. In this case, you do not have a set of planned comparisons to make. In contrast, in a decision-making mode, you would try to determine which treatment is preferred. In the above example, because you do not know which factors influence gasoline additive performance, you should use the exploration mode to identify those. A decision-making emphasis would choose the gasoline that provides the highest miles per gallon.

Choosing a Comparison Procedure

You should consider two items here. First, will you know before or after experimentation which comparisons are of interest? Second, are you interested in some or all possible comparisons? Your choice of an MCP will depend on how you answer these two questions.

Error Rates

You will need to consider two types of error rates: comparisonwise and experimentwise.

1. Comparisonwise error rate. In this case, you consider each comparison of the means as if it were the only test you were conducting. This is commonly denoted as α . The conceptual unit is the comparison. Other tests that might be conducted are ignored during the calculation of the error rate. If we perform several tests, the probability of a type I error on each test is α .
2. Experimentwise, or familywise, error rate. In this situation, the error rate relates to a group of independent tests. This is the probability of making one or more type I errors in a group of independent comparisons. We will denote this error rate as α_f .

One-Way Analysis of Covariance (ANCOVA)

The relationship between these two error rates is:

$$\alpha_f = 1 - (1 - \alpha)^c$$

where c is the total number of comparisons in the family. The following table shows these error rates for a few values of c and α . The body of the table consists of the calculated values of α_f .

Calculated Experimentwise Error Rates

α	k				
	2	3	5	10	20
0.20	0.360	0.488	0.672	0.893	0.988
0.10	0.190	0.271	0.410	0.651	0.878
0.05	0.098	0.143	0.226	0.401	0.642
0.02	0.040	0.059	0.096	0.183	0.332
0.01	0.020	0.030	0.049	0.096	0.182

As you can see, the possibility of at least one erroneous result goes up markedly as the number of tests increases. For example, to obtain an α_f of 0.05 with a c of 5, you would need to set α to 0.01.

Multiple Comparison Procedures

The multiple comparison procedures (MCPs) considered here assume that there is independence between treatments or samples, equal variance for each treatment, and normality. In addition, unless stated otherwise, the significance tests are assumed to be two-tailed.

Let \bar{y}_i represent the least squares mean of the i^{th} treatment group, $i = 1, \dots, k$. Let s^2 represent the mean square error for these least squares means based on v degrees of freedom. As described above, simultaneous confidence intervals are of the form

$$\mathbf{c}_i' \mathbf{b} \pm c_\alpha SE(\mathbf{c}_i' \mathbf{b}),$$

where c_α is the critical value, usually selected to keep the experimentwise error rate equal to α for the collection of all comparisons. Significance tests are of the form

$$T_i = \frac{|\mathbf{c}_i' \mathbf{b}|}{SE(\mathbf{c}_i' \mathbf{b})} \geq c_\alpha$$

Alpha

This is the α_f , or α , specified for the multiple comparison test. It may be comparisonwise or experimentwise, depending on the test. This alpha usually ranges from 0.01 to 0.10.

One-Way Analysis of Covariance (ANCOVA)

All-Pairs Comparison Procedures

For a group variable with k levels, when comparing all possible pairs, there are $c = k(k - 1)/2$ comparisons.

Tukey-Kramer

The Tukey-Kramer method (also known as Tukey's HSD (Honest Significant Difference) method) uses the Studentized Range distribution to compute the adjustment to c_α . The Tukey-Kramer method achieves the exact alpha level (and simultaneous confidence level $(1 - \alpha)$) if the group sample sizes are equal and is conservative if the sample sizes are unequal. The Tukey-Kramer test is one of the most powerful all-pairs testing procedures and is widely used.

The Tukey-Kramer adjusted critical value for tests and simultaneous confidence intervals is

$$c_\alpha = \frac{q_{1-\alpha, k, v}}{\sqrt{2}}$$

where $q_{1-\alpha, k, v}$ is the $1 - \alpha$ quantile of the studentized range distribution.

Scheffe

This method controls the overall experimentwise error rate but is generally less powerful than the Tukey-Kramer method. Scheffe's method might be preferred if the group sample sizes are unequal or if the number of comparisons is larger than, say, 20.

Scheffe's adjusted critical value for tests and simultaneous confidence intervals is

$$c_\alpha = \sqrt{(k - 1)F_{1-\alpha, k-1, v}}$$

where $F_{1-\alpha, k-1, v}$ is the $1 - \alpha$ quantile of the F distribution with $k - 1$ numerator and v denominator degrees of freedom.

Bonferroni

This conservative method always controls the overall experimentwise error rate (alpha level), even when tests are not independent. P-values are adjusted by multiplying each individual test p-value by the number of comparisons ($c = k(k - 1)/2$) (if the result is greater than 1, then the adjusted p-value is set to 1). Simultaneous confidence limits are adjusted by simply dividing the overall alpha level by the number of comparisons (α/c) and computing each individual interval at $1 - \alpha/c$. Generally, this MCP is run after the fact to find out which pairs are different.

Bonferroni's adjusted critical value for tests and simultaneous confidence intervals is

$$c_\alpha = T_{1-\frac{\alpha}{2c}, v}$$

where $T_{1-\frac{\alpha}{2c}, v}$ is the $1 - \frac{\alpha}{2c}$ quantile of the T distribution with v degrees of freedom.

One-Way Analysis of Covariance (ANCOVA)

Sidak

This method is like Bonferroni's method, but is more powerful if the tests are independent.

Sidak's adjusted critical value for tests and simultaneous confidence intervals is

$$c_{\alpha} = T_{(1-\alpha)^{\frac{1}{c}}/2, v}$$

where $T_{(1-\alpha)^{\frac{1}{c}}/2, v}$ is the $(1 - \alpha)^{\frac{1}{c}}/2$ quantile of the T distribution with v degrees of freedom.

Fisher's LSD (No Adjustment)

The only difference between this test and a regular two-sample T-test is that the degrees of freedom here is based on the whole-model error degrees of freedom, not the sample sizes from the individual groups. This method is not recommended since the overall alpha level is not protected.

The unadjusted critical value for tests and simultaneous confidence intervals is

$$c_{\alpha} = T_{1-\frac{\alpha}{2}, v}$$

where $T_{1-\frac{\alpha}{2}, v}$ is the $1 - \frac{\alpha}{2}$ quantile of the T distribution with v degrees of freedom.

Each vs. Reference Value (or Control) Comparison Procedures

For a group variable with k levels, when comparing each versus the reference value or control, there are $c = k - 1$ comparisons. Often, it is of interest to find only those treatments that are better (or worse) than the control, so both one- and two-sided versions of these tests are provided.

Dunnett's (*Only available for models without covariates*)

This method uses Dunnett's Range Distribution (either two- or one-sided depending on the test direction) to compute the adjustment (see Hsu(1996)). Dunnett's method controls the overall experimentwise error rate and is the most widely used method for all treatments versus control comparisons.

Dunnett's adjusted critical value for tests and simultaneous confidence intervals is

$$c_{\alpha} = q_{1-\alpha, v}$$

where $q_{1-\alpha, v}$ is the $1 - \alpha$ quantile of Dunnett's Range distribution.

One-Way Analysis of Covariance (ANCOVA)

Bonferroni

This conservative method always controls the overall experimentwise error rate (alpha level), even when tests are not independent. P-values are adjusted by multiplying each individual test p-value by the number of comparisons ($c = k - 1$) (if the result is greater than 1, then the adjusted p-value is set to 1).

Simultaneous confidence limits are adjusted by simply dividing the overall alpha level by the number of comparisons (α/c) and computing each individual interval at $1 - \alpha/c$. Generally, this MCP is run after the fact to find out which pairs are different.

Bonferroni's adjusted critical value for tests and simultaneous confidence intervals is

$$c_{\alpha} = T_{1-\frac{\alpha}{2c},v}$$

where $T_{1-\frac{\alpha}{2c},v}$ is the $1 - \frac{\alpha}{2c}$ quantile of the T distribution with v degrees of freedom.

Fisher's LSD (No Adjustment)

The only difference between this test and a regular two-sample T-test is that the degrees of freedom here is based on the whole-model error degrees of freedom, not the sample sizes from the individual groups. This method is not recommended since the overall alpha level is not protected.

The unadjusted critical value for tests and simultaneous confidence intervals is

$$c_{\alpha} = T_{1-\frac{\alpha}{2},v}$$

where $T_{1-\frac{\alpha}{2},v}$ is the $1 - \frac{\alpha}{2}$ quantile of the T distribution with v degrees of freedom.

Recommendations

These recommendations assume that normality and equal variance are valid.

1. Planned all-possible pairs. If you are interested in paired comparisons only and you know this in advance, use either the Bonferroni for pairs or the Tukey-Kramer MCP.
2. Unplanned all-possible pairs. Use Scheffe's MCP.
3. Each versus a control. Use Dunnett's test.

Data Structure

The data must be entered in a format that puts the responses in one column, the group values in a second column, and the covariate values in a third column. An example of data that might be analyzed using this procedure is shown next. The data contains a response variable (Yield), a group variable (TRT), and a covariate (Height). The data could be analyzed as a one-way ANOVA by including only the group variable in the model. Or, the data could be analyzed as an ANCOVA model by including the covariate along with the group variable.

ANCOVA2 Dataset

Yield	Height	TRT
13.2	46	C
13.1	49	C
13.3	49	C
17.6	64	SL
16.6	46	SL
16	36	SL
10.5	53	PS
10.8	63	PS
9.6	42	PS
.	.	.
.	.	.
.	.	.

Example 1 – ANCOVA Model Assuming Unequal Slopes (with Covariate-by-Group Interaction)

This section presents an example of how to run an analysis of the data presented above. These data are contained in the ANCOVA2 dataset. In this example, the two treatments, supplemental lighting (SL) and partial shading (PS), are compared to a control (C) in terms of soybean yield. The comparisons are adjusted for the initial height of the plants. Each treatment is replicated 15 times in a greenhouse study.

This example will run all reports and plots so that they may be documented.

Setup

To run this example, complete the following steps:

1 Open the ANCOVA2 example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **ANCOVA2** and click **OK**.

2 Specify the One-Way Analysis of Covariance (ANCOVA) procedure options

- Find and open the **One-Way Analysis of Covariance (ANCOVA)** 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)	Yield
Group Variable.....	TRT
Covariate Variable	Height
Calc Group Means at.....	Covariate Mean

Reports Tab

Run Summary	Checked
Descriptive Statistics.....	Checked
ANOVA Table	Checked
Coefficient T-Tests.....	Checked
Coefficient C.I.'s	Checked
Least Squares Means.....	Checked
Least Squares Means with Hypothesis	Checked
Tests of H0 Mean = 0	
Compare All Pairs	Checked
Compare Each vs. Reference Group.....	Checked
Residual Normality Tests.....	Checked
Residuals	Checked
Predicted Values for Means.....	Checked
Predicted Values for Individuals.....	Checked

One-Way Analysis of Covariance (ANCOVA)

Plots Tab

Response vs Covariate by Group Scatter.....**Checked**
 Plot
 Means Plots.....**Checked**
 Multiple Comparisons Plots.....**Checked**
 Histogram.....**Checked**
 Probability Plot.....**Checked**
 Resids vs Yhat.....**Checked**
 Resids vs X.....**Checked**

3 Run the procedure

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

Run Summary**Run Summary**

Response Variable: Yield
 Group Variable: TRT
 Reference Group: "C"
 Covariate Variable: Height
 Slopes Assumed to be: Unequal
 Model: Height + TRT + Height*TRT

Parameter	Value	Rows	Value
R ²	0.998	Rows Processed	45
Adjusted R ²	0.9978	Rows Filtered Out	0
Coefficient of Variation	0.0095	Rows with Response Missing	0
Mean Square Error	0.01637464	Rows with Group or Covariate Missing	0
Square Root of MSE	0.1279634	Rows Used in Estimation	45
Average Percent Error	0.74	Completion Status	Normal Completion
Error Degrees of Freedom	39		

This report summarizes the results. It presents the variables used, the model, the number of rows used, and basic summary results.

R²

R^2 , officially known as the *coefficient of determination*, is defined as

$$R^2 = \frac{SS_{Model}}{SS_{Total(Adjusted)}}$$

R^2 is probably the most popular measure of how well a model fits the data. R^2 may be defined either as a ratio or a percentage. Since we use the ratio form, its values range from zero to one. A value of R^2 near zero indicates no linear relationship, while a value near one indicates a perfect linear fit. Although popular, R^2 should not be used indiscriminately or interpreted without scatter plot support. Following are some qualifications on its interpretation:

One-Way Analysis of Covariance (ANCOVA)

1. *Additional independent variables.* It is possible to increase R^2 by adding more independent variables, but the additional independent variables may cause an increase in the mean square error, an unfavorable situation. This usually happens when the sample size is small.
2. *Range of the independent variables.* R^2 is influenced by the range of the independent variables. R^2 increases as the range of the X 's increases and decreases as the range of the X 's decreases.
3. *Slope magnitudes.* R^2 does not measure the magnitude of the slopes.
4. *Linearity.* R^2 does not measure the appropriateness of a linear model. It measures the strength of the linear component of the model. Suppose the relationship between X and Y was a perfect sphere. Although there is a perfect relationship between the variables, the R^2 value would be zero.
5. *Predictability.* A large R^2 does not necessarily mean high predictability, nor does a low R^2 necessarily mean poor predictability.
6. *Sample size.* R^2 is highly sensitive to the number of observations. The smaller the sample size, the larger its value.

Adjusted R^2

This is an adjusted version of R^2 . The adjustment seeks to remove the distortion due to a small sample size. The formula for adjusted R^2 is

$$\bar{R}^2 = 1 - \frac{(N - 1)(1 - R^2)}{N - p - 1}$$

Coefficient of Variation

The coefficient of variation is a relative measure of dispersion, computed by dividing root mean square error by the mean of the response variable. By itself, it has little value, but it can be useful in comparative studies.

$$CV = \frac{\sqrt{MSE}}{\bar{y}}$$

Average |Percent Error|


This is the average of the absolute percent errors. It is another measure of the goodness of fit of the linear model to the data. It is calculated using the formula

$$AAPE = \frac{100 \sum_{j=1}^N \left| \frac{y_j - \hat{y}_j}{y_j} \right|}{N}$$

Note that when the response variable is zero, its predicted value is used in the denominator.

Response vs Covariate by Group Scatter Plot

Response vs. Covariate by Group Scatter Plot



This is a scatter plot with the response variable, Yield, on the Y-axis, the covariate variable, Height, on the X-axis, and the group variable, TRT, in the legend. The slopes appear to be quite equal among the groups, which is one indication that the height*TRT interaction is probably not significant.

Descriptive Statistics

Descriptive Statistics

Variable	Count	Mean	Standard Deviation	Minimum	Maximum
Height	45	52.13334	9.198814	34	68
(TRT="PS")	45	0.3333333	0.4767313	0	1
(TRT="SL")	45	0.3333333	0.4767313	0	1
Height*(TRT="PS")	45	18.71111	27.18675	0	68
Height*(TRT="SL")	45	16.71111	24.44718	0	64
Yield	45	13.52889	2.697947	9.5	17.7

For each variable, the count, arithmetic mean, standard deviation, minimum, and maximum are computed. This report is particularly useful for checking that the correct variables were selected. Recall that the group variable with K levels is represented by $K - 1$ binary indicator variables. The reference group is not listed.

One-Way Analysis of Covariance (ANCOVA)

Analysis of Variance

Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F-Ratio	P-Value	Significant at $\alpha = 0.05$?
Model	5	319.6338	63.92677	3904.011	0.0000	Yes
Height	1	11.4385	11.4385	698.550	0.0000	Yes
TRT	2	10.16219	5.081096	310.303	0.0000	Yes
Height*TRT	2	0.07937068	0.03968534	2.424	0.1018	No*
Error	39	0.6386109	0.01637464			
Total(Adjusted)	44	320.2724	7.278919			

* The covariate-by-group interaction term is not significant at $\alpha = 0.05$. You may want to run the analysis again with the assumption that the slopes are equal, which excludes the interaction.

An analysis of variance (ANOVA) table summarizes the information related to the variation in data. Both Height and TRT are significant. The Height*TRT interaction is not significant, indicating that it is probably safe to assume that the slopes are equal. Example 2 shows the results with the assumption of equal slopes (without the interaction term in the model).

Source

This represents a partition of the variation in Y .

DF

The degrees of freedom are the number of dimensions associated with this term. Note that each observation can be interpreted as a dimension in n -dimensional space. The degrees of freedom for the intercept, model, error, and adjusted total are 1, p , $n-p-1$, and $n-1$, respectively.

Sum of Squares

These are the sums of squares associated with the corresponding sources of variation. Note that these values are in terms of the dependent variable. The formulas for each are

$$SS_{Model} = \sum (\hat{y}_j - \bar{y})^2$$

$$SS_{Error} = \sum (y_j - \hat{y}_j)^2$$

$$SS_{Total} = \sum (y_j - \bar{y})^2$$

Mean Square

The mean square is the sum of squares divided by the degrees of freedom. This mean square is an estimated variance. For example, the mean square error is the estimated variance of the residuals.

One-Way Analysis of Covariance (ANCOVA)

F-Ratio

This is the F -statistic for testing the null hypothesis that all $\beta_j = 0$. This F -statistic has p degrees of freedom for the numerator variance and $n-p-1$ degrees of freedom for the denominator variance.

P-Value

This is the p -value for the above F -test. The p -value is the probability that the test statistic will take on a value at least as extreme as the observed value, if the null hypothesis is true. If the p -value is less than α , say 0.05, the null hypothesis is rejected. If the p -value is greater than α , then the null hypothesis is accepted.

Significant at $\alpha = 0.05$?

This is the decision based on the p -value and the user-entered Tests Alpha value. The default is Tests Alpha = 0.05.

Model Coefficient T-Tests**Model Coefficient T-Tests**

Hypotheses Tested: $H_0: \beta(i) = 0$ vs. $H_1: \beta(i) \neq 0$

Independent Variable	Model Coefficient b(i)	Standard Error Sb(i)	Hypothesis Tests		
			T-Statistic	P-Value	Reject H0 at $\alpha = 0.05$?
Intercept	10.44549	0.1890552	55.251	0.0000	Yes
Height	0.05614055	0.003713012	15.120	0.0000	Yes
(TRT="PS")	-3.649902	0.2963626	-12.316	0.0000	Yes
(TRT="SL")	3.714274	0.2686974	13.823	0.0000	Yes
Height*(TRT="PS")	0.009258767	0.005474596	1.691	0.0988	No
Height*(TRT="SL")	-0.002279411	0.005277996	-0.432	0.6682	No

This section reports the values and significance tests of the model coefficients.

Independent Variable

The names of the independent variables are listed here. The intercept is the value of the Y intercept.

Note that the name may become very long, especially for interaction terms. These long names may misalign the report. You can force the rest of the items to be printed on the next line by using the Stagger label ... option on the Report Options tab. This should create a better-looking report when the names are extra-long.

Model Coefficient b(i)

The coefficients are the least squares estimates of the parameters. The value indicates how much change in Y occurs for a one-unit change in a particular X when the remaining X 's are held constant. These coefficients are often called partial-regression coefficients since the effect of the other X 's is removed. These coefficients are the values of b_0, b_1, \dots, b_p .

One-Way Analysis of Covariance (ANCOVA)

Standard Error Sb(i)

The standard error of the coefficient, s_{b_j} , is the standard deviation of the estimate. It is used in hypothesis tests and confidence limits.

T-Statistic

This is the t-test value for testing the hypothesis that $\beta_j = 0$ versus the alternative that $\beta_j \neq 0$ after removing the influence of all other X 's. This t -value has $n-p-1$ degrees of freedom.

To test for a value other than zero, use the formula below. There is an easier way to test hypothesized values using confidence limits. See the discussion below under Confidence Limits. The formula for the t -test is

$$t_j = \frac{b_j - \beta_j^*}{s_{b_j}}$$

P-Value

This is the p -value for the significance test of the coefficient. The p -value is the probability that this t -statistic will take on a value at least as extreme as the observed value, assuming that the null hypothesis is true (i.e., the coefficient estimate is equal to zero). If the p -value is less than alpha, say 0.05, the null hypothesis of equality is rejected. This p -value is for a two-tail test.

Reject H0 at $\alpha = 0.05$?

This is the decision based on the p -value and the user-entered Tests Alpha value. The default is Tests Alpha = 0.05.

Model Coefficient Confidence Intervals**Model Coefficient Confidence Intervals**

Independent Variable	Model Coefficient b(i)	Standard Error Sb(i)	95% Confidence Interval Limits for $\beta(i)$	
			Lower	Upper
Intercept	10.44549	0.1890552	10.06309	10.82789
Height	0.05614055	0.003713012	0.04863027	0.06365082
(TRT="PS")	-3.649902	0.2963626	-4.249352	-3.050452
(TRT="SL")	3.714274	0.2686974	3.170783	4.257766
Height*(TRT="PS")	0.009258767	0.005474596	-0.001814648	0.02033218
Height*(TRT="SL")	-0.002279411	0.005277996	-0.01295517	0.008396343

Note: The T-Value used to calculate these confidence interval limits was 2.023.

This section reports the values and confidence intervals of the model coefficients.

One-Way Analysis of Covariance (ANCOVA)

Independent Variable

The names of the independent variables are listed here. The intercept is the value of the Y intercept.

Note that the name may become very long, especially for interaction terms. These long names may misalign the report. You can force the rest of the items to be printed on the next line by using the Stagger label ... option on the Report Options tab. This should create a better-looking report when the names are extra-long.

Model Coefficient

The coefficients are the least squares estimates of the parameters. The value indicates how much change in Y occurs for a one-unit change in a particular X when the remaining X 's are held constant. These coefficients are often called partial-regression coefficients since the effect of the other X 's is removed. These coefficients are the values of b_0, b_1, \dots, b_p .

Standard Error

The standard error of the coefficient, s_{b_j} , is the standard deviation of the estimate. It is used in hypothesis tests and confidence limits.

95% Confidence Interval Limits for $\beta(i)$

These are the lower and upper values of a $100(1 - \alpha)\%$ interval estimate for β_j based on a t -distribution with $n-p-1$ degrees of freedom. This interval estimate assumes that the residuals for the regression model are normally distributed.

These confidence limits may be used for significance testing values of β_j other than zero. If a specific value is not within this interval, it is significantly different from that value. Note that these confidence limits are set up as if you are interested in each regression coefficient separately.

The formulas for the lower and upper confidence limits are:

$$b_j \pm t_{1-\alpha/2, n-p-1} s_{b_j}$$

Note: The T-Value ...

This is the value of $t_{1-\alpha/2, n-p-1}$ used to construct the confidence limits.

Least Squares Means

Least Squares Means

Error Degrees of Freedom (DF): 39

Means Calculated at: Height = 52.13334 (Mean)

Name	Count	Least Squares Mean	Standard Error	95% Confidence Interval Limits for the Mean	
				Lower	Upper
Intercept					
All	45	13.51502	0.02012547	13.47432	13.55573
TRT					
C	15	13.37228	0.03386426	13.30378	13.44078
PS	15	10.20507	0.03675049	10.13073	10.27940
SL	15	16.96772	0.03388105	16.89919	17.03625

Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

This section reports the least squares means and associated confidence intervals. In this example, the least squares means are calculated at the mean of the covariate, Height = 52.13334. The results are based on $n-p-1 = 39$ degrees of freedom for error.

Name

The name of the group variable and its individual group names are listed here. The intercept is the value of the Y intercept.

Note that the name may become very long, especially for interaction terms. These long names may misalign the report. You can force the rest of the items to be printed on the next line by using the Stagger label ... option on the Report Options tab. This should create a better-looking report when the names are extra-long.

Count

This column specifies the number of observations in each group.

Least Squares Mean

This is the least squares mean estimate, $\hat{\mu}_j$. The least squares means are adjusted based on the model. In balanced designs with no covariates, the least squares group means will be equal to the raw group means. In unbalanced designs or when covariates are present, the least squares means usually are different from the raw means.

Standard Error

The standard error of the mean, $SE(\hat{\mu}_j)$, is the standard deviation of the estimate. It is used in hypothesis tests and confidence limits.

One-Way Analysis of Covariance (ANCOVA)

95% Confidence Interval Limits for the Mean

These are the lower and upper values of a $100(1 - \alpha)\%$ interval estimate for the mean, μ_j , based on a t -distribution with $n-p-1$ degrees of freedom.

The formulas for the lower and upper confidence limits are:

$$\hat{\mu}_j \pm t_{1-\frac{\alpha}{2}, n-p-1} \times SE(\hat{\mu}_j)$$

Least Squares Means with Hypothesis Tests of H0: Mean = 0

Least Squares Means with Hypothesis Tests of H0: Mean = 0

Error Degrees of Freedom (DF): 39

Means Calculated at: Height = 52.13334 (Mean)

Hypotheses Tested: H0: Mean = 0 vs. H1: Mean \neq 0

				Hypothesis Tests		
Name	Count	Least Squares Mean	Standard Error	T-Statistic	P-Value	Reject H0 at $\alpha = 0.05$?
Intercept						
All	45	13.51502	0.02012547	671.538	0.0000	Yes
TRT						
C	15	13.37228	0.03386426	394.879	0.0000	Yes
PS	15	10.20507	0.03675049	277.685	0.0000	Yes
SL	15	16.96772	0.03388105	500.803	0.0000	Yes

Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

This section reports the least squares means and associated hypothesis tests. In this example, the least squares means are calculated at the mean of the covariate, Height = 52.13334. The results are based on $n-p-1 = 39$ degrees of freedom for error.

Name

The name of the group variable and its individual group names are listed here. The intercept is the value of the Y intercept.

Note that the name may become very long, especially for interaction terms. These long names may misalign the report. You can force the rest of the items to be printed on the next line by using the Stagger label ... option on the Report Options tab. This should create a better-looking report when the names are extra-long.

Count

This column specifies the number of observations in each group.

One-Way Analysis of Covariance (ANCOVA)

Least Squares Mean

This is the least squares mean estimate, $\hat{\mu}_j$. The least squares means are adjusted based on the model. In balanced designs with no covariates, the least squares group means will be equal to the raw group means. In unbalanced designs or when covariates are present, the least squares means usually are different from the raw means.

Standard Error

The standard error of the mean, $SE(\hat{\mu}_j)$, is the standard deviation of the estimate. It is used in hypothesis tests and confidence limits.

T-Statistic

This is the t-test value for testing the hypothesis that the mean is equal to 0 versus the alternative that it is not equal to 0. This t -value has $n-p-1$ degrees of freedom and is calculated as

$$t_j = \frac{\hat{\mu}_j}{SE(\hat{\mu}_j)}$$

P-Value


This is the p -value for the significance test of the mean. The p -value is the probability that this t -statistic will take on a value at least as extreme as the observed value, if the null hypothesis is true (i.e., the mean estimate is equal to zero). If the p -value is less than alpha, say 0.05, the null hypothesis of equality is rejected. This p -value is for a two-tail test.

Reject H0 at $\alpha = 0.05$?

This is the decision based on the p -value and the user-entered Tests Alpha value. The default is Tests Alpha = 0.05.

Means Plots

Means Plots



The means plot displays the least squares means along with user-selected variability lines, in this case \pm SE.

All-Pairs Comparisons of Least Squares Means

All-Pairs Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 39
 Means Calculated at: Height = 52.13334 (Mean)
 Multiple Comparison Type: Tukey-Kramer
 Hypotheses Tested: H0: Diff = 0 vs. H1: Diff \neq 0
 Number of Comparisons: 3

Comparison	Least Squares Mean Difference	Standard Error	Hypothesis Tests			
			T-Statistic	P-Value		Reject H0 at $\alpha = 0.05$?†
				Unadjusted	Adjusted*	
TRT						
C - PS	3.167212	0.04997386	63.377	0.0000	0.0000	Yes
C - SL	-3.595441	0.04790317	-75.056	0.0000	0.0000	Yes
PS - SL	-6.762653	0.04998524	-135.293	0.0000	0.0000	Yes

* Adjusted p-values are computed using the number of comparisons (3) and the adjustment type (Tukey-Kramer).

† Rejection decisions are based on adjusted p-values.

Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

This section reports the least squares mean differences and associated multiple comparison hypothesis tests for all pairs. In this example, the least squares means are calculated at the mean of the covariate, Height = 52.13334. The results are based on $n-p-1 = 39$ degrees of freedom for error.

You should only consider these tests if the group variable was found to be significant in the ANOVA table. All treatment group means (C, PS, and SL) are found here to be significantly different from one another based on the Tukey-Kramer-adjusted p -values, which are based on 3 comparisons.

Comparison

The name of the group variable and the individual comparisons are listed here. The multiple comparison adjustment is based on the selected Multiple Comparison Adjustment Type and the number of simultaneous comparisons.

Note that the name may become very long, especially for interaction terms. These long names may misalign the report. You can force the rest of the items to be printed on the next line by using the Stagger label ... option on the Report Options tab. This should create a better-looking report when the group names are extra-long.

Least Squares Mean Difference

This is the least squares mean difference estimate for group i minus group j , $\hat{\mu}_i - \hat{\mu}_j$.

Standard Error

The least squares mean difference estimate, $SE(\hat{\mu}_i - \hat{\mu}_j)$, is the standard deviation of the estimate. It is used in hypothesis tests and confidence limits.

One-Way Analysis of Covariance (ANCOVA)

T-Statistic

This is the t-test value for testing the hypothesis that the mean difference is equal to 0 versus the alternative that it is not equal to 0. This t -value has $n-p-1$ degrees of freedom and is calculated as

$$t_j = \frac{\hat{\mu}_i - \hat{\mu}_j}{SE(\hat{\mu}_i - \hat{\mu}_j)}$$

Unadjusted P-Value

This is the unadjusted p -value for the significance test of the mean difference and assumes no multiple comparisons. This p -value is valid if you are only interested in one of the comparisons. You can use this p -value to make your own adjustment (e.g. Bonferroni) if needed.

Adjusted P-Value

This is the adjusted p -value for the significance test of the mean difference. Adjusted p -values are computed using the number of comparisons and the multiple comparison adjustment type, which in this example is Tukey-Kramer.

Reject H0 at $\alpha = 0.05$?

This is the decision based on the adjusted p -value and the user-entered Tests Alpha value. The default is Tests Alpha = 0.05.

Simultaneous Confidence Intervals for All-Pairs Comparisons of Least Squares Means

Simultaneous Confidence Intervals for All-Pairs Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 39
 Means Calculated at: Height = 52.13334 (Mean)
 Multiple Comparison Type: Tukey-Kramer
 Number of Comparisons: 3

Comparison	Least Squares Mean Difference	Standard Error	95% Simultaneous Confidence Interval Limits*	
			Lower	Upper
TRT				
C - PS	3.167212	0.04997386	3.045453	3.28897
C - SL	-3.595441	0.04790317	-3.712155	-3.478728
PS - SL	-6.762653	0.04998524	-6.884439	-6.640867

* Confidence interval limits are adjusted based on the number of comparisons (3) and the adjustment type (Tukey-Kramer).
 Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

This section reports the least squares mean differences and associated multiple comparison simultaneous confidence intervals for all pairs. In this example, the least squares means are calculated at the mean of the covariate, Height = 52.13334. The results are based on $n-p-1 = 39$ degrees of freedom for error.

One-Way Analysis of Covariance (ANCOVA)

Since the simultaneous confidence intervals are adjusted for the multiplicity of tests, there is a one-to-one correspondence between the intervals and the hypothesis tests--- all differences for which the $100(1 - \alpha)\%$ simultaneous confidence interval does not include zero will be significant at level α . As you can see, the 95% simultaneous confidence intervals for all group mean differences do not include 0, so the corresponding tests (C, PS, and SL) are found here to be significantly different from one another.

Comparison

The name of the group variable and the individual comparisons are listed here. The multiple comparison adjustment is based on the selected Multiple Comparison Adjustment Type and the number of simultaneous comparisons.

Note that the name may become very long, especially for interaction terms. These long names may misalign the report. You can force the rest of the items to be printed on the next line by using the Stagger label ... option on the Report Options tab. This should create a better-looking report when the group names are extra-long.

Least Squares Mean Difference

This is the least squares mean difference estimate for group i minus group j , $\hat{\mu}_i - \hat{\mu}_j$.

Standard Error

The least squares mean difference estimate, $SE(\hat{\mu}_i - \hat{\mu}_j)$, is the standard deviation of the estimate. It is used in hypothesis tests and confidence limits.

95% Simultaneous Confidence Interval Limits

These are the lower and upper values of a $100(1 - \alpha)\%$ simultaneous confidence interval estimates for the mean difference, $\mu_i - \mu_j$. The formulas for the lower and upper confidence limits are:

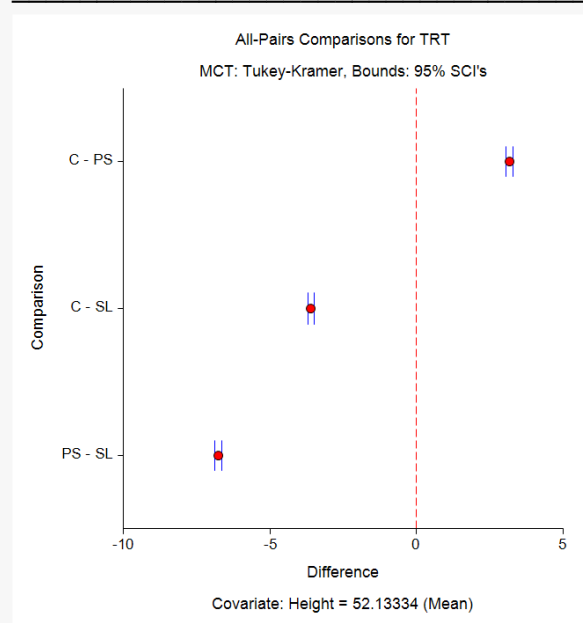
$$\hat{\mu}_i - \hat{\mu}_j \pm c_\alpha \times SE(\hat{\mu}_i - \hat{\mu}_j)$$

where c_α is the adjusted critical value, usually selected to keep the experimentwise error rate equal to α for the collection of all comparisons.

One-Way Analysis of Covariance (ANCOVA)

All-Pairs Comparisons Plots

All-Pairs Comparisons Plots



This multiple comparisons plot displays the mean differences along with 95% simultaneous confidence intervals. All comparisons for which the interval does not contain zero are significant.

Each vs. Reference Value Comparisons of Least Squares Means

Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 39
 Means Calculated at: Height = 52.13334 (Mean)
 Multiple Comparison Type: Bonferroni
 Hypotheses Tested: H0: Diff = 0 vs. H1: Diff \neq 0
 Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	Hypothesis Tests			
			T-Statistic	P-Value		Reject H0 at $\alpha = 0.05$?†
				Unadjusted	Adjusted*	
TRT						
PS - C	-3.167212	0.04997386	-63.377	0.0000	0.0000	Yes
SL - C	3.595441	0.04790317	75.056	0.0000	0.0000	Yes

* Adjusted p-values are computed using the number of comparisons (2) and the adjustment type (Bonferroni).

† Rejection decisions are based on adjusted p-values.

Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

One-Way Analysis of Covariance (ANCOVA)

This section reports the least squares mean differences and associated multiple comparison hypothesis tests for each group versus the reference group. In this example, the least squares means are calculated at the mean of the covariate, Height = 52.13334. The results are based on $n-p-1 = 39$ degrees of freedom for error.

You should only consider these tests if the group variable was found to be significant in the ANOVA table. Both PS and SL are found here to be significantly different from the control, C, based on the Bonferroni-adjusted p -values, which are based on 2 comparisons.

Often, Dunnett's test is used to compare individual groups to the control, but when there are covariates in the model, only the Bonferroni adjustment is available in NCSS.

Simultaneous Confidence Intervals for Each v. Reference Value Comparisons of Least Squares Means

Simultaneous Confidence Intervals for Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 39
 Means Calculated at: Height = 52.13334 (Mean)
 Multiple Comparison Type: Bonferroni
 Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	95% Simultaneous Confidence Interval Limits*	
			Lower	Upper
TRT				
PS - C	-3.167212	0.04997386	-3.283714	-3.050709
SL - C	3.595441	0.04790317	3.483766	3.707116


* Confidence interval limits are adjusted based on the number of comparisons (2) and the adjustment type (Bonferroni).
 Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

This section reports the least squares mean differences and associated multiple comparison simultaneous confidence intervals for each group versus the reference group. In this example, the least squares means are calculated at the mean of the covariate, Height = 52.13334. The results are based on $n-p-1 = 39$ degrees of freedom for error.

Since the simultaneous confidence intervals are adjusted for the multiplicity of tests, there is a one-to-one correspondence between the intervals and the hypothesis tests--- all differences for which the $100(1 - \alpha)\%$ simultaneous confidence interval does not include zero will be significant at level α . As you can see, the 95% simultaneous confidence intervals for all group mean differences do not include 0, so the corresponding tests (PS vs. C and SL vs. C) are found here to be significant.

Each vs. Reference Group Comparisons Plots

Each vs. Reference Group Comparisons Plots



This multiple comparisons plot displays the mean differences along with 95% simultaneous confidence intervals. All comparisons for which the interval does not contain zero are significant.

Residual Normality Assumption Tests

Residual Normality Assumption Tests

Normality Test	Test Statistic	P-Value	Reject Residual Normality at $\alpha = 0.2$?
Shapiro-Wilk	0.983	0.7316	No
Anderson-Darling	0.237	0.7867	No
D'Agostino Skewness	0.967	0.3335	No
D'Agostino Kurtosis	-0.079	0.9370	No
D'Agostino Omnibus (Skewness and Kurtosis)	0.942	0.6245	No

This report gives the results of applying several normality tests to the residuals. The Shapiro-Wilk test is probably the most popular, so it is given first. These tests are discussed in detail in the Normality Tests section of the Descriptive Statistics procedure.

Graphic Residual Analysis

The residuals can be graphically analyzed in numerous ways. You should examine all of the basic residual graphs: the histogram, the density trace, the normal probability plot, the scatter plot of the residuals versus the predicted value of the dependent variable, and the scatter plot of the residuals versus each of the independent variables.

For the basic scatter plots of residuals versus either the predicted values of Y or the independent variables, Hoaglin (1983) explains that there are several patterns to look for. You should note that these patterns are very difficult, if not impossible, to recognize for small data sets.

Point Cloud

A point cloud, basically in the shape of a rectangle or a horizontal band, would indicate no relationship between the residuals and the variable plotted against them. This is the preferred condition.

Wedge

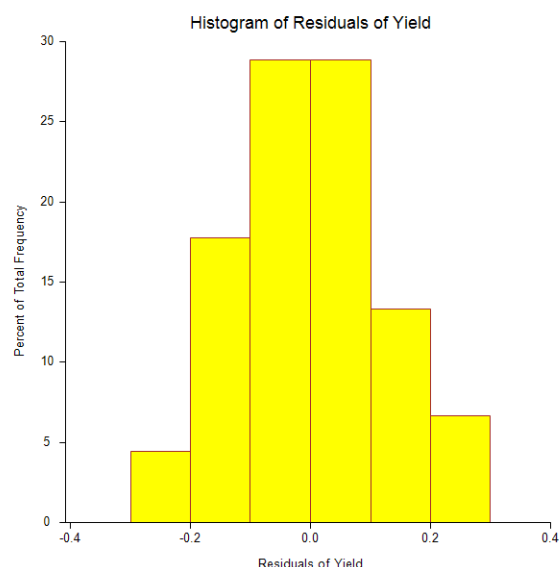
An increasing or decreasing wedge would be evidence that there is increasing or decreasing (non-constant) variation. A transformation of Y may correct the problem.

Bowtie

This is similar to the wedge above in that the residual plot shows a decreasing wedge in one direction while simultaneously having an increasing wedge in the other direction. A transformation of Y may correct the problem.

Histogram of Residuals

Residual Analysis Plots




One-Way Analysis of Covariance (ANCOVA)

The purpose of the histogram and density trace of the residuals is to evaluate whether they are normally distributed. Unless you have a large sample size, it is best not to rely on the histogram for visually evaluating normality of the residuals. The better choice would be the normal probability plot.

Probability Plot of Residuals

Residual Analysis Plots




If the residuals are normally distributed, the data points of the normal probability plot will fall along a straight line through the origin with a slope of 1.0. Major deviations from this ideal picture reflect departures from normality. Stragglers at either end of the normal probability plot indicate outliers, curvature at both ends of the plot indicates long or short distributional tails, convex or concave curvature indicates a lack of symmetry, and gaps or plateaus or segmentation in the normal probability plot may require a closer examination of the data or model. Of course, use of this graphic tool with very small sample sizes is not recommended.

If the residuals are not normally distributed, then the t-tests on regression coefficients, the F-tests, and any interval estimates are not valid. This is a critical assumption to check.

Residuals vs. Yhat (Predicted) Plot

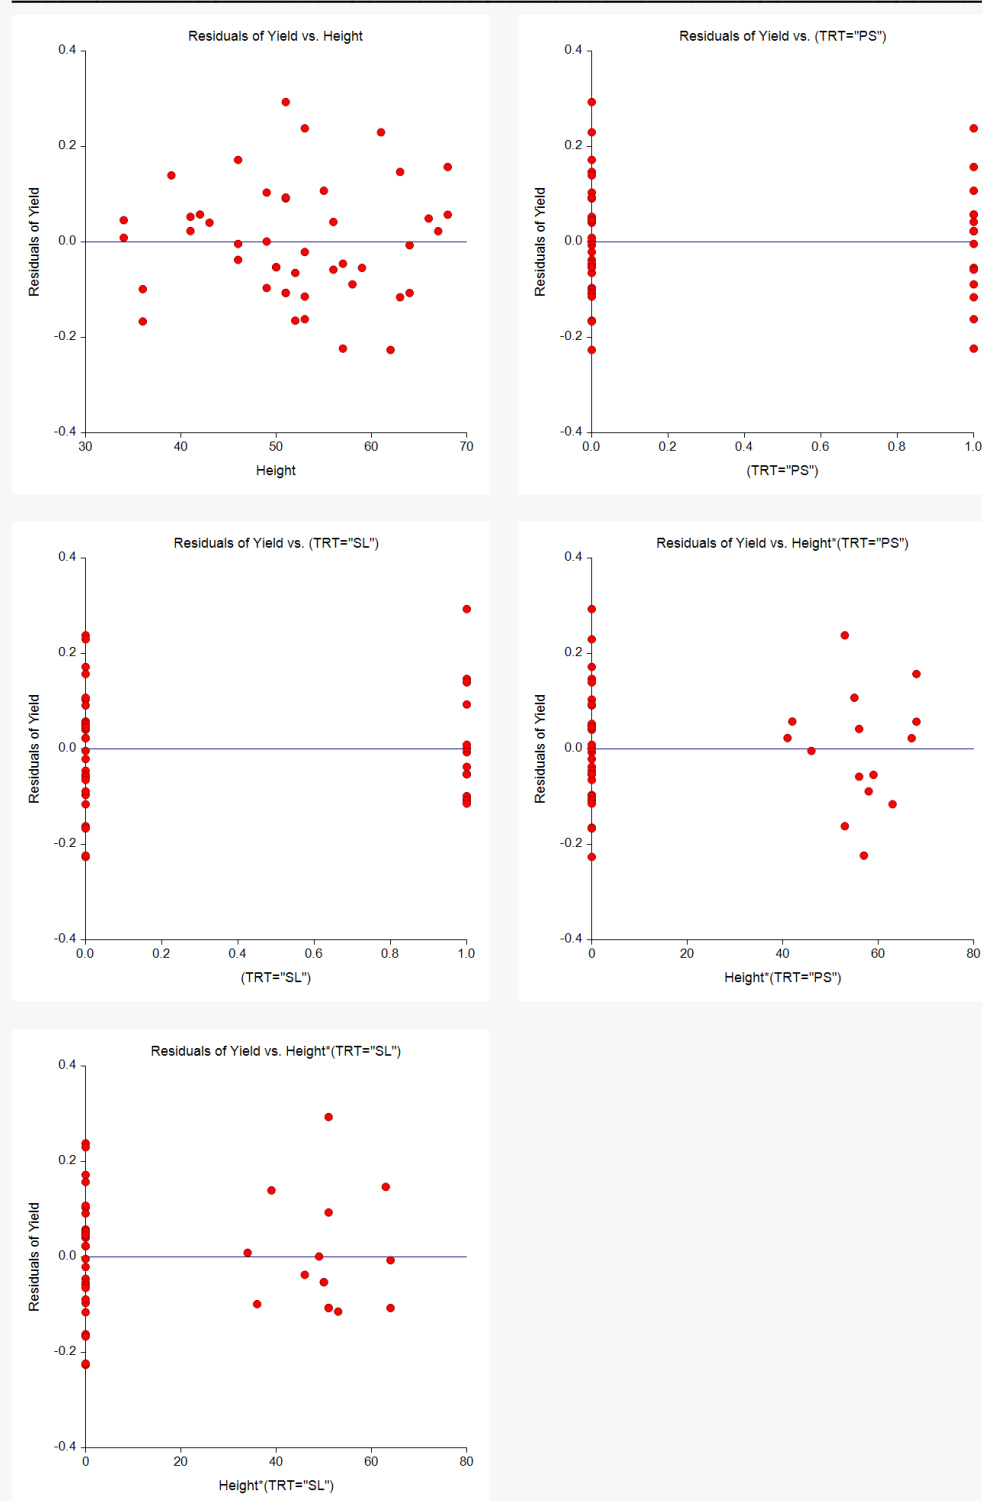
Residual Analysis Plots



This plot should always be examined. The preferred pattern to look for is a point cloud or a horizontal band. A wedge or bowtie pattern is an indicator of non-constant variance, a violation of a critical assumption. The sloping or curved band signifies inadequate specification of the model. The sloping band with increasing or decreasing variability suggests non-constant variance and inadequate specification of the model.

Residuals vs. X Plots

Residual Analysis Plots



These are scatter plots of the residuals versus each independent variable. Again, the preferred pattern is a rectangular shape or point cloud. Any other nonrandom pattern may require a redefining of the model.

One-Way Analysis of Covariance (ANCOVA)

Residuals List

Residuals List

Row	Yield		Residual	Absolute Percent Error	Sqrt(MSE) Without This Row
	Actual	Predicted			
1	13.2	13.027950	0.172047600	1.303	0.1263245
2	13.1	13.196370	-0.096374040	0.736	0.1286210
3	13.3	13.196370	0.103626000	0.779	0.1284618
4	17.6	17.606870	-0.006874427	0.039	0.1296300
5	16.6	16.637370	-0.037373970	0.225	0.1294818
6	16.0	16.098760	-0.098762600	0.617	0.1283299
7	10.5	10.261750	0.238251200	2.269	0.1232406
8	10.8	10.915740	-0.115742000	1.072	0.1280937
9	9.6	9.542356	0.057643650	0.600	0.1291771
10	13.4	13.420940	-0.020936240	0.156	0.1295882
11	14.1	13.870060	0.229939400	1.631	0.1230328
12	13.2	13.364800	-0.164795700	1.248	0.1266389
13	16.8	16.906680	-0.106679600	0.635	0.1283918
14	16.8	16.906680	-0.106679600	0.635	0.1283918
15	17.2	16.906680	0.293320400	1.705	0.1199084
16	10.5	10.392550	0.107452600	1.023	0.1283727
17	10.1	10.261750	-0.161748800	1.601	0.1267286
18	11.4	11.242740	0.157261500	1.379	0.1264359
19	12.9	12.859530	0.040469240	0.314	0.1294494
20	13.7	13.926200	-0.226201200	1.651	0.1230973
.
.
.

This section reports on the sample residuals, or *e*'s.

Actual

This is the actual value of *Y*.

Predicted

The predicted value of *Y* using the values of the independent variables given on this row.

Residual

This is the error in the predicted value. It is equal to the *Actual* minus the *Predicted*.

Absolute Percent Error

This is percentage that the absolute value of the *Residual* is of the *Actual* value. Scrutinize rows with the large percent errors.

Sqrt(MSE) Without This Row

This is the value of the square root of the mean square error that is obtained if this row is deleted. A perusal of this statistic for all observations will highlight observations that have an inflationary impact on mean square error and could be outliers.

Predicted Values with Confidence Intervals for the Means

Predicted Values with Confidence Intervals for the Means

Row	Yield		Standard Error of Predicted	95% Confidence Interval Limits for the Mean	
	Actual	Predicted		Lower	Upper
1	13.2	13.027950	0.03643043	12.954260	13.101640
2	13.1	13.196370	0.03330692	13.129000	13.263740
3	13.3	13.196370	0.03330692	13.129000	13.263740
4	17.6	17.606870	0.06162171	17.482230	17.731520
5	16.6	16.637370	0.03649704	16.563550	16.711200
6	16.0	16.098760	0.06246837	15.972410	16.225120
7	10.5	10.261750	0.03536298	10.190220	10.333280
8	10.8	10.915740	0.04306705	10.828630	11.002850
9	9.6	9.542356	0.06576140	9.409342	9.675371
10	13.4	13.420940	0.03471220	13.350720	13.491150
11	14.1	13.870060	0.05214987	13.764580	13.975540
12	13.2	13.364800	0.03375916	13.296510	13.433080
13	16.8	16.906680	0.03319957	16.839530	16.973830
14	16.8	16.906680	0.03319957	16.839530	16.973830
15	17.2	16.906680	0.03319957	16.839530	16.973830
16	10.5	10.392550	0.03335312	10.325080	10.460010
17	10.1	10.261750	0.03536298	10.190220	10.333280
18	11.4	11.242740	0.05805809	11.125310	11.360170
19	12.9	12.859530	0.04234571	12.773880	12.945180
20	13.7	13.926200	0.05507287	13.814810	14.037600
.
.
.

Confidence intervals for the mean response of Y given specific levels for the group and covariate variables are provided here. It is important to note that violations of any assumptions will invalidate these interval estimates.

Actual

This is the actual value of Y .

Predicted

The predicted value of Y . It is predicted using the values of the group and covariate variables for this row. If the input data had both group and covariate values but no value for Y , the predicted value is still provided.

Standard Error of Predicted

This is the standard error of the mean response for the specified values of the group and covariate variables. Note that this value is not constant for all variable values. In fact, it is a minimum at the average value of each group and covariate variable.

95% Confidence Interval Limits for the Mean

These are the lower and upper limits of a 95% confidence interval estimate of the mean of Y for this observation.

Predicted Values with Prediction Intervals for Individuals

Predicted Values with Prediction Intervals for Individuals

Row	Yield		Standard Error of Predicted	95% Prediction Interval Limits for an Individual	
	Actual	Predicted		Lower	Upper
1	13.2	13.027950	0.1330482	12.758840	13.297070
2	13.1	13.196370	0.1322270	12.928920	13.463830
3	13.3	13.196370	0.1322270	12.928920	13.463830
4	17.6	17.606870	0.1420277	17.319600	17.894150
5	16.6	16.637370	0.1330664	16.368220	16.906530
6	16.0	16.098760	0.1423971	15.810740	16.386790
7	10.5	10.261750	0.1327599	9.993217	10.530280
8	10.8	10.915740	0.1350163	10.642650	11.188840
9	9.6	9.542356	0.1438722	9.251348	9.833365
10	13.4	13.420940	0.1325880	13.152750	13.689120
11	14.1	13.870060	0.1381819	13.590560	14.149560
12	13.2	13.364800	0.1323417	13.097110	13.632480
13	16.8	16.906680	0.1322000	16.639280	17.174080
14	16.8	16.906680	0.1322000	16.639280	17.174080
15	17.2	16.906680	0.1322000	16.639280	17.174080
16	10.5	10.392550	0.1322387	10.125070	10.660030
17	10.1	10.261750	0.1327599	9.993217	10.530280
18	11.4	11.242740	0.1405182	10.958510	11.526960
19	12.9	12.859530	0.1347880	12.586900	13.132160
20	13.7	13.926200	0.1393114	13.644420	14.207980
.
.
.

A prediction interval for the individual response of Y given specific values of the group and covariate variables is provided here for each row.

Actual

This is the actual value of Y .

Predicted

The predicted value of Y . It is predicted using the values of the group and covariate variables for this row. If the input data had both group and covariate values but no value for Y , the predicted value is still provided.

Standard Error of Predicted

This is the standard error of the mean response for the specified values of the group and covariate variables. Note that this value is not constant for all variable values. In fact, it is a minimum at the average value of the group and covariate variable.

95% Prediction Interval Limits for the Mean

These are the lower and upper limits of a 95% prediction interval of the individual value of Y for this observation.

Example 2 – ANCOVA Model Assuming Equal Slopes (No Covariate-by-Group Interaction)

In this example, the two treatments, supplemental lighting (SL) and partial shading (PS), are compared to a control (C) in terms of soybean yield. The comparisons are adjusted for the initial height of the plants. Each treatment is replicated 15 times in a greenhouse study.

In Example 1 we found the Age*TRT interaction to be not significant; this section presents an example of how to run an analysis on the same data assuming equal slopes.

Setup

To run this example, complete the following steps:

1 Open the ANCOVA2 example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **ANCOVA2** and click **OK**.

2 Specify the One-Way Analysis of Covariance (ANCOVA) procedure options

- Find and open the **One-Way Analysis of Covariance (ANCOVA)** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 2** 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
Group Variable.....	TRT
Covariate Variable	Height
Calc Group Means at.....	Covariate Mean
Assume Slopes are.....	Equal (No Covariate-by-Group Interaction)

Reports Tab

Compare All Pairs.....	Unchecked
Compare Each vs. Reference Group.....	Checked

3 Run the procedure

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

One-Way Analysis of Covariance (ANCOVA)


Output

Run Summary

Response Variable: Yield
 Group Variable: TRT
 Reference Group: "C"
 Covariate Variable: Height
 Slopes Assumed to be: Equal
 Model: Height + TRT

Parameter	Value	Rows	Value
R ²	0.9978	Rows Processed	45
Adjusted R ²	0.9976	Rows Filtered Out	0
Coefficient of Variation	0.0098	Rows with Response Missing	0
Mean Square Error	0.01751175	Rows with Group or Covariate Missing	0
Square Root of MSE	0.132332	Rows Used in Estimation	45
Average Percent Error	0.82	Completion Status	Normal Completion
Error Degrees of Freedom	41		

Response vs. Covariate by Group Scatter Plot



Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F-Ratio	P-Value	Significant at $\alpha = 0.05$?
Model	3	319.5545	106.5182	6082.669	0	Yes
Height	1	11.36735	11.36735	649.127	0	Yes
TRT	2	317.8324	158.9162	9074.833	0	Yes
Error	41	0.7179816	0.01751175			
Total(Adjusted)	44	320.2724	7.278919			

One-Way Analysis of Covariance (ANCOVA)

Model Coefficient T-TestsHypotheses Tested: $H_0: \beta(i) = 0$ vs. $H_1: \beta(i) \neq 0$

Independent Variable	Model Coefficient b(i)	Standard Error Sb(i)	Hypothesis Tests		
			T-Statistic	P-Value	Reject H0 at $\alpha = 0.05$?
Intercept	10.34539	0.1193907	86.652	0.0000	Yes
Height	0.05813709	0.002281857	25.478	0.0000	Yes
(TRT="PS")	-3.142156	0.05022297	-62.564	0.0000	Yes
(TRT="SL")	3.6	0.0483208	74.502	0.0000	Yes

Least Squares Means

Error Degrees of Freedom (DF): 41


Means Calculated at: Height = 52.13334 (Mean)

Name	Count	Least Squares Mean	Standard Error	95% Confidence Interval Limits for the Mean	
				Lower	Upper
Intercept					
All	45	13.52889	0.01972688	13.48905	13.56873
TRT					
C	15	13.37627	0.03447140	13.30666	13.44589
PS	15	10.23412	0.03536608	10.16269	10.30554
SL	15	16.97627	0.03447140	16.90666	17.04589

Note: These results assume that the slopes for all groups are equal (i.e., the covariate-by-group interaction is not significant). To check this assumption, run the model with unequal slopes that includes the covariate-by-group interaction and review the test of the interaction term in the Analysis of Variance report.

One-Way Analysis of Covariance (ANCOVA)

Means Plots



Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 41
 Means Calculated at: Height = 52.13334 (Mean)
 Multiple Comparison Type: Bonferroni
 Hypotheses Tested: H0: Diff = 0 vs. H1: Diff \neq 0
 Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	Hypothesis Tests			
			T-Statistic	P-Value		Reject H0 at $\alpha = 0.05$?†
				Unadjusted	Adjusted*	
TRT						
PS - C	-3.142156	0.05022297	-62.564	0.0000	0.0000	Yes
SL - C	3.6	0.0483208	74.502	0.0000	0.0000	Yes

* Adjusted p-values are computed using the number of comparisons (2) and the adjustment type (Bonferroni).

† Rejection decisions are based on adjusted p-values.

Note: These results assume that the slopes for all groups are equal (i.e., the covariate-by-group interaction is not significant). To check this assumption, run the model with unequal slopes that includes the covariate-by-group interaction and review the test of the interaction term in the Analysis of Variance report.

One-Way Analysis of Covariance (ANCOVA)

Simultaneous Confidence Intervals for Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 41

Means Calculated at: Height = 52.13334 (Mean)

Multiple Comparison Type: Bonferroni


Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	95% Simultaneous Confidence Interval Limits*	
			Lower	Upper
TRT				
PS - C	-3.142156	0.05022297	-3.259011	-3.025301
SL - C	3.6	0.0483208	3.487571	3.712429

* Confidence interval limits are adjusted based on the number of comparisons (2) and the adjustment type (Bonferroni).

Note: These results assume that the slopes for all groups are equal (i.e., the covariate-by-group interaction is not significant). To check this assumption, run the model with unequal slopes that includes the covariate-by-group interaction and review the test of the interaction term in the Analysis of Variance report.

Each vs. Reference Group Comparisons Plots

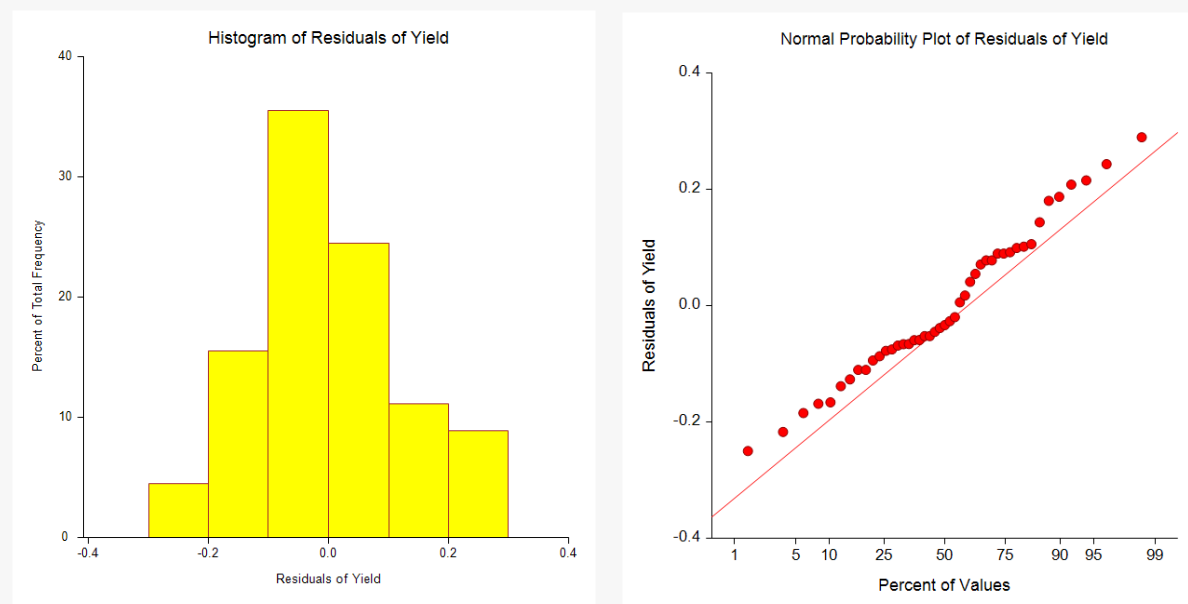


One-Way Analysis of Covariance (ANCOVA)

Residual Normality Assumption Tests

Normality Test	Test Statistic	P-Value	Reject Residual Normality at $\alpha = 0.2?$
Shapiro-Wilk	0.975	0.4405	No
Anderson-Darling	0.493	0.2167	No
D'Agostino Skewness	0.907	0.3645	No
D'Agostino Kurtosis	-0.664	0.5064	No
D'Agostino Omnibus (Skewness and Kurtosis)	1.264	0.5316	No

Residual Analysis Plots



The scatter plot shows the three regression lines with equal slopes. The lines seem to fit the data quite well and the assumptions of equal slopes and residual normality appear to be valid. Both Height and TRT are significant. The multiple comparison tests indicate that both PS and SL are different from C, with partial shade (PS) resulting in a significantly lower yield than control (C) and supplemental lighting (SL) resulting in a significantly higher yield than control (C).

Example 3 – ANCOVA Model with One-Sided Multiple Comparison Tests and Simultaneous Confidence Intervals

In this example we will analyze a dataset ANCOVA3 consisting of a treatment variable with 3 levels (C, T1, and T2), a covariate, X, and a response variable. Assume that we want to determine if T1 and/or T2 are significantly greater than C at an alpha of 0.05. We'll assume that the slopes are unequal and perform multiple comparison tests at various values of the covariate.

Setup

To run this example, complete the following steps:

1 Open the ANCOVA3 example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **ANCOVA3** and click **OK**.

2 Specify the One-Way Analysis of Covariance (ANCOVA) procedure options

- Find and open the **One-Way Analysis of Covariance (ANCOVA)** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 3** 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**
 Group Variable..... **Treatment**
 Covariate Variable **X**
 Calc Group Means at..... **45 Mean 65**

Reports Tab

Run Summary..... **Checked**
 ANOVA Table **Checked**
 Least Squares Means..... **Checked**
 Compare Each vs. Reference Group..... **Checked**
 Residual Normality Tests..... **Checked**
 Hypothesis Test and SCI Direction **One-Sided Upper**

3 Run the procedure

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

One-Way Analysis of Covariance (ANCOVA)


Output

Run Summary

Response Variable: Response
 Group Variable: Treatment
 Reference Group: "C"
 Covariate Variable: X
 Slopes Assumed to be: Unequal
 Model: $X + \text{Treatment} + X * \text{Treatment}$

Parameter	Value	Rows	Value
R ²	0.917	Rows Processed	27
Adjusted R ²	0.8972	Rows Filtered Out	0
Coefficient of Variation	0.0919	Rows with Response Missing	0
Mean Square Error	494.3325	Rows with Group or Covariate Missing	0
Square Root of MSE	22.23359	Rows Used in Estimation	27
Average Percent Error	6.849	Completion Status	Normal Completion
Error Degrees of Freedom	21		

Response vs. Covariate by Group Scatter Plot



Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F-Ratio	P-Value	Significant at $\alpha = 0.05$?
Model	5	114631.7	22926.34	46.378	0.0000	Yes
X	1	1376.559	1376.559	2.785	0.1100	No
Treatment	2	6752.797	3376.398	6.830	0.0052	Yes
X*Treatment	2	7365.478	3682.739	7.450	0.0036	Yes
Error	21	10380.98	494.3325			
Total(Adjusted)	26	125012.7	4808.18			

One-Way Analysis of Covariance (ANCOVA)

Least Squares Means

Error Degrees of Freedom (DF): 21





Means Calculated at: Covariate Values 1 to 3 (See below)

Name	Count	Least Squares Mean	Standard Error	95% Confidence Interval Limits for the Mean	
				Lower	Upper
Covariate Value 1: X = 45					
Intercept					
All	27	234.8158	8.247410	217.6644	251.9672
Treatment					
C	10	155.2916	12.523030	129.2485	181.3346
T1	7	305.0636	16.385300	270.9885	339.1387
T2	10	244.0921	13.670170	215.6635	272.5208
Covariate Value 2: X = 54.85185 (Mean)					
Intercept					
All	27	246.3570	4.355235	237.2998	255.4142
Treatment					
C	10	160.9761	7.069115	146.2751	175.6772
T1	7	286.7843	8.439433	269.2336	304.3351
T2	10	291.3107	7.036776	276.6769	305.9444
Covariate Value 3: X = 65					
Intercept					
All	27	258.2454	8.274477	241.0377	275.4531
Treatment					
C	10	166.8316	14.050920	137.6111	196.0521
T1	7	267.9553	15.404220	235.9205	299.9901
T2	10	339.9493	13.471610	311.9336	367.9651

Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

One-Way Analysis of Covariance (ANCOVA)

Means Plots



One-Way Analysis of Covariance (ANCOVA)

Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 21
Means Calculated at: Covariate Values 1 to 3 (See below)
Multiple Comparison Type: Bonferroni
Hypotheses Tested: $H_0: \text{Diff} \leq 0$ vs. $H_1: \text{Diff} > 0$
Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	Hypothesis Tests			
			T-Statistic	P-Value		Reject H0 at $\alpha = 0.05$?†
				Unadjusted	Adjusted*	
Covariate Value 1: X = 45						
Treatment						
T1 - C	149.772	20.62291	7.262	0.0000	0.0000	Yes
T2 - C	88.80054	18.53914	4.790	0.0000	0.0001	Yes
Covariate Value 2: X = 54.85185 (Mean)						
Treatment						
T1 - C	125.8082	11.00892	11.428	0.0000	0.0000	Yes
T2 - C	130.3345	9.974397	13.067	0.0000	0.0000	Yes
Covariate Value 3: X = 65						
Treatment						
T1 - C	101.1237	20.8499	4.850	0.0000	0.0001	Yes
T2 - C	173.1177	19.46568	8.893	0.0000	0.0000	Yes

* Adjusted p-values are computed using the number of comparisons (2) and the adjustment type (Bonferroni).

† Rejection decisions are based on adjusted p-values.

Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

One-Way Analysis of Covariance (ANCOVA)

Simultaneous Confidence Intervals for Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 21

Means Calculated at: Covariate Values 1 to 3 (See below)

Multiple Comparison Type: Bonferroni

Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	95% Simultaneous Confidence Interval Limits*	
			Lower	Upper

Covariate Value 1: X = 45**Treatment**

T1 - C	149.772	20.62291	106.8844	Infinity
T2 - C	88.80054	18.53914	50.24629	Infinity

Covariate Value 2: X = 54.85185 (Mean)**Treatment**

T1 - C	125.8082	11.00892	102.9139	Infinity
T2 - C	130.3345	9.974397	109.5916	Infinity

Covariate Value 3: X = 65**Treatment**

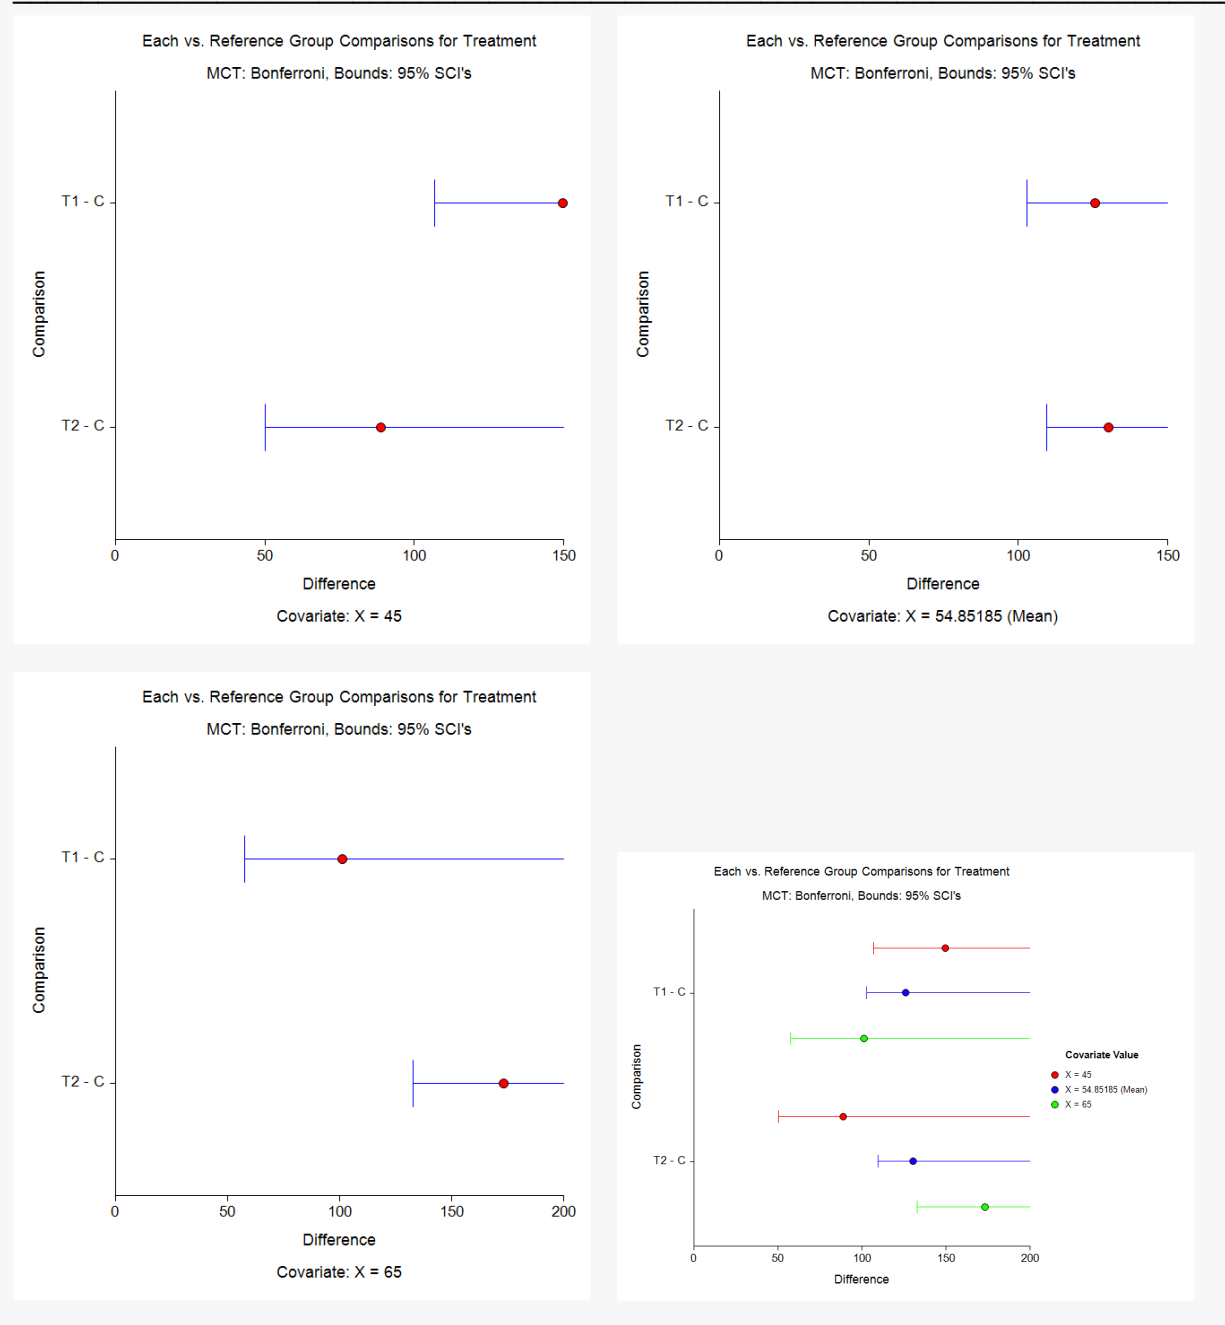
T1 - C	101.1237	20.8499	57.76394	Infinity
T2 - C	173.1177	19.46568	132.6366	Infinity

* Confidence interval limits are adjusted based on the number of comparisons (2) and the adjustment type (Bonferroni).

Note: When the model includes a significant covariate-by-group interaction, you may want to calculate and compare means at various values of the covariate and consider the results collectively. If you calculate and compare means at only one covariate value, the results may be misleading.

One-Way Analysis of Covariance (ANCOVA)

Each vs. Reference Group Comparisons Plots

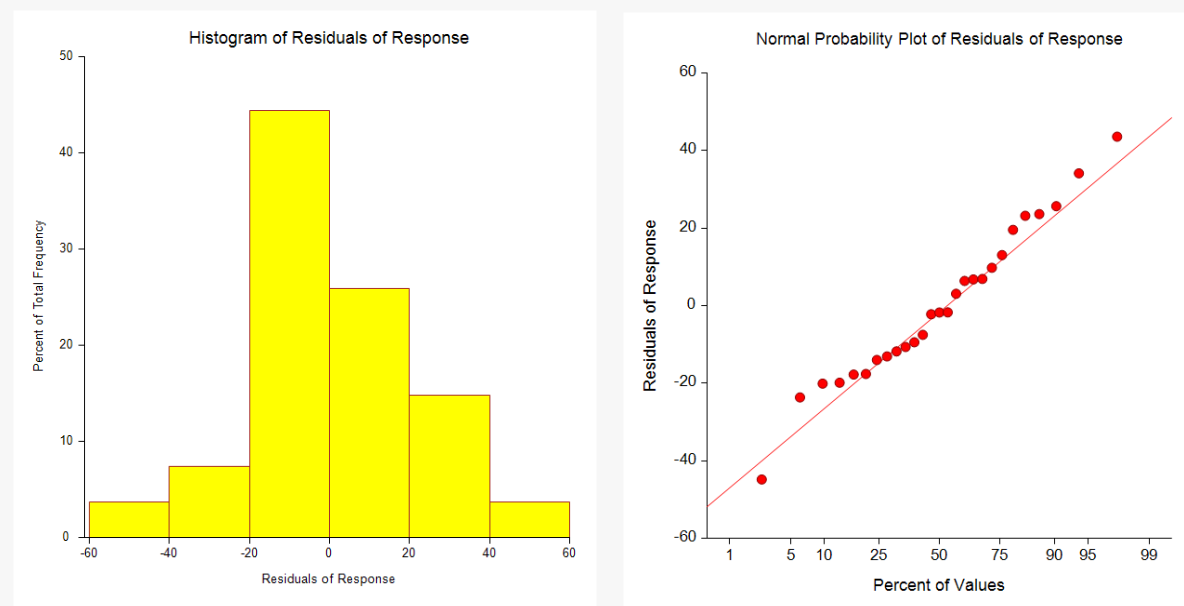


One-Way Analysis of Covariance (ANCOVA)

Residual Normality Assumption Tests

Normality Test	Test Statistic	P-Value	Reject Residual Normality at $\alpha = 0.2$?
Shapiro-Wilk	0.979	0.8487	No
Anderson-Darling	0.286	0.6264	No
D'Agostino Skewness	0.463	0.6437	No
D'Agostino Kurtosis	0.253	0.8000	No
D'Agostino Omnibus (Skewness and Kurtosis)	0.278	0.8702	No

Residual Analysis Plots



The scatter plot indicates that a model with unequal slopes is needed here. This is supported by the fact that the $X \times \text{Treatment}$ interaction is significant. The one-sided multiple comparison tests at the 3 different covariate values all indicate that the means for both T1 and T2 are significantly higher than C, regardless of the covariate value in this range. The multiple comparison plots indicate, however, that the one-sided confidence intervals for the difference are affected by the covariate value used. The assumption of residual normality is also confirmed.

Example 4 – Checking the Parallel Slopes Assumption in ANCOVA

Example 3 of the Multiple Regression procedure documentation and Example 4 of the General Linear Models (GLM) procedure documentation discuss how to check the parallel slopes assumption in ANCOVA. This example will show how to do this very quickly using this procedure.

The ANCOVA dataset contains three variables: State, Age, and IQ. The researcher wants to test for IQ differences across the three states while controlling for each subjects age. An analysis of covariance should include a preliminary test of the assumption that the slopes between age and IQ are equal across the three states. Without parallel slopes, differences among mean state IQ's depend on age.

ANCOVA Dataset

State	Age	IQ
Iowa	12	100
Iowa	13	102
.	.	.
.	.	.
.	.	.
Utah	14	104
Utah	11	105
.	.	.
.	.	.
.	.	.
Texas	15	105
Texas	16	106
.	.	.
.	.	.
.	.	.

Setup

To run this example, complete the following steps:

1 Open the ANCOVA example dataset

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

2 Specify the One-Way Analysis of Covariance (ANCOVA) procedure options

- Find and open the **One-Way Analysis of Covariance (ANCOVA)** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 4** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

One-Way Analysis of Covariance (ANCOVA)

Variables Tab

Response Variable(s) **IQ**
 Group Variable..... **State**
 Covariate Variable **Age**

Reports Tab

ANOVA Table **Checked**

Plots Tab

All Plots..... **Unchecked**

3 Run the procedure

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

Output**Analysis of Variance**

Source	DF	Sum of Squares	Mean Square	F-Ratio	P-Value	Significant at $\alpha = 0.05$?
Model	5	80.15984	16.03197	1.547	0.2128	No
Age	1	9.740934	9.740934	0.940	0.3419	No
State	2	46.57466	23.28733	2.248	0.1274	No
Age*State	2	38.72052	19.36026	1.869	0.1761	No*
Error	24	248.6402	10.36001			
Total(Adjusted)	29	328.8	11.33793			

* The covariate-by-group interaction term is not significant at $\alpha = 0.05$. You may want to run the analysis again with the assumption that the slopes are equal, which excludes the interaction.

The F-Value for the Age*State interaction term is 1.869. This matches the result that was obtained in Multiple Regression Example 3 and by hand calculations in General Linear Models (GLM) Example 4. Since the p-value of 0.1761 is not significant, we cannot reject the assumption that the three slopes are equal.

Example 5 – One-Way ANOVA Model

If you run an ANCOVA analysis and find the covariate to be non-significant, you may want to remove the covariate from the analysis and run a simple one-way ANOVA model. This example will show you how to perform a one-way ANOVA analysis using this procedure and compare the treatments to the control using Dunnett's multiple comparison test.

Note: The one-way ANOVA options are limited in this procedure. For additional options specifically related to the one-way ANOVA scenario, we suggest you use the One-Way Analysis of Variance procedure in **NCSS** instead.

Setup

To run this example, complete the following steps:

1 Open the ANCOVA2 example dataset

- From the File menu of the NCSS Data window, select **Open Example Data**.
- Select **ANCOVA2** and click **OK**.

2 Specify the One-Way Analysis of Covariance (ANCOVA) procedure options

- Find and open the **One-Way Analysis of Covariance (ANCOVA)** procedure using the menus or the Procedure Navigator.
- The settings for this example are listed below and are stored in the **Example 5** 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**
Group Variable..... **TRT**

Reports Tab

Run Summary **Checked**
ANOVA Table **Checked**
Compare Each vs. Reference Group..... **Checked**
Show Simultaneous Confidence Intervals **Checked**
(SCIs)
Multiple Comparison Test Type **Dunnett**

Plots Tab

Multiple Comparisons Plots **Checked**

3 Run the procedure

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

One-Way Analysis of Covariance (ANCOVA)

Output

Run Summary

Response Variable: Yield
 Group Variable: TRT
 Reference Group: "C"
 Covariate Variable: [None]
 Slopes Assumed to be: Unequal
 Model: TRT

Parameter	Value	Rows	Value
R ²	0.9623	Rows Processed	45
Adjusted R ²	0.9605	Rows Filtered Out	0
Coefficient of Variation	0.0396	Rows with Response Missing	0
Mean Square Error	0.287746	Rows with Group or Covariate Missing	0
Square Root of MSE	0.5364196	Rows Used in Estimation	45
Average Percent Error	3.031	Completion Status	Normal Completion
Error Degrees of Freedom	42		

Analysis of Variance

Source	DF	Sum of Squares	Mean Square	F-Ratio	P-Value	Significant at $\alpha = 0.05$?
Model	2	308.1871	154.0936	535.519	0	Yes
TRT	2	308.1871	154.0936	535.519	0	Yes
Error	42	12.08533	0.287746			
Total(Adjusted)	44	320.2724	7.278919			

Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 42
 Multiple Comparison Type: Dunnett
 Hypotheses Tested: H0: Diff = 0 vs. H1: Diff \neq 0
 Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	Hypothesis Tests			Reject H0 at $\alpha = 0.05$?†
			T-Statistic	P-Value		
				Unadjusted	Adjusted*	
TRT						
PS - C	-2.793333	0.1958728	-14.261	0.0000	0.0000	Yes
SL - C	3.6	0.1958728	18.379	0.0000	0.0000	Yes

* Adjusted p-values are computed using the number of comparisons (2) and the adjustment type (Dunnett).

† Rejection decisions are based on adjusted p-values.

One-Way Analysis of Covariance (ANCOVA)


Simultaneous Confidence Intervals for Each vs. Reference Group Comparisons of Least Squares Means

Error Degrees of Freedom (DF): 42
 Multiple Comparison Type: Dunnett
 Number of Comparisons: 2

Comparison	Least Squares Mean Difference	Standard Error	95% Simultaneous Confidence Interval Limits*	
			Lower	Upper
TRT				
PS - C	-2.793333	0.1958728	-3.241659	-2.345008
SL - C	3.6	0.1958728	3.151674	4.048326

* Confidence interval limits are adjusted based on the number of comparisons (2) and the adjustment type (Dunnett).

Each vs. Reference Group Comparisons Plots



These results are equivalent to those that would be obtained using the One-Way Analysis of Variance procedure in **NCSS**.