

Analysis of Variance

- Analysis of Variance Overview, 3-2
- One-Way Analysis of Variance, 3-5
- Two-Way Analysis of Variance, 3-11
- Analysis of Means, 3-13
- Overview of Balanced ANOVA and GLM, 3-18
- Balanced ANOVA, 3-25
- General Linear Model, 3-36
- Fully Nested ANOVA, 3-47
- Balanced MANOVA, 3-50
- General MANOVA, 3-56
- Test for Equal Variances, 3-59
- Interval Plot for Mean, 3-62
- Main Effects Plot, 3-65
- Interactions Plot, 3-67

See also

Residual Plots, Chapter 2



3-2

Analysis of Variance Overview

Analysis of Variance Overview

Analysis of variance (ANOVA) is similar to regression in that it is used to investigate and model the relationship between a response variable and one or more independent variables. However, analysis of variance differs from regression in two ways: the independent variables are qualitative (categorical), and no assumption is made about the nature of the relationship (that is, the model does not include coefficients for variables). In effect, analysis of variance extends the two-sample t-test for testing the equality of two population means to a more general null hypothesis of comparing the equality of more than two means, versus them not all being equal. Several of MINITAB's ANOVA procedures, however, allow models with both qualitative and quantitative variables.

MINITAB'S ANOVA capabilities include procedures for fitting ANOVA models to data collected from a number of different designs, for fitting MANOVA models to designs with multiple responses, for fitting ANOM (analysis of means) models, and specialty graphs for testing equal variances, for error bar or confidence interval plots, and graphs of main effects and interactions.

One-way and two-way ANOVA models

- One-way analysis of variance tests the equality of population means when classification is by one variable. The classification variable, or factor, usually has three or more levels (one-way ANOVA with two levels is equivalent to a t-test), where the *level* represents the treatment applied. For example, if you conduct an experiment where you measure durability of a product made by one of three methods, these methods constitute the levels. The one-way procedure also allows you to examine differences among means using multiple comparisons.
- Two-way analysis of variance performs an analysis of variance for testing the equality of populations means when classification of treatments is by two variables or factors. In two-way ANOVA, the data must be balanced (all cells must have the same number of observations) and factors must be fixed.

If you wish to specify certain factors to be random, use **Balanced ANOVA** if your data are balanced; use **General Linear Models** if your data are unbalanced or if you wish to compare means using multiple comparisons.

Analysis of Means

Analysis of Means (ANOM) is a graphical analog to ANOVA for the testing of the equality of population means. ANOM [15] was developed to test main effects from a designed experiment in which all factors are fixed. This procedure is used for one-way designs. MINITAB uses an extension of ANOM or Analysis of Mean treatment Effects (ANOME) [23] to test the significance of mean treatment effects for two-way designs.



Analysis of Variance

ANOM can be used if you assume that the response follows a normal distribution (similar to ANOVA) and the design is one-way or two-way. You can also use ANOM when the response follows either a binomial or Poisson distribution.

More complex ANOVA models

MINITAB offers a choice of three procedures for fitting models based upon designs more complicated than one- or two-way designs. Balanced ANOVA and General Linear Model are general procedures for fitting ANOVA models that are discussed more completely in *Overview of Balanced ANOVA and GLM* on page 3-18.

- Balanced ANOVA performs univariate (one response) analysis of variance when you have a balanced design (though one-way designs can be unbalanced). Balanced designs are ones in which all cells have the same number of observations. Factors can be crossed or nested, fixed or random. You can also use General Linear Models to analyze balanced, as well as unbalanced, designs.
- General linear model (GLM) fits the general linear model for univariate responses. In matrix form, this model is Y = XB + E, where Y is the response vector, X contains the predictors, B contains parameters to be estimated, and E represents errors assumed to be normally distributed with mean vector **0** and variance Σ. Using the general linear model, you can perform a univariate analysis of variance with balanced and unbalanced designs, analysis of covariance, and regression. GLM also allows you to examine differences among means using multiple comparisons.
- Fully nested ANOVA fits a fully nested (hierarchical) analysis of variance and estimates variance components. All factors are implicitly assumed to be random.

Testing the equality of means from multiple response

Balanced MANOVA and general MANOVA are procedures for testing the equality of vectors of means from multiple responses. Your choice between these two procedures depends upon the experimental design and the available options. Both procedures can fit MANOVA models to balanced data with up to 9 factors.

- Balanced MANOVA is used to perform multivariate analysis of variance with balanced designs. See Balanced designs on page 3-19. You can also specify factors to be random and obtain expected means squares. Use general MANOVA with unbalanced designs.
- General MANOVA is used to perform multivariate analysis of variance with either balanced or unbalanced designs that can also include covariates. You cannot specify factors to be random as you can for balanced MANOVA, although you can work around this restriction by specifying the error term for testing different model terms.



Analysis of Variance Overview

The table below summarizes the differences between Balanced and General MANOVA:

	Balanced MANOVA	General MANOVA
Can fit unbalanced data	no	yes
Can specify factors as random and obtain expected means squares	yes	no
Can fit covariates	no	yes
Can fit restricted and unrestricted forms of a mixed model	yes	no; unrestricted only

Special analytical graphs

- Test for equal variances performs Bartlett's (or F-test if 2 levels) and Levene's hypothesis tests for testing the equality or homogeneity of variance. Many statistical procedures, including ANOVA, are based upon the assumption that samples from different populations have the same variance.
- Interval plot for mean creates a plot of means with either error bars or confidence intervals when you have a one-way design.
- Main effects plot creates a main effects plot for either raw response data or fitted values from a model-fitting procedure. The points in the plot are the means at the various levels of each factor with a reference line drawn at the grand mean of the response data. Use the main effects plot to compare magnitudes of marginal means.
- Interactions plot creates a single interaction plot if two factors are entered, or a matrix of interaction plots if 3 to 9 factors are entered. An interactions plot is a plot of means for each level of a factor with the level of a second factor held constant. Interactions plots are useful for judging the presence of interaction, which means that the difference in the response at two levels of one factor depends upon the level of another factor. Parallel lines in an interactions plot indicate no interaction. The greater the departure of the lines from being parallel, the higher the degree of interaction. To use an interactions plot, data must be available from all combinations of levels.

Use the main effects plot and the interactions plot in Chapter 19 to generate main effects plots and interaction plots specifically for 2-level factorial designs, such as those generated by Create Factorial Design and Create RS Design.



One-Way Analysis of Variance

One-Way Analysis of Variance

One-way analysis of variance tests the equality of population means when classification is by one variable. There are two ways to organize your data in the worksheet. You can enter the response in one column (stacked) or in different columns (unstacked). If your response is in one column, you can examine differences among means using multiple comparisons.

Data

The response variable must be numeric. You can enter the sample data from each population into separate columns of your worksheet (*unstacked* case), or you can stack the response data in one column with another column of level values identifying the population (*stacked* case). In the stacked case, the factor level column can be numeric, text, or date/time. If you wish to change the order in which text levels are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories in the Manipulating Data chapter of MINITAB User's Guide 1. You do not need to have the same number of observations in each level. You can use Calc > Make Patterned Data to enter repeated factor levels. See the Generating Patterned Data chapter in MINITAB User's Guide 1.

To perform a one-way analysis of variance with stacked data

- Choose Stat > ANOVA > One-way.

- 2 In Response, enter the column containing the responses.
- 3 In Factor, enter the column containing the factor levels.
- 4 If you like, use one or more of the options described below, then click OK.



Analysis of Variance



One-Way Analysis of Variance

To perform a one-way analysis of variance with unstacked data

1 Choose Stat > ANOVA > One-way (Unstacked).

One-way Analysis of Variance			
	Response	es (in separate col	umns):
	I		<u>_</u>
			-
	,		_
· · · · · · · · · · · · · · · · · · ·			
Select			G <u>r</u> aphs
Help		<u>0</u> K	Cancel

- 2 In **Responses (in separate columns)**, enter the columns containing the separate response variables.
- 3 If you like, use one or more of the options described below, then click OK.

Options with stacked data

One-way dialog box

store residuals and fitted values (the means for each level).

Comparisons subdialog box

 display confidence intervals for the differences between means, using four different multiple comparison methods: Fisher's LSD, Tukey's, Dunnett's, and Hsu's MCB (multiple comparisons with the best). See *Multiple comparisons of means* on page 3-7.

Graphs subdialog box

- draw boxplots, dotplots, and residual plots. You can draw five different residual plots:
 - histogram.
 - normal probability plot.
 - plot of residuals versus the fitted values (Y).
 - plot of residuals versus data order. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.
 - separate plot for the residuals versus each specified column.

For a discussion of the residual plots, see Residual plots on page 2-6.



One-Way Analysis of Variance

Options with unstacked data

Graphs subdialog box

draw boxplots and dotplots that display the sample mean for each sample.

Multiple comparisons of means

Multiple comparisons of means allow you to examine which means are different and to estimate by how much they are different. When you have a single factor and your data are stacked, you can obtain multiple comparisons of means by choosing the Stat > ANOVA ➤ One-way and then clicking the Comparisons subdialog box.

The choice of method

The multiple comparison methods compare different means and use different error rates. If you wish to examine all pairwise comparisons of means, use either Fisher's least significant difference (LSD) or Tukey's (also called Tukey-Kramer in the unbalanced case) method. The choice depends on whether you wish to control the individual (comparison-wise) error rate or the family (experiment-wise) error rate. The danger in using the individual error rate with Fisher's method is having an unexpectedly high probability of making at least one Type I error (declaring a difference when there is none) among all the comparisons. MINITAB displays both error rates. In most cases, the Tukey method is probably the choice that you should make when you want to judge all pairwise differences, because you can control the family error rate.

Choose the Dunnett method if you are comparing treatments to a control. When this method is suitable, it is inefficient to use the Tukey all-pairwise approach, because the Tukey confidence intervals will be wider and the hypothesis tests less powerful for a given family error rate. You will need to specify which level represents the control. If this level is text or date/time, enclose it with double quotes.

Choose Hsu's MCB (multiple comparison with the best) method if it makes sense to compare each mean only with the "best" among all of the other ones. This procedure allows you to judge how much worse a level might be if it is not the best or how much better it might be than its closest competitor. You will need to specify if the "best" is smallest or largest. If you are mainly interested in comparing each level to the "best" it is inefficient to use the Tukey all-pairwise approach because you will waste your error rate comparing pairs of level means which do not include the best mean.

MINITAB User's Guide 2



Analysis of Variance

One-Way Analysis of Variance



Comparison method	Purpose	Error rate
Fisher's LSD	all pairwise differences	individual
Tukey	all pairwise differences	family
Dunnett	comparison to a control	family
Hsu's MCB	comparison with the best	family

Interpreting confidence intervals

MINITAB presents results in confidence interval form to allow you to assess the practical significance of differences among means, in addition to statistical significance. As usual, the null hypothesis of no difference between means is rejected if and only if zero is not contained in the confidence interval.

Specify error rates as percents between 0.1 and 50%. The default error rate of 5% is the family error rate for the Tukey, Dunnett, and MCB methods and the individual error rate for the Fisher method. Individual error rates are exact in all cases, meaning that they can be calculated by an explicit formula. Family error rates are exact for equal group sizes. If group sizes are unequal, the true family error rate for the Tukey, Fisher, and MCB methods will be slightly smaller than stated, resulting in conservative confidence intervals [4], [21]. The Dunnett family error rates are exact for unequal sample sizes.

The F-test and multiple comparisons

The results of the F-test and multiple comparisons can conflict. For example, it is possible for the F-test to reject the null hypothesis of no differences among the level means, and yet all the Tukey pairwise confidence intervals may contain zero. Conversely, it is possible for the F-test to fail to reject the null hypothesis, and yet have one or more of the Tukey pairwise confidence intervals not include zero. The F-test has been used to protect against the occurrence of false positive differences in means. However, the Tukey, Dunnett, and MCB methods have protection against false positives built in, while the Fisher method only benefits from this protection when all means are equal. If the use of multiple comparisons is conditioned upon the significance of the F-test, the error rate can be higher than the error rate in the unconditioned application of multiple comparisons [14].

See Help for computational details of the multiple comparison methods.



One-Way Analysis of Variance

Analysis of Variance

Example of a one-way ANOVA with multiple comparisons

You design an experiment to assess the durability of four experimental carpet products. You place a sample of each of the carpet products in four homes and you measure durability after 60 days. Because you wish to test the equality of means and to assess the differences in means, you use the one-way ANOVA procedure (data in stacked form) with multiple comparisons. Generally, you would choose one multiple comparison method as appropriate for your data. However, two methods are selected here to demonstrate MINITAB's capabilities.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > One-way.
- 3 In Response, enter Durability. In Factor, enter Carpet.
- 4 Click **Comparisons**. Check **Tukey's, family error rate** and enter *10* in the text box. Check **Hsu's MCB, family error rate** and enter *10* in the text box. Click **OK** in each dialog box.





One-Way Analysis of Variance

Tukey's pairwise comparisons

I

C

Chapter 3

Family ndividual	error rate error rate	= 0.100 = 0.0250		
ritical v	alue = 3.62			
ntervals	for (column	level mean)	- (row level m	ıe
	1	2	3	
2	-2.106 11.601			
3	-5.178 8.528	-9.926 3.781		
4	-9.376 4.331	-14.123 -0.417	-11.051 2.656	

Interpreting the results

The default one-way output contains an analysis of variance table, a table of level means, individual 95% confidence intervals, and the pooled standard deviation. The F-test p-value of 0.101 indicates that there is not quite sufficient evidence (at $\alpha = 0.10$ or less) to claim that not all the means are equal. However, you should examine the multiple comparison results, which use family error rates of 0.10, because the methods used (Tukey, MCB) have built in protection against false positive results.

an)

The output labeled "Hsu's MCB" compares each mean with the best of the other means. Here, "best" is the default or largest of the others. The means of carpets 1, 2, and 3 were compared to the level 4 mean because the carpet 4 mean is the largest of the rest. The level 4 mean was compared to the carpet 1 mean. Carpets 1, 3, or 4 may be best, since the corresponding confidence intervals contain positive values. There is no evidence that carpet 2 is the best because the upper interval endpoint is 0, the smallest it can be.

In addition, it is possible to describe the potential advantage or disadvantage of any of the contenders for the best. For example, if carpet 3 is best, it is no more than 0.809 better than its closest competitor, and it may be as much as 9.204 worse than the best of the other level means. If carpet 1 is not the best, it is no more than 2.484 worse than the best of the other means, and it may be as much as 7.529 better than the best of the others.

The first pair of numbers in the Tukey output table, (-2.106, 11.601), gives the confidence interval for the mean of carpet 1 minus the mean of carpet 2. Confidence intervals for entries not in the table can be found from entries in the table. For example, the confidence interval for the mean of level 2 minus the mean of carpet 1 is (-11.601, 2.106). Carpets 2 and 4 are the only ones for which the means can be declared as different, since the confidence interval for this combination of means is the only one that excludes zero.



Two-Way Analysis of Variance

Analysis of Variance

By not conditioning upon the F-test, differences in treatment means appear to have occurred at family error rates of 0.10. If the MCB method is a good choice for these data, carpet 2 might be eliminated as a choice for the "best". By the Tukey method, the mean durability from carpets 2 and 4 appears to be different.

Two-Way Analysis of Variance

A two-way analysis of variance tests the equality of populations means when classification of treatments is by two variables or factors. For this procedure, the data must be balanced (all cells must have the same number of observations) and factors must be fixed.

If you wish to specify certain factors to be random, use **Balanced ANOVA** if your data are balanced, and use **General Linear Model** if your data are unbalanced or if you wish to compare means using multiple comparisons.

Data

The response variable must be numeric and in one worksheet column. You must have a single factor level column for each of the two factors. These can be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories in the Manipulating Data chapter of MINITAB User's Guide 1. You must have a balanced design (same number of observations in each treatment combination) with fixed and crossed factors. See Balanced designs on page 3-19, Fixed vs. random factors on page 3-19, and Crossed vs. nested factors on page 3-19. You can use Calc ➤ Make Patterned Data to enter repeated factor levels. See the Generating Patterned Data chapter in MINITAB User's Guide 1.

To perform a two-way analysis of variance

1 Choose Stat > ANOVA > Two-way.

Two-way Analysis of Varian	ce		×
	Response:		
	Ro <u>w</u> factor:		🗆 Display means
	<u>C</u> olumn factor:		🗖 Display <u>m</u> eans
	☐ Store r <u>e</u> sidua ☐ Store fits	ls	
I	□ Fit <u>a</u> dditive m	odel	
Select			G <u>r</u> aphs
Help		<u>0</u> K	Cancel

MINITAB User's Guide 2

3-11

Two-Way Analysis of Variance

- 2 In **Response**, enter the column containing the response.
- 3 In **Row Factor**, enter one of the factor level columns.
- 4 In Column Factor, enter the other factor level column.
- 5 If you like, use one or more of the options described below, then click OK.

Options

Two-way dialog box

- print sample means and 95% confidence intervals for factor levels means.
- store residuals and fits.
- fit an additive model, that is, a model without the interaction term. In this case, the fitted value for cell (i, j) is (mean of observations in row i) + (mean of observations in row j) (mean of all observations).

Graphs subdialog box

- draw five different residual plots. You can display the following plots:
 - histogram.
 - normal probability plot.
 - plot of residuals versus the fitted values (Y).
 - plot of residuals versus data order. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.
 - separate plot for the residuals versus each specified column.

For a discussion of the residual plots, see Residual plots on page 2-6.

Example of two-way analysis of variance

You are a biologist who is studying how zooplankton live in two lakes. You set up twelve tanks in your laboratory, six each with water from a different lake. You add one of three nutrient supplements to each tank and after 30 days you count the zooplankton in a unit volume of water. You use two-way ANOVA to test if the population means are equal, or equivalently, to test whether there is significant evidence of interactions and main effects.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Two-way.
- 3 In Response, enter Zooplankton.
- 4 In Row factor, enter Supplement. Check Display means.
- 5 In Column factor, enter Lake. Check Display means. Click OK.

D

Analysis of Variance

Analysis of Means

Session window output

Two-way ANOVA: Zooplankton versus Supplement, Lake						
Analysis of Source Suppleme	Variance DF 2	for Zoop SS 1919	lank MS 959	F 9.25	P 0.015	
Lake Interaction Error	1 2 6	21 561 622	21 281 104	0.21 2.71	0.666 0.145	
Total	11	3123 Indivio	dual 95% C	I		
Suppleme 1 2 3	Mean 43.5 68.3 39.8	+ ((* *) (+) +	
		30.0 Individ	45.0 dual 95% C	60.0	75.0	
Lake Dennison Rose	Mean 51.8 49.2	((42.0	+ +) 48.	+*- * +	+)) -++- 4.0 60.0	-)

Interpreting the results

The default output for two-way ANOVA is the analysis of variance table. For the zooplankton data, there is no significant evidence for a supplement*lake interaction effect or a lake main effect if your acceptable α value is less than 0.145 (the p-value for the interaction F-test). There is significant evidence for supplement main effects, as the F-test p-value is 0.015.

As requested, the means are displayed with individual 95% confidence intervals. Supplement 2 appears to have provided superior plankton growth in this experiment. These are t-distribution confidence intervals calculated using the error degrees of freedom and the pooled standard deviation (square root of the mean square error). If you want to examine simultaneous differences among means using multiple comparisons, use General Linear Model (page 3-36).

Analysis of Means

Analysis of Means (ANOM), a graphical analog to ANOVA, tests the equality of population means. ANOM [15] was developed to test main effects from a designed experiment in which all factors are fixed. This procedure is used for one-factor designs. MINITAB uses an extension of ANOM or ANalysis Of Mean treatment Effects (ANOME) [23] to test the significance of mean treatment effects for two-factor designs.



An ANOM chart can be described in two ways: by its appearance and by its function. In appearance, it resembles a Shewhart control chart. In function, it is similar to ANOVA for detecting differences in population means [12]. There are some important differences between ANOM and ANOVA, however. The hypotheses they test are not identical [16]. ANOVA tests whether the treatment means are different from each other; ANOM tests whether the treatment means differ from the grand mean.

For most cases, ANOVA and ANOM will likely give similar results. However, there are some scenarios where the two methods might be expected to differ: 1) if one group of means is above the grand mean and another group of means is below the grand mean, ANOVA's F-test might indicate evidence for differences where ANOM might not; 2) if the mean of one group is separated from the other means, the ANOVA F-test might not indicate evidence for differences whereas ANOM might flag this group as being different from the grand mean. Refer to [20], [21], [22], and [23] for an introduction to the analysis of means.

ANOM can be used if you assume that the response follows a normal distribution, similar to ANOVA, and the design is one-way or two-way. You can also use ANOM when the response follows either a binomial distribution or a Poisson distribution.

Data

Response data from a normal distribution

Your response data must be numeric and entered into one column. Factor columns may be numeric, text, or date/time and may contain any values. MINITAB's capability to enter patterned data can be helpful in entering numeric factor levels. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories in the Manipulating Data chapter of MINITAB User's Guide 1. You can use Calc > Make Patterned Data to enter repeated factor levels. See the Generating Patterned Data chapter in MINITAB User's Guide 1.

One-way designs may be balanced or unbalanced and can have up to 100 levels. Two-way designs must be balanced and can have up to 50 levels for each factor. All factors must be fixed. See *Fixed vs. random factors* on page 3-19.

Rows with missing data are automatically omitted from calculations. If you have two factors, the design must be balanced after omitting rows with missing values.

Response data from a binomial distribution

The response data are the numbers of defectives (or defects) found in each sample, with a maximum of 500 samples. These data must be entered into one column.

Since the decision limits in the ANOM chart are based upon the normal distribution, one of the assumptions that must be met when the response data are binomial is that the sample size must be large enough to ensure that the normal approximation to the

MINITAB User's Guide 2

3-14

Analysis of Means



Analysis of Variance

binomial is valid. A general rule of thumb is to only use ANOM if np > 5 and n(1 - p) > 5, where n is the sample size and p is the proportion of defectives. The second assumption is that all of the samples are the same size. See [23] for more details.

A sample with a missing response value (*) is automatically omitted from the analysis.

Response data from a Poisson distribution

The response data are the numbers of defects found in each sample. You can have up to 500 samples.

The Poisson distribution can be adequately approximated by a normal distribution if the mean of the Poisson distribution is at least 5. Hence, when the Poisson mean is large enough, analysis of means can be applied to data from a Poisson distribution to test if the population means are equal to the grand mean.

A sample with a missing response value (*) is automatically omitted from the analysis.

To perform an analysis of means

1 Choose Stat > ANOVA > Analysis of Means.

Analysis of Means	×
	Response:
	Distribution of Data
	• Normal
	Factor 1:
	Factor <u>2</u> : (Optional)
	C <u>B</u> inomial
	Sample size:
	O <u>P</u> oisson
	Alpha level: 0.05
	☐ Include a summary table
·	
Select	<u>T</u> itle:
Help	<u>O</u> K Cancel

- 2 In **Response**, enter a numeric column containing the response variable.
- 3 Under Distribution of Data, choose Normal, Binomial, or Poisson.
 - If you choose Normal, do one of the following:
 - for a one-way design, enter the column containing the factor levels in Factor 1
 - for a two-way design, enter the columns containing the factor levels in Factor 1 and Factor 2
 - If you choose Binomial, enter a number in Sample size.
- 4 If you like, use one or more of the options described below, then click OK.

Options

- change the experiment wide error rate, or alpha level (default is 0.05). This will change the location of the decision lines on the graph.
- print a summary table of level statistics for normal (prints means, standard errors, sample size) or binomial data (prints number, proportion of defectives).
- replace the default graph title with your own title.

Example of a two-way analysis of means (ANOM)

You perform an experiment to assess the effect of three process time levels and three strength levels on density. You use analysis of means for normal data and a two-way design to identify any significant interactions or main effects.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Analysis of Means.
- 3 In **Response**, enter *Density*.
- 4 Choose Normal.
- 5 In Factor 1, enter Minutes. In Factor 2, enter Strength. Click OK.

Graph window output



Interpreting the results

Three plots are displayed in one graph with a two-way ANOM: one showing the interaction effects, one showing the means for the first factor, and one showing the means for the second factor. Control charts have a center line and control limits. If a point falls outside the control limits, then there is significant evidence that the mean represented by that point is different from the grand mean. With a two-way ANOM, look at the interaction effects first. If there is significant evidence for interaction, it usually does not make sense to consider main effects, because the effect of one factor depends upon the level of the other.

MINITAB User's Guide 2



Analysis of Means

Analysis of Means

In our example, the interaction effects are well within the control limits, signifying no evidence of interaction. Now you can look at the main effects. The lower two plots show the means for the levels of the two factors, with the main effect being the difference between the mean and the center line. The point representing the level 3 mean of the factor Minutes is displayed by a red asterisk, which indicates that there is significant evidence for the level 3 mean being different from the grand mean at $\alpha = 0.05$. You may wish to investigate any point near or above the control limits. The main effects for levels 1 and 3 of factor Strength are well outside the control limits of the lower left plot, signifying that there is evidence for these means being different from the grand mean at $\alpha = 0.05$.

Example of an ANOM for binomial response data

You count the number of rejected welds from samples of size 80 in order to identify samples whose proportions of rejects are out of line with the other samples. Because the data are binomial (two possible outcomes, constant proportion of success, and independent samples) you use analysis of means for binomial data.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Analysis of Means.
- 3 In Response, enter WeldRejects.
- 4 Choose **Binomial** and enter *80* in **Sample size**. Click **OK**. One-Way Binomial Analysis of Means for WeldRejects

Graph window output



Interpreting the results

A single plot displays the proportion of defects, a center line representing the average proportion, and upper and lower decision limits. A similar plot is displayed for one-way normal data or for Poisson data. As with the two-way ANOM plot, you can judge if there is significant evidence for a sample mean being different from the average if the point representing that sample falls outside the control limits.

MINITAB User's Guide 2

3-17



D

Chapter 3

3-18

Overview of Balanced ANOVA and GLM

In this example, the proportion of defective welds in sample four is identified as being unusually high because the point representing this sample falls outside the control limits.

Overview of Balanced ANOVA and GLM

Balanced ANOVA and general linear model (GLM) are ANOVA procedures for analyzing data collected with many different experimental designs. Your choice between these procedures depends upon the experimental design and the available options. The experimental design refers to the selection of units or subjects to measure, the assignment of treatments to these units or subjects, and the sequence of measurements taken on the units or subjects. Both procedures can fit univariate models to balanced data with up to 9 factors. Here are some of the other options:

	Balanced ANOVA	GLM
Can fit unbalanced data	no	yes
Can specify factors as random and obtain expected means squares	yes	yes
Fits covariates	no	yes
Performs multiple comparisons	no	yes
Fits restricted/unrestricted forms of mixed model	yes	unrestricted only

You can use balanced ANOVA to analyze data from balanced designs—see *Balanced designs* on page 3-19. You can use GLM to analyze data from any balanced design, though you cannot choose to fit the restricted case of the mixed model, which only balanced ANOVA can fit—see *Restricted and unrestricted form of mixed models* on page 3-27.

To determine how to classify your variables, see *Crossed vs. nested factors* on page 3-19, *Fixed vs. random factors* on page 3-19, and *Covariates* on page 3-20.

For information on how to specify the model, see *Specifying the model terms* on page 3-20, *Specifying terms involving covariates* on page 3-21, *Specifying reduced models* on page 3-21, and *Specifying models for some specialized designs* on page 3-22.

For easy entering of repeated factor levels into your worksheet, see *Using patterned data* to set up factor levels on page 3-24.

 $\mathbf{\Sigma}$

Overview of Balanced ANOVA and GLM

Balanced designs

Your design must be balanced to use balanced ANOVA, with the exception of a one-way design. A *balanced* design is one with equal numbers of observations at each combination of your treatment levels. A quick test to see whether or not you have a balanced design is to use **Stat > Tables > Cross Tabulation**. Enter your classification variables and see if you have equal numbers of observations in each cell, indicating balanced data.

Crossed vs. nested factors

When two or more factors are present in a design, they may be *crossed* or *nested*, depending upon how the levels of one factor appear with the levels of the other factor. This concept can be demonstrated by the following example. Suppose that there are two factors: plant and operator. If observations are made with each operator at each plant, then these are *crossed* factors. If observations are made with different operators at each plant, then operator is *nested* within plant. In general, if each level of factor A occurs with each level of factor B, factors A and B are *crossed*. If each level of factor B occurs within only one level of factor A, then factor B is *nested* within factor A. The designation of whether a factor is *crossed* or nested within MINITAB occurs with the specification of the model. See *Specifying the model terms* on page 3-20. It is important make the correct designation in order to obtain the correct error term for factors.

Fixed vs. random factors

MEET MTB

In addition to the crossed or nested designation for pairs of factors, a factor can be either *fixed* or *random*. Designating a factor as fixed or random depends upon how you view that factor in a larger context. Suppose that one factor is machine operator. If one is truly interested in each operator and may, for example, initiate a training procedure for specific operators depending upon the tests results, then the operator factor is *fixed*. If the operators are considered to be drawn at random from a population of all operators and interest is in the population and not the individuals, then the factor is *random*. In MINITAB, factors are assumed to be fixed unless specified otherwise. It is important to make the correct designation in order to obtain the correct error term for factors.

The terms *fixed* and *random* often modify the word *effect*. What is usually meant by *effect* for a fixed factor is the difference between the mean corresponding to a factor level and the overall mean. The testing of fixed effects being zero is equivalent to the testing of treatment means being equal. An effect for a random factor is not defined as the difference in means because the interest is in estimation and testing of variance components. An effect is *mixed* if it is the interaction effect of fixed and random factors. With balanced ANOVA, you can choose whether to restrict the sum of mixed effects. See *Restricted and unrestricted form of mixed models* on page 3-27.

MINITAB User's Guide 2

3-19

UGUIDE 1 UGUIDE 2 SC QREF





Analysis of Variance

D

Overview of Balanced ANOVA and GLM

Chapter 3

Covariates

A *covariate* is a quantitative variable included in an ANOVA model. A covariate may be a variable for which the level is not controlled as part of the design, but has been measured and it is entered into the model to reduce the error variance. A covariate may also be a quantitative variable for which the levels have been controlled as part of the experiment. Regardless of the origin, the statistical model contains a coefficient for the covariate was a predictor in a regression model.

Specifying the model terms

You must specify the model terms in the **Model** box. This is an abbreviated form of the statistical model that you may see in textbooks. Because you enter the response variable(s) in **Responses**, in **Model** you enter only the variables or products of variables that correspond to terms in the statistical model. MINITAB uses a simplified version of a statistical model as it appears in many textbooks. Here are some examples of statistical models and the terms to enter in **Model**. A, B, and C represent factors.

Case	Statistical model	Terms in model
Factors A, B crossed	$y_{ijk} = \mu + a_i + b_j + ab_{ij} + e_{k(ij)}$	A B A*B
Factors A, B, C crossed	$y_{ijkl} = \mu + a_i + b_j + c_k + ab_{ij} + ac_{ik} + bc_{jk} + abc_{ijk} + e_{l(ijk)}$	A B C A*B A*C B*C A*B*C
3 factors nested (B within A, C within A and B)	$y_{ijkl} = \mu + a_i + b_{j(i)} + c_{k(ij)} + e_{l(ijk)}$	A B(A) C(AB)
Crossed and nested (B nested within A, both crossed with C)	$y_{ijkl} = \mu + a_i + b_{j(i)} + c_k + ac_{ik} + bc_{jk(i)} + e_{l(ijk)}$	A B(A) C A*C B*C

In MINITAB's models you omit the subscripts, μ , e, and +'s that appear in textbook models. An * is used for an interaction term and parentheses are used for nesting. For example, when B is nested within A, you enter B (A), and when C is nested within both A and B, you enter C (A B). Enter B(A) C(B) for the case of 3 sequentially nested factors. Terms in parentheses are always factors in the model and are listed with blanks between them. Thus, D*F (A B E) is correct but D*F (A*B E) and D (A*B*C) are not. Also, one set of parentheses cannot be used inside another set. Thus, C (A B) is correct but C (A B (A)) is not. An interaction term between a nested factor and the factor it is nested within is invalid.

See *Specifying terms involving covariates* on page 3-21 for details on specifying models with covariates.

Several special rules apply to naming columns. You may omit the quotes around variable names. Because of this, variable names must start with a letter and contain only letters and numbers. Alternatively, you can use C notation (C1, C2, etc.) to denote data



Overview of Balanced ANOVA and GLM

Analysis of Variance

columns. You can use special symbols in a variable name, but then you must enclose the name in single quotes.

You can specify multiple responses. In this case, a separate analysis of variance will be performed for each response.

Specifying terms involving covariates

You can specify variables to be covariates in GLM. You must specify the covariates in **Covariates**, but you can enter the covariates in **Model**, though this is not necessary unless you cross or nest the covariates (see below table).

In an unbalanced design or a design involving covariates, GLM's sequential sums of squares (the additional model sums of squares explained by a variable) will depend upon the order in which variables enter the model. If you do not enter the covariates in **Model** when using GLM, they will be fit first, which is what you usually want when a covariate contributes background variability. The subsequent order of fitting is the order of terms in **Model**. The sequential sums of squares for unbalanced terms A B will be different depending upon the order that you enter them in the model. The default adjusted sums of squares (sums of squares with all other terms in the model), however, will be the same, regardless of model order.

GLM allows terms containing covariates crossed with each other and with factors, and covariates nested within factors. Here are some examples of these models, where A is a factor.

Case	Covariates	Terms in model
test homogeneity of slopes (covariate crossed with factor)	Х	A X A*X
same as previous	Х	A X
quadratic in covariate (covariate crossed with itself)	Х	A X X*X
full quadratic in two covariates (covariates crossed)	ХZ	A X Z X*X Z*Z X*Z
separate slopes for each level of A (covariate pested within a factor)	х	A X(A)

Specifying reduced models

You can fit *reduced* models. For example, suppose you have a three factor design, with factors, A, B, and C. The *full* model would include all one factor terms: A, B, C, all two-factor interactions: A*B, A*C, B*C, and the three-factor interaction: A*B*C. It becomes a reduced model by omitting terms. You might reduce a model if terms are

MINITAB User's Guide 2

3-21

CONTENTS INDEX MEET MTB UGUIDE 1 UGUIDE 2 SC QREF HOW TO US

(D)

Chapter 3

3-22

Overview of Balanced ANOVA and GLM

not significant or if you need additional error degrees of freedom and you can assume that certain terms are zero. For this example, the model with terms A B C A*B is a reduced three-factor model.

One rule about specifying reduced models is that they must be hierarchical. That is, for a term to be in the model, all lower order terms contained in it must also be in the model. For example, suppose there is a model with four factors: A, B, C, and D. If the term A*B*C is in the model, then the terms A B C A*B A*C B*C must also be in the model, though any terms involving D do not have to be in the model. The hierarchical structure applies to nesting as well. If B(A) is in the model, then A must be also.

Because models can be quite long and tedious to type, two shortcuts have been provided. A vertical bar indicates crossed factors, and a minus sign removes terms.

Long form	Short form
A B C A*B A*C B*C A*B*C	A B C
A B C A*B A*C B*C	A B C – A*B*C
A B C B*C E	A B C E
A B C D A*B A*C A*D B*C B*D C*D A*B*D A*C*D B*C*D	A B C D - A*B*C - A*B*C*D
A B(A) C A*C B*C	A B(A) C

In general, all crossings are done for factors separated by bars unless the cross results in an illegal term. For example, in the last example, the potential term A*B(A) is illegal and MINITAB automatically omits it. If a factor is nested, you must indicate this when using the vertical bar, as in the last example with the term B(A).

Specifying models for some specialized designs

Some experimental designs can effectively provide information when measurements are difficult or expensive to make or can minimize the effect of unwanted variability on treatment inference. The following is a brief discussion of three commonly used designs that will show you how to specify the model terms in MINITAB. To illustrate these designs, two treatment factors (A and B) and their interaction (A*B) are considered. These designs are not restricted to two factors, however. If your design is balanced, you can use balanced ANOVA to analyze your data. Otherwise, use GLM.

Randomized block design

A *randomized block* design is a commonly used design for minimizing the effect of variability when it is associated with discrete units (e.g. location, operator, plant, batch, time). The usual case is to randomize one replication of each treatment combination within each block. There is usually no intrinsic interest in the blocks and these are considered to be random factors. The usual assumption is that the block by treatment

Overview of Balanced ANOVA and GLM

Analysis of Variance

interaction is zero and this interaction becomes the error term for testing treatment effects. If you name the block variable as Block, enter *Block A B A***B* in **Model** and enter *Block* in **Random Factors**.

Split-plot design

A *split-plot* design is another blocking design, which you can use if you have two or more factors. You might use this design when it is more difficult to randomize one of the factors compared to the other(s). For example, in an agricultural experiment with the factors variety and harvest date, it may be easier to plant each variety in contiguous rows and to randomly assign the harvest dates to smaller sections of the rows. The block, which can be replicated, is termed the *main plot* and within these the smaller plots (variety strips in example) are called *subplots*.

This design is frequently used in industry when it is difficult to randomize the settings on machines. For example, suppose that factors are temperature and material amount, but it is difficult to change the temperature setting. If the blocking factor is operator, observations will be made at different temperatures with each operator, but the temperature setting is held constant until the experiment is run for all material amounts. In this example, the plots under operator constitute the main plots and temperatures constitute the subplots.

There is no single error term for testing all factor effects in a split-plot design. If the levels of factor A form the subplots, then the mean square for Block*A will be the error term for testing factor A. There are two schools of thought for what should be the error term to use for testing B and A*B. If you enter the term Block*B, the expected mean squares show that the mean square for Block*B is the proper term for testing factor B and that the remaining error (which is Block*A*B) will be used for testing A*B. However, it is often assumed that the Block*B and Block*A*B interactions do not exist and these are then lumped together into error [6]. You might also pool the two terms if the mean square for Block*B is small relative to Block*A*B. In you don't pool, enter *Block A Block*A B A*B* in **Model** and what is labeled as Error is really Block*A*B. If you do pool terms, enter *Block A Block*A B A*B* in **Model** and what is labeled as Error is the set of pooled terms. In both cases enter *Block* in **Random Factors**.

Latin square with repeated measures design

A *repeated measures* design is a design where repeated measurements are made on the same subject. There are a number of ways in which treatments can be assigned to subjects. With living subjects especially, systematic differences (due to learning, acclimation, resistance, etc.) between successive observations may be suspected. One common way to assign treatments to subjects is to use a Latin square design. An advantage of this design for a repeated measures experiment is that it ensures a balanced fraction of a complete factorial (i.e. all treatment combinations represented) when subjects are limited and the sequence effect of treatment can be considered to be negligible.

Ð

Chapter 3

3 - 24

Overview of Balanced ANOVA and GLM

A *Latin square* design is a blocking design with two orthogonal blocking variables. In an agricultural experiment there might be perpendicular gradients that might lead you to choose this design. For a repeated measures experiment, one blocking variable is the group of subjects and the other is time. If the treatment factor B has three levels, b_1 , b_2 , and b_3 , then one of twelve possible Latin square randomizations of the levels of B to subjects groups over time is:

	Time 1	Time 2	Time 3
Group 1	b ₂	b ₃	b ₁
Group 2	b ₃	b ₁	b ₂
Group 3	b ₁	b ₂	b ₃

The subjects receive the treatment levels in the order specified across the row. In this example, group 1 subjects would receive the treatments levels in order b_2 , b_3 , b_1 . The interval between administering treatments should be chosen to minimize carryover effect of the previous treatment.

This design is commonly modified to provide information on one or more additional factors. If each group was assigned a different level of factor A, then information on the A and A*B effects could be made available with minimal effort if an assumption about the sequence effect given to the groups can be made. If the sequence effects are negligible compared to the effects of factor A, then the group effect could be attributed to factor A. If interactions with time are negligible, then partial information on the A*B interaction may be obtained [27]. In the language of repeated measures designs, factor A is called a *between-subjects* factor and factor B a *within-subjects* factor.

Let's consider how to enter the model terms into MINITAB. If the group or A factor, subject, and time variables were named A, Subject, and Time, respectively, enter A *Subject(A) Time B A**B in **Model** and enter *Subject* in **Random Factors**.

It is not necessary to randomize a repeated measures experiments according to a Latin square design. See *Example of a repeated measures design* on page 3-30 for a repeated measures experiment where the fixed factors are arranged in a complete factorial design.

Using patterned data to set up factor levels

MINITAB's set patterned data capability can be helpful when entering numeric factor levels. For example, to enter the level values for a three-way crossed design with a, b, and c (a, b, and c represent numbers) levels of factors A, B, C, and n observations per cell, fill out the **Calc > Make Patterned Data > Simple Set of Numbers** dialog box and execute 3 times, once for each factor, as shown: (See the *Generating Patterned Data* chapter in *MINITAB User's Guide 1*.)



Balanced ANOVA

	Factor					
From first value	1	1	1			
From last value	а	b	с			
List each value	bcn	cn	n			
List the whole sequence	1	а	ab			

Analysis of Variance

Balanced ANOVA

Use Balanced ANOVA to perform univariate analysis of variance for each response variable.

Your design must be balanced, with the exception of one-way designs. *Balanced* means that all treatment combinations (cells) must have the same number of observations. See *Balanced designs* on page 3-19. Use General Linear Model (page 3-36) to analyze balanced and unbalanced designs.

Factors may be crossed or nested, fixed or random. See *Crossed vs. nested factors* on page 3-19 and *Fixed vs. random factors* on page 3-19. You may include up to 50 response variables and up to 9 factors at one time.

Data

You need one column for each response variable and one column for each factor, with each row representing an observation. Regardless of whether factors are crossed or nested, use the same form for the data. Factor columns may be numeric, text, or date/ time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories in the Manipulating Data chapter in MINITAB User's Guide 1.

Balanced data are required except for one-way designs. The requirement for balanced data extends to nested factors as well. Suppose A has 3 levels, and B is nested within A. If B has 4 levels within the first level of A, B must have 4 levels within the second and third levels of A. MINITAB will tell you if you have unbalanced nesting. In addition, the subscripts used to indicate the 4 levels of B within each level of A must be the same. Thus, the four levels of B cannot be (1 2 3 4) in level 1 of A, (5 6 7 8) in level 2 of A, and (9 10 11 12) in level 3 of A. However, you can use GLM to analyze data coded in this way.

If any response or factor column specified contains missing data, that entire observation (row) is excluded from all computations. The requirement that data be balanced must be preserved after missing data are omitted. If an observation is missing for one response variable, that row is eliminated for all responses. If you want to eliminate

UGUIDE 2

CONTENTS INDEX

MEET MTB

HOW TO USE

Balanced ANOVA

missing rows separately for each response, perform balanced ANOVA separately for each response.

To perform a balanced ANOVA

1 Choose Stat > ANOVA > Balanced ANOVA.

Balanced Analysis of Variar	x X
	R <u>e</u> sponses:
	Mo <u>d</u> el:
	<u>v</u>
	Random factors:
	×
	Ogtions
Select	G <u>r</u> aphs Res <u>u</u> lts <u>S</u> torage
Help	<u>O</u> K Cancel

- 2 In **Responses**, enter up to 50 numeric columns containing the response variables.
- 3 In Model, type the model terms you want to fit. See *Specifying the model terms* on page 3-20.
- 4 If you like, use one or more of the options described below, then click OK.

Options

Balanced Analysis of Variance dialog box

specify which factors are random factors—see Fixed vs. random factors on page 3-19.

Options subdialog box

use the restricted form of the mixed models (both fixed and random effects). The
restricted model forces mixed interaction effects to sum to zero over the fixed effects.
By default, MINITAB fits the unrestricted model. See *Restricted and unrestricted form of
mixed models* on page 3-27.

Graphs subdialog box

- draw five different residual plots. You can display the following plots:
 - histogram.
 - normal probability plot.
 - plot of residuals versus the fitted values (Y).



Balanced ANOVA

Analysis of Variance

- plot of residuals versus data order. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.
- separate plot for the residuals versus each specified column.

For a discussion of the residual plots, see *Residual plots* on page 2-6.

Results subdialog box

- display expected means squares, estimated variance components, and error terms used in each F-test. See *Expected mean squares* on page 3-27.
- display a table of means corresponding to specified terms from the model. For example, if you specify A B D A*B*D, four tables of means will be printed, one for each main effect, A, B, D, and one for the three-way interaction, A*B*D.

Storage subdialog box

store the fits and residuals separately for each response. If you fit a full model, fits are cell means. If you fit a reduced model, fits are least squares estimates. See Specifying reduced models on page 3-21.

Restricted and unrestricted form of mixed models

A mixed model is one with both fixed and random factors. There are two forms of this model: one requires the crossed, mixed terms to sum to zero over subscripts corresponding to fixed effects (this is called the restricted model), and the other does not. See *Example of both restricted and unrestricted forms of the mixed model* on page 3-32. Many textbooks use the restricted model. Most statistics programs (e.g., SAS, JMP, and SPSS) use the unrestricted model. MINITAB fits the unrestricted model by default, but you can choose to fit the restricted form. The reasons to choose one form over the other have not been clearly defined in the statistical literature. Searle et al. [24] say "that question really has no definitive, universally acceptable answer," but also say that one "can decide which is more appropriate to the data at hand," without giving guidance on how to do so.

Your choice of model form does not affect the sums of squares, degrees of freedom, mean squares, or marginal and cell means. It does affect the expected mean squares, error terms for F-tests, and the estimated variance components. See *Example of both restricted and unrestricted forms of the mixed model* on page 3-32.

Expected mean squares

If you do not specify any factors to be random, MINITAB will assume that they are fixed. In this case, the denominator for F-statistics will be the MSE. However, for models which include random terms, the MSE is not always the correct error term. You can examine the expected means squares to determine the error term that was used in the F-test.

3 - 28

Balanced ANOVA

When you select **Display expected mean squares and variance components** in the Results subdialog box, MINITAB will print a table of expected mean squares, estimated variance components, and the error term (the denominator mean squares) used in each F-test. The *expected mean squares* are the expected values of these terms with the specified model. If there is no exact F-test for a term, MINITAB solves for the appropriate error term in order to construct an approximate F-test. This test is called a *synthesized test*.

The estimates of variance components are the usual unbiased analysis of variance estimates. They are obtained by setting each calculated mean square equal to its expected mean square, which gives a system of linear equations in the unknown variance components that is then solved. Unfortunately, this method can result in negative estimates, which should be set to zero. MINITAB, however, prints the negative estimates because they sometimes indicate that the model being fit is inappropriate for the data. Variance components are not estimated for fixed terms.

Example of ANOVA with two crossed factors

An experiment was conducted to test how long it takes to use a new and an older model of calculator. Six engineers each work on both a statistical problem and an engineering problem using each calculator model; the time in minutes to solve the problem is recorded. The engineers can be considered as blocks in the experimental design. There are two factors—type of problem, and calculator model—each with two levels. Because each level of one factor occurs in combination with each level of the other factor, these factors are crossed. The example and data are from Neter, Wasserman, and Kutner [18], page 936.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Balanced ANOVA.
- 3 In Responses, enter SolveTime.
- 4 In Model, type Engineer ProbType | Calculator.
- 5 In Random Factors, enter Engineer.
- 6 Click **Results**. In **Display means corresponding to the terms**, type *ProbType* | *Calculator*. Click **OK** in each dialog box.



Analysis of Variance

Balanced ANOVA

Session window	ANOVA: SolveTime versus Engineer, ProbType, Calculator	
output	Factor Type Levels Values Engineer random 6 Adams Dixon Erickson Jones Mayr Williams	ies
	ProbType fixed 2 Eng Stat Calculat fixed 2 New Old	
	Analysis of Variance for SolveTim	
	Source DF SS MS F P Engineer 5 1.053 0.211 3.13 0.039 ProbType 1 16.667 16.667 247.52 0.000 Calculat 1 72.107 72.107 1070.89 0.000 ProbType*Calculat 1 3.682 3.682 54.68 0.000 Error 15 1.010 0.067 1071 1071 1071 Total 23 94.518 94.518 94.518 94.518 94.518	
	Means	
	ProbType N SolveTim Eng 12 3.8250 Stat 12 5.4917	
	Calculat N SolveTim New 12 2.9250 Old 12 6.3917	
	ProbType Calculat N SolveTim Eng New 6 2.4833 Eng Old 6 5.1667 Stat New 6 3.3667 Stat Old 6 7.6167	

Interpreting the results

MINITAB displays a list of factors, with their type (fixed or random), number of levels, and values. Next displayed is the analysis of variance table. The analysis of variance indicates that there is a significant calculator by problem type interaction, which implies that the decrease in mean compilation time in switching from the old to the new calculator depends upon the problem type.

Because you requested means for all factors and their combinations, the means of each factor level and factor level combinations are also displayed. These show that the mean compilation time decreased in switching from the old to new calculator type.

MINITAB User's Guide 2

3-29



CONTENTS INDEX MEET MTB UGUIDE 1 UGUIDE 2 SC QREF HOW TO USE

Balanced ANOVA

Chapter 3

3-30

Example of a repeated measures design

The following example contains data from Winer [27], p. 546, to illustrate a complex repeated measures model. An experiment was run to see how several factors affect subject accuracy in adjusting dials. Three subjects perform tests conducted at one of two noise levels. At each of three time periods, the subjects monitored three different dials and made adjustments as needed. The response is an accuracy score. The noise, time, and dial factors are crossed, fixed factors. Subject is a random factor, nested within noise. Noise is a between-subjects factor, time (variable ETime) and dial are within-subjects factors.

The model terms are entered in a certain order so that the error terms used for the fixed factors are just below the terms for whose effects they test. (With a single random factor, the interaction of a fixed factor with the random factor becomes the error term for that fixed effect.) Because Subject was specified as Subject(Noise) the first time, you do not need to repeat "(Noise)" in the interactions involving Subject. The interaction ETime*Dial*Subject, the error term for ETime*Dial, is not entered in the model because there would be zero degrees of freedom left over for error. By not entering ETime*Dial*Subject in the model, it is labeled as Error and you have the error term that is needed.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Balanced ANOVA.
- 3 In Responses, enter Score.
- 4 In Model, enter Noise Subject(Noise) ETime Noise*ETime ETime*Subject Dial Noise*Dial Dial*Subject ETime*Dial Noise*ETime*Dial.
- 5 In Random Factors, enter Subject.
- 6 Click Options.
- 7 Check Use the restricted form of the model. Click OK.
- 8 Click Results.
- 9 Check Display expected mean squares and variance components. Click OK in each dialog box.

D

Analysis of Variance

Balanced ANOVA

Session window output

ANOVA: Score versus Noise, ETime, Dial, Subject												
	Factor Noise Subject(Noise) ETime Dial	Type L fixed random fixed fixed	evels. 2 3 3 3	Value 1 1 1 1	S	2 2 2 2	3 3 3					
	Analysis of Variance for Score											
	Source Noise Subject(Noise) ETime Noise*ETime ETime*Subject(No Dial Noise*Dial Dial*Subject(No ETime*Dial Noise*ETime*Dia Error Total	Noise) Dise) NI	DF 1 4 2 8 2 8 2 8 4 4 4 16 53	468 2491 3722 333 234 2370 50 105 105 11 127 9924	SS .17 .11 .33 .00 .89 .33 .33 .56 .67 .33 .11 .83		MS 468.17 622.78 1861.17 166.50 29.36 1185.17 25.17 13.19 2.67 2.83 7.94	F 0.75 78.39 63.39 5.67 3.70 89.82 1.91 1.66 0.34 0.36	P 0.435 0.000 0.029 0.013 0.000 0.210 0.184 0.850 0.836			
	Source 1 Noise 2 Subject(Nois 3 ETime 4 Noise*ETime 5 ETime*Subject 6 Dial 7 Noise*Dial 8 Dial*Subject 9 ETime*Dial 10 Noise*ETime* 11 Error	se) ct(Noise) c(Noise) *Dial	Varia compor 68. 7. 1.	.315 .139 .750	11 5 11 8 11 11 11 11	Ex (us (1 (1 (1 (1 (1 (1 (1 (1 (1 (1)	pected Me ing restr 1) + 9(2) 1) + 3(5) 1) + 3(5) 1) + 3(5) 1) + 3(5) 1) + 3(8) 1) + 3(8) 1) + 3(8) 1) + 3(8) 1) + 6Q[9 1) + 3Q[1 1)	an Squa icted m + 27Q[+ 18Q[+ 9Q[4 + 18Q[+ 9Q[7] 0]	re for odel) 1] 3]] 6]]	Each	Term	

Interpreting the results

MINITAB displays the table of factor levels, the analysis of variance table, and the expected mean squares. Important information to gain from the expected means squares are the estimated variance components and discovering which error term is used for testing the different model terms.

The term labeled Error is in row 11 of the expected mean squares table. The column labeled "Error term" indicates that term 11 was used to test terms 2, 5, and 8 to 10. Dial*Subject is numbered 8 and was used to test the sixth and seventh terms. You can follow the pattern for other terms.

You can gain some idea about how the design affected the sensitivity of F-tests by viewing the variance components. The variance components used in testing within-subjects factors are smaller (7.139, 1.750, 7.944) than the between-subjects

MINITAB User's Guide 2



E 1 UGUIDE 2 SC QREF HOW TO USE



3-32



variance (68.315). It is typical that a repeated measures model can detect smaller differences in means within subjects as compared to between subjects.

Of the four interactions among fixed factors, the noise by time interaction was the only one with a low p-value (0.029). This implies that there is significant evidence for judging that a subjects' sensitivity to noise changed over time. Because this interaction is significant, at least at $\alpha = 0.05$, the noise and time main effects are not examined. There is also significant evidence for a dial effect (p-value < 0.0005). Among random terms, there is significant evidence for time by subject (p-value = 0.013) and subject (p-value < 0.0005) effects.

Example of both restricted and unrestricted forms of the mixed model

A company ran an experiment to see how several conditions affect the thickness of a coating substance that it manufactures. The experiment was run at two different times, in the morning and in the afternoon. Three operators were chosen from a large pool of operators employed by the company. The manufacturing process was run at three settings, 35, 44, and 52. Two determinations of thickness were made by each operator at each time and setting. Thus, the three factors are crossed. One factor, operator, is random; the other two, time and setting, are fixed.

The statistical model is

 $Y_{ijkl} = \mu + T_i + O_j + S_k + TO_{ij} + TS_{ik} + OS_{ik} + TOS_{ijk} + e_{ijkl},$

where T_i is the time effect, O_j is the operator effect, and S_k is the setting effect, and TO_{ii} , TS_{ik} , OS_{ik} , and TOS_{iik} are the interaction effects.

Operator, all interactions with operator, and error are random. The random terms are:

O_j TO_{ij} OS_{jk} TOS_{ijk} e_{ijkl}

These terms are all assumed to be normally distributed random variables with mean zero and variances given by

$var\left(O_{j}\right) = V(O)$	var $(TO_{ij}) = V(TO)$	$var(OS_{jk}) = V(OS)$
$var \ (TOS_{jkl}) = V(TOS)$	var $(e_{iikl}) = V(e) = \sigma^2$	

These variances are called variance components. The output from expected means squares contains estimates of these variances.

In the unrestricted model, all these random variables are independent. The remaining terms in this model are fixed.

Balanced ANOVA

Analysis of Variance

In the restricted model, any term which contains one or more subscripts corresponding to fixed factors is required to sum to zero over each fixed subscript. In the example, this means

$$\begin{split} \sum_{i} (T_i) &= 0 \qquad \sum_{k} (S_k) &= 0 \qquad \sum_{i} (TO_{ij}) &= 0 \\ \sum_{k} (TS_{ik}) &= 0 \qquad \sum_{k} (OS_{jk}) &= 0 \qquad \sum_{i} (TOS_{ijk}) &= 0 \end{split}$$

Your choice of model does not affect the sums of squares, degrees of freedom, mean squares, or marginal and cell means. However, it does affect the expected mean squares, error term for the F-tests, and the estimated variance components.

Step 1: Fit the restricted form of the model

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Balanced ANOVA.
- 3 In **Responses**, enter *Thickness*.
- 4 In Model, type Time | Operator | Setting.
- 5 In Random Factors, enter Operator.
- 6 Click Options. Check Use the restricted form of the model. Click OK.
- 7 Click Results. Check Display expected mean squares and variance components.
- 8 Click **OK** in each dialog box.

Step 2: Fit the unrestricted form of the model

1 Repeat steps 1-8 above except that in 6), uncheck Use the restricted form of the model.

Balanced ANOVA

Output for restricted case

Session	ANOVA: Thickness versus	Time, C	perator	; Setting			
viridow	Eactor Type Loyals V	aluos					
output	Time fixed 2	1	2				
	Operator random 3	1	2	3			
	Setting fixed 3	35	44	52			
	Analysis of Variance for	Thickn	es				
	Source	DF	SS	MS	F	Р	
	Time	1	9.0	9.0	0.29	0.644	
	Operator	2	1120.9	560.4	165.38	0.000	
	Setting	2	15676.4	7838.2	73.18	0.001	
	Time*Operator	2	62.0	31.0	9.15	0.002	
	Time*Setting	2	114.5	57.3	2.39	0.208	
	Operator*Setting	4	428.4	107.1	31.61	0.000	
	Time*Operator*Setting	4	96.0	24.0	7.08	0.001	
	Error	18	61.0	3.4			
	Total	35	17568.2				
	Source	Varian	ce Erro	r Expected M	ean Squa	re for	Each Term
		compone	nt term	(using rest	ricted m	odel)	
	1 Time		4	(8) + 6(4)	+ 18Q[1]	
	2 Operator	46.4	21 8	(8) + 12(2))		
	3 Setting		6	(8) + 4(6)	+ 12Q[3]	
	4 Time*Operator	4.6	02 8	(8) + 6(4)			
	5 Time*Setting		7	(8) + 2(7)	+ 6Q[5]		
	6 Operator*Setting	25.9	31 8	(8) + 4(6)			
	7 Time*Operator*Setting	10.3	06 8	(8) + 2(7)			
	8 Error	3.3	89	(8)			

MINITAB User's Guide 2

 \geq



Balanced ANOVA

Analysis of Variance

Output for unrestricted case

Session	ANOVA: Thickness versus	Time,	Operato	or, Sett	ing			
window output	Factor Type Levels V Time fixed 2 Operator random 3 Setting fixed 3	alues 1 1 35	2 2 44	3 52				
	Analysis of Variance for	Thick	nes					
	Source Time Operator Setting Time*Operator Time*Setting Operator*Setting Time*Operator*Setting Error Total x Not an exact F-test.	DF 1 2 2 2 2 4 4 18 35	S: 9.0 1120.9 15676.4 62.0 114.9 428.4 96.0 61.0 17568.2	S O 9 4 0 5 4 0 0 2	MS 9.0 560.4 7838.2 31.0 57.3 107.1 24.0 3.4	F 0.29 4.91 73.18 1.29 2.39 4.46 7.08	P 0.644 0.090 x 0.001 0.369 0.208 0.088 0.001	
	Source 1 Time 2 Operator 3 Setting 4 Time*Operator 5 Time*Setting 6 Operator*Setting 7 Time*Operator*Setting 8 Error	Varia compon 37. 1. 20. 10. 3.	nce Erro ent terr 194 * 167 7 778 7 306 8 389	or Exp n (usi (8) (8) (8) (8) (8) (8) (8) (8) (8)	ected Me ng unres + 2(7) + 2(7) + 2(7) + 2(7) + 2(7) + 2(7) + 2(7)	an Squa tricted + 6(4) + 4(6) + 4(6) + 6(4) + Q[5] + 4(6)	re for Ea model) + Q[1,5] + 6(4) + + Q[3,5]	ch Term 12(2)
	* Synthesized Test. Error Terms for Synthesi	zed Te	sts					
	Source 2 Operator	Erro	r DF E 3.73	rror M 114.	S Synth 1 (4)	esis of + (6) -	Error MS (7)	

Interpreting the results

The organization of the output is the same for restricted and unrestricted models: a table of factor levels, the analysis of variance table, and as requested, the expected mean squares. The differences in the output are in the expected means squares and the F-tests for some model terms. In this example, the F-test for Operator is synthesized for the unrestricted model because it could not be calculated exactly.

Examine the 3 factor interaction, Time*Operator*Setting. The F-test is the same for both forms of the mixed model, giving a p-value of 0.001. This implies that the coating thickness depends upon the combination of time, operator, and setting. Many analysts



General Linear Model

would go no further than this test. If an interaction is significant, any lower order interactions and main effects involving terms of the significant interaction are not considered meaningful.

Let's examine where these models give different output. The Operator*Setting F-test is different, because the error terms are Error in the restricted case and Time*Operator*Setting in the unrestricted case, giving p-values of < 0.0005 and 0.088, respectively. Likewise, the Time*Operator differs for the same reason, giving p-values of 0.002 and 0.369, for the restricted and unrestricted cases, respectively. The estimated variance components for Operator, Time*Operator, and Operator*Setting also differ.

General Linear Model

Use General Linear Model (GLM) to perform univariate analysis of variance with balanced and unbalanced designs, analysis of covariance, and regression, for each response variable.

Calculations are done using a regression approach. A "full rank" design matrix is formed from the factors and covariates and each response variable is regressed on the columns of the design matrix.

Factors may be crossed or nested, fixed or random. Covariates may be crossed with each other or with factors, or nested within factors. You can analyze up to 50 response variables with up to 9 factors and 50 covariates at one time.

Data

Set up your worksheet in the same manner as with balanced ANOVA: one column for each response variable, one column for each factor, and one column for each covariate, so that there is one row for each observation. The factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See Ordering Text Categories in the Manipulating Data chapter in MINITAB User's Guide 1.

Although models can be unbalanced in GLM, they must be "full rank," that is, there must be enough data to estimate all the terms in your model. For example, suppose you have a two-factor crossed model with one empty cell. Then you can fit the model with terms A B, but not A B A*B. MINITAB will tell you if your model is not full rank. In most cases, eliminating some of the high order interactions in your model (assuming, of course, they are not important) can solve this problem.

Nesting does not need to be balanced. A nested factor must have at least 2 levels at some level of the nesting factor. If factor B is nested within factor A, there can be unequal levels of B within each level of A. In addition, the subscripts used to identify the


Ø

General Linear Model

Analysis of Variance

B levels can differ within each level of A. This means, for example, that the B levels can be (1 2 3 4) in level 1 of A, (5 6 7 8) in level 2 of A, and (9 10 11 12) in level 3 of A.

If any response, factor, or covariate column contains missing data, that entire observation (row) is excluded from all computations. If you want to eliminate missing rows separately for each response, perform GLM separately for each response.

To perform an analysis using general linear model

1 Choose Stat > ANOVA > General Linear Model.

General Linear Model	×
	R <u>e</u> sponses:
	Mo <u>d</u> el:
	<u>À</u>
	v
	Random <u>f</u> actors:
	×
	Councilator Ontions Comparisons
	<u>Gr</u> aphs <u>Storage</u>
Select	Factor Plots
Help	<u>QK</u> Cancel

- 2 In Responses, enter up to 50 numeric columns containing the response variables.
- **3** In **Model**, type the model terms you want to fit. See *Specifying the model terms* on page 3-20.
- 4 If you like, use one or more of the options described below, then click OK.

Options

General Linear Model dialog box

specify which factors are random factors—see *Fixed vs. random factors* on page 3-19.

Covariates subdialog box

include up to 50 covariates in the model.

Options subdialog box

- enter a column containing weights to perform weighted regression—see Weighted regression on page 2-6.
- select adjusted (Type III) or sequential (Type I) sums of squares for calculations. See Adjusted vs. sequential sums of squares on page 3-42.

General Linear Model

Comparisons subdialog box

- perform multiple comparison of treatment means with the mean of a control level.
 You can also choose
 - pairwise comparisons or comparisons with a control
 - the term(s) you wish to compare the means
 - from among three methods (Tukey, Bonferroni, and/or Sidak for pairwise comparisons or Dunnett, Bonferroni, and/or Sidak for comparisons with a control)
 - the alternative (less than, equal to, greater than) when you choose comparisons with a control. *Equal to* is the default.
 - whether to display the comparisons by confidence intervals or hypothesis tests (both are given by default). You can specify the family confidence level of intervals (the default is 95%).

For a discussion of multiple comparisons, see *Multiple comparisons of means* on page 3-39.

Graphs subdialog box

- draw five different residual plots for regular, standardized, or deleted residuals—see Choosing a residual type on page 2-5. Available residual plots include a
 - histogram.
 - normal probability plot.
 - plot of residuals versus the fitted values (Y).
 - plot of residuals versus data order. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n.
 - separate plot for the residuals versus each specified column.

For a discussion of the residual plots, see *Residual plots* on page 2-6.

Results subdialog box

- display the following in the Session window:
 - no output.
 - the table of factor levels and the analysis of variance table.
 - the default output, which includes the above plus estimated coefficients for covariate terms and a table of unusual observations
 - the above plus estimated coefficients for all terms
- display expected means squares, estimated variance components, and error terms used in each F-test—see *Expected mean squares* on page 3-27.
- display the adjusted or least squares means (fitted values) corresponding to specified terms from the model.



Analysis of Variance

Storage subdialog box

- store coefficients for the model, in separate columns for each response.
- store fits and regular, standardized, and deleted residuals separately for each response—see Choosing a residual type on page 2-5.
- store leverages, Cook's distances, and DFITS, for identifying outliers—see Identifying outliers on page 2-9.
- store the design matrix. The design matrix multiplied by the coefficients will yield the fitted values. See Design matrix used by General Linear Model on page 3-42.

Factor Plots subdialog box

- enter factors to construct a main effects plot—see Main Effects Plot on page 3-65.
- enter factors to construct an interactions plot—see Interactions Plot on page 3-67.

Multiple comparisons of means

Multiple comparisons of means allow you to examine which means are different and to estimate by how much they are different. When you have multiple factors, you can obtain multiple comparisons of means through GLM's Comparisons subdialog box.

There are some common pitfalls to the use of multiple comparisons. If you have a quantitative factor you should probably examine linear and higher order effects rather than performing multiple comparisons (see [12] and *Example of using GLM to fit linear and quadratic effects* on page 3-43). In addition, performing multiple comparisons for those factors which appear to have the greatest effect or only those with a significant F-test can result in erroneous conclusions (see *Which means to compare?* on page 3-40).

You have the following choices when using multiple comparisons:

- pairwise comparisons or comparisons with a control
- which means to compare
- the method of comparison
- display comparisons in confidence interval or hypothesis test form
- the confidence level, if you choose to display confidence intervals
- the alternative, if you choose comparisons with a control

Following are some guidelines for making these choices.

Pairwise comparisons or comparison with a control

Choose **Pairwise Comparisons** when you do not have a control level but you would like to examine which pairs of means are different.





General Linear Model

Choose **Comparisons with a Control** when you are comparing treatments to a control. When this method is suitable, it is inefficient to use the all-pairwise approach, because the all-pairwise confidence intervals will be wider and the hypothesis tests less powerful for a given family error rate. If you do not specify a level that represents the control, MINITAB will assume that the lowest level of the factors is the control. If you wish to change which level is the control, specify a level that represents the control for each term that you are comparing the means of. If these levels are text or date/time, enclose each with double quotes.

Which means to compare?

Choosing which means to compare is an important consideration when using multiple comparisons; a poor choice can result in confidence levels that are not what you think. Issues that should be considered when making this choice might include: 1) should you compare the means for only those terms with a significant F-test or for those sets of means for which differences appear to be large? 2) how deep into the design should you compare means—only within each factor, within each combination of first-level interactions, or across combinations of higher level interactions?

It is probably a good idea to decide which means you will compare before collecting your data. If you compare only those means with differences that appear to be large, which is called data snooping, then you are increasing the likelihood that the results suggest a real difference where no difference exists [9], [18]. Similarly, if you condition the application of multiple comparisons upon achieving a significant F-test, then the error rate of the multiple comparisons can be higher than the error rate in the unconditioned application of multiple comparisons [9], [14]. The multiple comparison methods have protection against false positives already built in.

In practice, however, many people commonly use F-tests to guide the choice of which means to compare. The ANOVA F-tests and multiple comparisons are not entirely separate assessments. For example, if the p-value of an F-test is 0.9, you probably will not find statistically significant differences among means by multiple comparisons.

How deep within the design should you compare means? There is a trade-off: if you compare means at all two-factor combinations and higher orders turn out to be significant, then the means that you compare might be a mix of effects; if you compare means at too deep a level, you lose power because the sample sizes become smaller and the number of comparisons become larger. In practice, you might decide to compare means for factor level combinations for which you believe the interactions are meaningful.

MINITAB restricts the terms that you can compare means for to fixed terms or interactions among fixed terms. Nesting is considered to be a form of interaction.

To specify which means to compare, enter terms from the model in the **Terms** box. If you have 2 factors named A and B, entering **A B** will result in multiple comparisons within each factor. Entering A*B will result in multiple comparisons for all level combination of factors A and B. You can use the notation A|B to indicate interaction for pairwise comparisons but not for comparisons with a control.



Analysis of Variance

The multiple comparison method

You can choose from among three methods for both pairwise comparisons and comparisons with a control. Each method provides simultaneous or joint confidence intervals, meaning that the confidence level applies to the set of intervals computed by each method and not to each one individual interval. By protecting against false positives with multiple comparisons, the intervals are wider than if there were no protection.

The Tukey (also called Tukey-Kramer in the unbalanced case) and Dunnett methods are extensions of the methods used by one-way ANOVA. The Tukey approximation has been proven to be conservative when comparing three means. "Conservative" means that the true error rate is less than the stated one. In comparing larger numbers of means, there is no proof that the Tukey method is conservative for the general linear model. The Dunnett method uses a factor analytic method to approximate the probabilities of the comparisons. Because it uses the factor analytic approximation, the Dunnett method is not generally conservative. The Bonferroni and Sidak methods are conservative methods based upon probability inequalities. The Sidak method is slightly less conservative than the Bonferroni method.

Comparison method	Properties
Dunnett	comparison to a control only, not proven to be conservative
Tukey	all pairwise differences only, not proven to be conservative
Bonferroni	most conservative
Sidak	conservative, but slightly less so than Bonferroni

Some characteristics of the multiple comparison methods are summarized below:

Display of comparisons in confidence interval or hypothesis test form

MINITAB presents multiple comparison results in confidence interval and/or hypothesis test form. Both are given by default.

When viewing confidence intervals, you can assess the practical significance of differences among means, in addition to statistical significance. As usual, the null hypothesis of no difference between means is rejected if and only if zero is not contained in the confidence interval. When you request confidence intervals, you can specify family confidence levels for the confidence intervals. The default level is 95%.

MINITAB calculates adjusted p-values for hypothesis test statistics. The adjusted p-value for a particular hypothesis within a collection of hypotheses is the smallest family wise α level at which the particular hypothesis would be rejected.

See Help for computational details of the multiple comparison methods.

S

Chapter 3

General Linear Model

Adjusted vs. sequential sums of squares

MINITAB by default uses adjusted (Type III) sums of squares for all GLM calculations. Adjusted sums of squares are the additional sums of squares determined by adding each particular term to the model given the other terms are already in the model. You also have the choice of using sequential (Type I) sums of squares in all GLM calculations. Sequential sums of squares are the sums of squares added by a term with only the previous terms entered in the model. These sums of squares can differ when your design is unbalanced or if you have covariates. Usually, you would probably use adjusted sums of squares. However, there may be cases where you might want to use sequential sums of squares.

Design matrix used by General Linear Model

General Linear Model uses a regression approach to fit the model that you specify. First MINITAB creates a design matrix, from the factors and covariates, and the model that you specify. The columns of this matrix are the predictors for the regression.

The design matrix has n rows, where n = number of observations, and one block of columns, often called dummy variables, for each term in the model. There are as many columns in a block as there are degrees of freedom for the term. The first block is for the constant and contains just one column, a column of all ones. The block for a covariate also contains just one column, the covariate column itself.

Suppose A is a factor with 4 levels. Then it has 3 degrees of freedom and its block contains 3 columns, call them A1, A2, A3. Each row is coded as one of the following:

level of A	A1	A2	A3
1	1	0	0
2	0	1	0
3	0	0	1
4	-1	-1	-1



Analysis of Variance

Suppose factor B has 3 levels nested within each level of A. Then its block contains $(3 - 1) \times 4 = 8$ columns, call them B11, B12, B21, B22, B31, B32, B41, B42, coded as follows:

level of A	level of B	B11	B12	B21	B22	B31	B32	B41	B42
1	1	1	0	0	0	0	0	0	0
1	2	0	1	0	0	0	0	0	0
1	3	-1	-1	0	0	0	0	0	0
2	1	0	0	1	0	0	0	0	0
2	2	0	0	0	1	0	0	0	0
2	3	0	0	-1	-1	0	0	0	0
3	1	0	0	0	0	1	0	0	0
3	2	0	0	0	0	0	1	0	0
3	3	0	0	0	0	-1	-1	0	0
4	1	0	0	0	0	0	0	1	0
4	2	0	0	0	0	0	0	0	1
4	3	0	0	0	0	0	0	-1	-1

To calculate the dummy variables for an interaction term, just multiply all the corresponding dummy variables for the factors and/or covariates in the interaction. For example, suppose factor A has 6 levels, C has 3 levels, D has 4 levels, and Z and W are covariates. Then the term A*C*D*Z*W*W has $5 \times 2 \times 3 \times 1 \times 1 \times 1 = 30$ dummy variables. To obtain them, multiply each dummy variable for A by each for C, by each for D, by the covariates Z once and W twice.

Example of using GLM to fit linear and quadratic effects

An experiment is conducted to test the effect of temperature and glass type upon the light output of an oscilloscope. There are three glass types and three temperature levels: 100, 125, and 150 degrees Fahrenheit. These factors are fixed because we are interested in examining the response at those levels. The example and data are from Montgomery [14], page 252.

When a factor is quantitative with three or more levels it is appropriate to partition the sums of squares from that factor into effects of polynomial orders [12]. If there are k levels to the factor, you can partition the sums of squares into k-1 polynomial orders. In this example, the effect due to the quantitative variable temperature can be partitioned into linear and quadratic effects. Similarly, you can partition the interaction. To do this, you must code the quantitative variable with the actual treatment values (that is, code Temperature levels as 100, 125, and 150), use GLM to analyze your data, and declare the quantitative variable to be a covariate.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > General Linear Model.
- 3 In **Responses**, enter *LightOutput*.
- 4 In Model, type Temperature Temperature * Temperature GlassType GlassType * Temperature GlassType * Temperature * Temperature.



5 Click Covariates. In Covariates, enter Temperature.

6 Click **OK** in each dialog box.

Session General Linear Model: LightOutput versus GlassType

window output

Factor Type Levels Values GlassTyp fixed 3 1 2 3

Analysis of Variance for LightOut, using Adjusted SS for Tests

Source	DF	Seq	SS	Adj SS	Adj MS	F	Р
Temperat	1	1779	756	262884	262884	719.21	0.000
Temperat*Temperat	1	190	579	190579	190579	521.39	0.000
GlassTyp	2	150	865	41416	20708	56.65	0.000
GlassTyp*Temperat	2	226	178	51126	25563	69.94	0.000
GlassTyp*Temperat*							
Temperat	2	64	374	64374	32187	88.06	0.000
Error	18	6	579	6579	366		
Total	26	2418	330				
Term	(Coef	SE Coe	f T	Р		
Constant	-496	58.8	191.	3 -25.97	0.000		
Temperat	83.	.867	3.12	7 26.82	0.000		
Temperat*Temperat	-0.28	3516	0.0124	9 -22.83	0.000		
Temperat*GlassTyp							
1	-24	.400	4.42	3 -5.52	0.000		
2	-27	.867	4.42	3 -6.30	0.000		
Temperat*Temperat*GlassTyp							
1	0.11	1236	0.0176	6 6.36	0.000		
2	0.12	2196	0.0176	6 6.91	0.000		
		-					

Unusual Observations for LightOut

0bs	LightOut	Fit	SE Fit	Residual	St Resid
11	1070.00	1035.00	11.04	35.00	2.24R
17	1000.00	1035.00	11.04	-35.00	-2.24R

R denotes an observation with a large standardized residual.

Interpreting the results

MINITAB first displays a table of factors, with their number of levels, and the level values. The second table gives an analysis of variance table. This is followed by a table of coefficients, and then a table of unusual observations.

The Analysis of Variance table gives, for each term in the model, the degrees of freedom, the sequential sums of squares (Seq SS), the adjusted (partial) sums of squares (Adj SS), the adjusted means squares (Adj MS), the F-statistic from the adjusted means squares, and its p-value. The sequential sums of squares is the added sums of squares given that prior terms are in the model. These values depend upon the model order. The adjusted sums of squares are the sums of squares given that all other terms are in the model. These values do not depend upon the model order. If you had selected

sequential sums of squares in the Options subdialog box, MINITAB would use these values for mean squares and F-tests.

In the example, all p-values were printed as 0.000, meaning that they are less than 0.0005. This indicates significant evidence of effects if your level of significance, α , is greater than 0.0005. The significant interaction effects of glass type with both linear and quadratic temperature terms implies that the coefficients of second order regression models of the effect of temperature upon light output depends upon the glass type.

The next table gives the estimated coefficients for the covariate, Temperature, and the interactions of Temperature with GlassType, their standard errors, t-statistics, and p-values. Following the table of coefficients is a table of unusual values. Observations with large standardized residuals or large leverage values are flagged. In our example, two values have standardized residuals whose absolute values are greater than 2.

Example of using GLM and multiple comparisons with an unbalanced nested design

Four chemical companies produce insecticides that can be used to kill mosquitoes, but the composition of the insecticides differs from company to company. An experiment is conducted to test the efficacy of the insecticides by placing 400 mosquitoes inside a glass container treated with a single insecticide and counting the live mosquitoes 4 hours later. Three replications are performed for each product. The goal is to compare the product effectiveness of the different companies. The factors are fixed because you are interested in comparing the particular brands. The factors are nested because each insecticide for each company is unique. The example and data are from Milliken and Johnson [13], page 414. You use GLM to analyze your data because the design is unbalanced and you will use multiple comparisons to compare the mean response for the company brands.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > General Linear Model.
- 3 In Responses, enter NMosquito.
- 4 In Model, enter Company Product(Company).
- 5 Click Comparisons. Under Pairwise Comparisons, enter *Company* in Terms. Click OK in each dialog box.

MINITAB User's Guide 2



CONTENTS INDEX MEET MTB UGUIDE 1 UGUIDE 2 SC QREF HOW TO USE



3								General Lir	near Model
Session window output	General Lin Factor Company Product (Co	near Mode T fi	l: NMosqu ype Level xed	i to vers s Value 4 A B C 1 A1 A2	sus Compa s D A3 B1 B2	ny, Produc	t n2 n3 n4		
	Analysis o	of Variance	for NMos	auit. u	sing Adjus	sted SS fo	r Tests		
	Source Company Product(Co Error Total	D ompany) 2 3	F Seq 3 2281 7 150 2 126 2 2557	SS 3.3 0.6 0.0 3.9	Adj SS 22813.3 1500.6 1260.0	Adj MS 7604.4 214.4 57.3	F 132.78 3.74	P 0.000 0.008	
	Tukey 95.0 Response V All Pairwi	Tukey 95.0% Simultaneous Confidence Intervals Response Variable NMosquit All Pairwise Comparisons among Levels of Company							
	Company =	A subtract	ed from:						
	Company B C D	Lower -2.92 -52.25 -61.69	Center 8.17 -41.17 -52.42	Uppe 19.2 -30.0 -43.1	r 5 8 4 ((* (* *)	+	+ (*)	
						+ -50	+ -25	+0	
	Company =	B subtract	ed from:						
	Company C D	Lower -61.48 -71.10	Center -49.33 -60.58	Uppe -37.1 -50.0	r 9 (7 (*-	+)) +	+	+	
						-50	-25	0	
	Company =	C subtract	ed from:						
	Company D	Lower -21.77	Center -11.25	Uppe -0.734	r 7 	·+	+ (+ -*) +	
						-50	-25	0	
	Tukey Simu Response V All Pairwi	iltaneous T Variable NM se Compari	ests osquit sons amon	g Level	s of Compa	iny			
	Company =	A subtract	ed from:						
	Level Company B C D	Differen of Mea 8. -41. -52.	ce ns Diffe 17 17 42	SE of rence 3.989 3.989 3.337	T-Value 2.05 -10.32 -15.71	Adjusted P-Value 0.2016 0.0000 0.0000			



Fully Nested ANOVA

Company = B subtracted from:

Level Company	Difference of Means	SE of	T-Value	Adjusted
C	-49.33	4.369	-11.29	0.0000
D	-60.58	3.784	-16.01	0.0000
Company = C	subtracted	from:		
Level	Difference	SE of		Adjusted
Company	of Means	Difference	T-Value	P-Value
D	-11.25	3.784	-2.973	0.0329

Interpreting the results

MINITAB displays a factor level table, an ANOVA table, multiple comparison confidence intervals for pairwise differences between companies, and the corresponding multiple comparison hypothesis tests. The ANOVA F-tests indicate that there is significant evidence for company effects.

Examine the multiple comparison confidence intervals. There are three sets: 1) for the company A mean subtracted from the company B, C, and D means; 2) for the company B mean subtracted from the company C and D means; and 3) for the company C mean subtracted from the company D mean. The first interval, for the company B mean minus the company A mean, contains zero is in the confidence interval. Thus, there is no significant evidence at $\alpha = 0.05$ for differences in means. However, there is evidence that all other pairs of means are different, because the confidence intervals for the differences in means do not contain zero. An advantage of confidence intervals is that you can see the magnitude of the differences between the means.

Examine the multiple comparison hypothesis tests. These are laid out in the same way as the confidence intervals. You can see at a glance the mean pairs for which there is significant evidence of differences. The adjusted p-values are small for all but one comparison, that of company A to company B. An advantage of hypothesis tests is that you can see what α level would be required for significant evidence of differences.

Fully Nested ANOVA

Use Fully Nested ANOVA to perform fully nested (hierarchical) analysis of variance and to estimate variance components for each response variable. All factors are implicitly assumed to be random. MINITAB uses sequential (Type I) sums of squares for all calculations.

UGUIDE 2

You can analyze up to 50 response variables with up to 9 factors at one time.

MINITAB User's Guide 2

CONTENTS INDEX

HOW TO USE



Analysis of Variance

MEET MTB

3-48



If your design is not hierarchically nested or if you have fixed factors, use either Balanced ANOVA or GLM. Use GLM if you want to use adjusted sums of squares for a fully nested model.

Data

Set up your worksheet in the same manner as with Balanced ANOVA or GLM: one column for each response variable and one column for each factor, so that there is one row for each observation. The factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See *Ordering Text Categories* in the *Manipulating Data* chapter in *MINITAB User's Guide 1*.

Nesting does not need to be balanced. A nested factor must have at least 2 levels at some level of the nesting factor. If factor B is nested within factor A, there can be unequal levels of B within each level of A. In addition, the subscripts used to identify the B levels can differ within each level of A.

If any response or factor column contains missing data, that entire observation (row) is excluded from all computations. If an observation is missing for one response variable, that row is eliminated for all responses. If you want to eliminate missing rows separately for each response, perform a fully nested ANOVA separately for each response.

To perform an analysis using fully nested ANOVA

Fully Nested ANOVA			×
	<u>R</u> esponses: 		A
	Eactors:		×
			V
Select			
нер		<u>U</u> K	Lancel

1 Choose Stat ➤ ANOVA ➤ Fully Nested ANOVA.

- 2 In **Responses**, enter up to 50 numeric columns containing the response variables.
- 3 In Factors, type in the factors in hierarchical order. See *The fully nested or hierarchical model* below.

Fully Nested ANOVA

Analysis of Variance

The fully nested or hierarchical model

MINITAB fits a fully nested or hierarchical model with the nesting performed according to the order of factors in the **Factors** box. If you enter factors **A B C**, then the model terms will be **A B(A) C(B)**. You do not need to specify these terms in model form as you would for Balanced ANOVA or GLM.

MINITAB uses sequential (Type I) sums of squares for all calculations of fully nested ANOVA. This usually makes sense for a hierarchical model. GLM offers the choice of sequential or adjusted (Type III) sums of squares and uses the adjusted sums of squares by default. These sums of squares can differ when your design is unbalanced. Use GLM if you want to use adjusted sums of squares for calculations.

Example of a fully nested ANOVA

You are an engineer trying to understand the sources of variability in the manufacture of glass jars. The process of making the glass requires mixing materials in small furnaces for which the temperature setting is to be 475 degrees F. Your company has a number of plants where the jars are made, so you select four as a random sample. You conduct an experiment and measure furnace temperature three times during a work shift for each of four operators from each plant over four different shifts. Because your design is fully nested, you use Fully Nested ANOVA to analyze your data.

- 1 Open the worksheet FURNTEMP.MTW.
- 2 Choose Stat > ANOVA > Fully Nested ANOVA.
- 3 In Responses, enter Temp.
- 4 In Factors, enter Plant-Batch. Click OK.

Nested ANOVA: Temp versus Plant, Operator, Shift, Batch

Session window output

Analysis of Variance for Temp

Source	DF	SS	MS	F	Р
Plant	3	731.5156	243.8385	5.854	0.011
Operator	12	499.8125	41.6510	1.303	0.248
Shift	48	1534.9167	31.9774	2.578	0.000
Batch	128	1588.0000	12.4062		
Total	191	4354.2448			

Variance Components

Source	Var Comp.	% of Total	StDev
Plant	4.212	17.59	2.052
Operator	0.806	3.37	0.898
Shift	6.524	27.24	2.554
Batch	12.406	51.80	3.522
Total	23.948		4.894

MINITAB User's Guide 2

3-49



Balanced MANOVA

Expected Mean Squares

```
1 Plant 1.00(4) + 3.00(3) + 12.00(2) + 48.00(1)
2 Operator 1.00(4) + 3.00(3) + 12.00(2)
3 Shift 1.00(4) + 3.00(3)
4 Batch 1.00(4)
```

Interpreting the results

MINITAB displays three tables of output: 1) the ANOVA table, 2) the estimated variance components, and 3) the expected means squares. There are four sequentially nested sources of variability in this experiment: plant, operator, shift, and batch. The ANOVA table indicates that there is significant evidence for plant and shift effects at $\alpha = 0.05$ (F-test p-values < 0.05). There is no significant evidence for an operator effect. The variance component estimates indicate that the variability attributable to batches, shifts, and plants was 52, 27, and 18 percent, respectively, of the total variability.

If a variance component estimate is less than zero, MINITAB displays what the estimate is, but sets the estimate to zero in calculating the percent of total variability.

Balanced MANOVA

Use Balanced MANOVA to perform multivariate analysis of variance (MANOVA) for balanced designs. You can take advantage of the data covariance structure to simultaneously test the equality of means from different responses.

Your design must be balanced, with the exception of one-way designs. *Balanced* means that all treatment combinations (cells) must have the same number of observations. Use General MANOVA (page 3-56) to analyze either balanced and unbalanced MANOVA designs or if you have covariates. You cannot designate factors to be random with general MANOVA, unlike for balanced ANOVA, though you can work around this restriction by supplying error terms to test the model terms.

Factors may be crossed or nested, fixed or random. See *Crossed vs. nested factors* on page 3-19 and *Fixed vs. random factors* on page 3-19.

Data

You need one column for each response variable and one column for each factor, with each row representing an observation. Regardless of whether factors are crossed or nested, use the same form for the data. Factor columns may be numeric, text, or date/ time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See *Ordering Text Categories* in the *Manipulating Data* chapter in *MINITAB User's Guide 1*. You may include up to 50 response variables and up to 9 factors at one time.

Analysis of Variance

Balanced data are required except for one-way designs. The requirement for balanced data extends to nested factors as well. Suppose A has 3 levels, and B is nested within A. If B has 4 levels within the first level of A, B must have 4 levels within the second and third levels of A. MINITAB will tell you if you have unbalanced nesting. In addition, the subscripts used to indicate the 4 levels of B within each level of A must be the same. Thus, the four levels of B cannot be (1 2 3 4) in level 1 of A, (5 6 7 8) in level 2 of A, and (9 10 11 12) in level 3 of A. You can use general MANOVA if you have different levels of B within the levels of A.

If any response or factor column specified contains missing data, that entire observation (row) is excluded from all computations. The requirement that data be balanced must be preserved after missing data are omitted.

To perform a balanced MANOVA

1 Choose Stat > ANOVA > Balanced MANOVA.

Balanced MANOVA		×
	R <u>e</u> sponses:	
	Mo <u>d</u> el:	
		<u></u>
		7
	Random <u>f</u> actors:	
		~
	,	_
	Options	
Select	G <u>r</u> aphs Res <u>u</u> lts	<u>S</u> torage
Help	<u>0</u> K	Cancel

- 2 In **Responses**, enter up to 50 numeric columns containing the response variables.
- 3 In Model, type the model terms that you want to fit. See Overview of Balanced ANOVA and GLM on page 3-18.
- 4 If you like, use one or more of the options described below, then click OK.

Options

Balanced MANOVA dialog box

specify which factors are random factors—see Fixed vs. random factors on page 3-19

Options subdialog box

use the restricted form of the mixed models (both fixed and random effects). The
restricted model forces mixed interaction effects to sum to zero over the fixed effects.
By default, MINITAB fits the unrestricted model. See *Restricted and unrestricted form of
mixed models* on page 3-27.

MINITAB User's Guide 2

3-51



3-52

Balanced MANOVA

Graphs subdialog box

- draw five different residual plots. You can display the following plots:
 - histogram
 - normal probability plot
 - plot of residuals versus the fitted values (Y)
 - plot of residuals versus data order. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n
 - separate plot for the residuals versus each specified column

For a discussion of the residual plots, see Residual plots on page 2-6.

Results subdialog box

- display different MANOVA output. You can request the display of the hypothesis matrix H, the error matrix E, and a matrix of partial correlations (see MANOVA tests on page 3-53), the eigenvalues and eigenvalues for the matrix E⁻¹ H, univariate analysis of variance for each response, and when you have requested univariate analyses of variance, the expected means squares.
- display a table of means corresponding to specified terms from the model. For example, if you specify A B D A*B*D, four tables of means will be printed, one for each main effect, A, B, D, and one for the three-way interaction, A*B*D.
- perform four multivariate tests for model terms that you specify. See Specifying terms to test on page 3-52. Default tests are performed for all model terms.

Storage subdialog box

 store the fits and residuals separately for each response. If you fit a full model, fits are cell means. If you fit a reduced model, fits are least squares estimates. See Specifying reduced models on page 3-21.

Specifying terms to test

In the Results subdialog box, you can specify model terms in **Custom multivariate test for the following terms** and designate an error term in **Error** and MINITAB will perform four multivariate tests (see below) for those terms. This option is probably less useful for balanced MANOVA than it is for general MANOVA; because you can specify factors to be random with balanced MANOVA, MINITAB will use the correct error terms. This option exists for special purpose tests.

If you specify an error term, it must be a single term that is in the model. This error term is used for all requested tests. If you do not specify an error term, MINITAB determines an appropriate error term.

MINITAB User's Guide 2

HOW TO USE

Analysis of Variance

MANOVA tests

MINITAB automatically performs four multivariate tests—Wilk's test, Lawley-Hotelling test, Pillai's test, and Roy's largest root test—for each term in the model and for specially requested terms (see above). All four tests are based on two SSCP (sums of squares and cross products) matrices: H, the hypothesis matrix and E, the error matrix. There is one H associated with each term. E is the matrix associated with the error for the test. These matrices are printed when you request the hypothesis matrices and are labeled by SSCP Matrix.

The test statistics can be expressed in terms of either H and/or E or the eigenvalues of E^{-1} H. You can request to have these eigenvalues printed. (If the eigenvalues are repeated, corresponding eigenvectors are not unique and in this case, the eigenvectors MINITAB prints and those in books or other software may not agree. The MANOVA tests, however, are always unique.) See Help for computational details on the tests.

You can also print the matrix of partial correlations, which are the correlations among the residuals, or alternatively, the correlations among the responses conditioned on the model. The formula for this matrix is $W^{-.5} \in W^{-.5}$, where E is the error matrix and W has the diagonal of E as its diagonal and 0's off the diagonal.

Hotelling's T² Test

Hotelling's T² test to compare the mean vectors of two groups is a special case of MANOVA, using one factor that has two levels. MINITAB'S MANOVA option can be used to do this test. The usual T² test statistic can be calculated from MINITAB's output using the relationship T²=(N-2)U, where N is the total number of observations and U is the Lawley-Hotelling trace. See Help for calculations.

Example of balanced MANOVA

You perform a study in order to determine optimum conditions for extruding plastic film. You measure three responses—tear resistance, gloss, and opacity—five times at each combination of two factors—rate of extrusion and amount of an additive—each set at low and high levels. The data and example are from Johnson and Wichern [10], page 266. You use Balanced MANOVA to test the equality of means because the design is balanced.

- 1 Open the worksheet EXH_MVAR.MTW.
- 2 Choose Stat > ANOVA > Balanced MANOVA.
- 3 In Responses, enter Tear Gloss Opacity.
- 4 In Model, enter Extrusion | Additive.
- 5 Click **Results**. Under **Display of Results**, check **Matrices (hypothesis, error, partial correlations)** and **Eigen analysis**. Click **OK** in each dialog box.

MINITAB User's Guide 2



ENTS INDEX MEET MTB UGUIDE 1 UGUIDE 2 SC QREF HOW



Session window	ANOVA: Tear, Gloss, Opacity versus Extrusion, Additive				
output	MANOVA for Extrusio $s = 1 m = 0.5 n = 6.0$				
	CriterionTest StatisticFDFPWilk's0.381867.554(3, 14)0.003Lawley-Hotelling1.618777.554(3, 14)0.003Pillai's0.618147.554(3, 14)0.003Roy's1.61877				
	SSCP Matrix for Extrusio				
	TearGlossOpacityTear1.740-1.5040.8555Gloss-1.5041.301-0.7395Opacity0.855-0.7390.4205				
	SSCP Matrix for Error				
	TearGlossOpacityTear1.7640.0200-3.070Gloss0.0202.6280-0.552Opacity-3.070-0.552064.924				
	Partial Correlations for the Error SSCP Matrix				
	Tear Gloss Opacity Tear 1.00000 0.00929 -0.28687 Gloss 0.00929 1.00000 -0.04226 Opacity -0.28687 -0.04226 1.00000				
	EIGEN Analysis for Extrusio				
	Eigenvalue 1.619 0.00000 0.00000 Proportion 1.000 0.00000 0.00000 Cumulative 1.000 1.00000 1.00000				
	Eigenvector 1 2 3 Tear 0.6541 0.4315 0.0604 Gloss -0.3385 0.5163 0.0012 Opacity 0.0359 0.0302 -0.1209				
	multivariate output for Additive and Extrusion*Additive would follow				
	Interpreting the results				
	By default, MINITAB displays a table of the four multivariate tests (Wilk's, Lawley-Hotelling, Pillai's, and Roy's) for each term in the model. The values s, m, and n are used in the calculations of the F-statistics for Wilk's, Lawley-Hotelling, and Pillai's tests. The F-statistic is exact if s = 1 or 2, otherwise it is approximate [10]. Because you requested the display of additional matrices (hypothesis, error, and partial correlations) and an eigen analysis, this information is also displayed. The output is shown only for one model term, Extrusion, and not for the terms Additive or Extrusion*Additive.				

MINITAB User's Guide 2

Chapter 3

Analysis of Variance

Examine the p-values for the Wilk's, Lawley-Hotelling, and Pillai's test statistic to judge whether there is significant evidence for model effects. These values are 0.003 for the model term Extrusion, indicating that there is significant evidence for Extrusion main effects at α levels greater than 0.003. The corresponding p-values for Additive and for Additive*Extrusion are 0.025 and 0.302, respectively (not shown), indicating that there is no significant evidence for interaction, but there is significant evidence for Extrusion and Additive main effects at α levels of 0.05 or 0.10.

You can use the SSCP matrices to assess the partitioning of variability in a similar way as you would look at univariate sums of squares. The matrix labeled as SSCP Matrix for Extrusn is the hypothesis sums of squares and cross-products matrix, or H, for the three response with model term Extrusion. The diagonal elements of this matrix, 1.740, 1.301, and 0.4205, are the univariate ANOVA sums of squares for the model term Extrusion when the response variables are Tear, Gloss, and Opacity, respectfully. The off-diagonal elements of this matrix are the cross products.

The matrix labeled as SSCP Matrix for Error is the error sums of squares and cross-products matrix, or E. The diagonal elements of this matrix, 1.764, 2.6280, and 64.924, are the univariate ANOVA error sums of squares when the response variables are Tear, Gloss, and Opacity, respectfully. The off-diagonal elements of this matrix are the cross products. This matrix is printed once, after the SSCP matrix for the first model term.

You can use the matrix of partial correlations, labeled as Partial Correlations for the Error SSCP Matrix, to assess how related the response variables are. These are the correlations among the residuals or, equivalently, the correlations among the responses conditioned on the model. Examine the off-diagonal elements. The partial correlations between Tear and Gloss of 0.00929 and between Gloss and Opacity of -0.04226 are small. The partial correlation of -0.28687 between Tear and Opacity is not large. Because the correlation structure is weak, you might be satisfied with performing univariate ANOVA for these three responses. This matrix is printed once, after the SSCP matrix for error.

You can use the eigen analysis to assess how the response means differ among the levels of the different model terms. The eigen analysis is of E⁻¹ H, where E is the error SCCP matrix and H is the response variable SCCP matrix. These are the eigenvalues that are used to calculate the four MANOVA tests.

Place the highest importance on the eigenvectors that correspond to high eigenvalues. In the example, the second and third eigenvalues are zero and therefore the corresponding eigenvectors are meaningless. For both factors, Extrusion and Additive, the first eigenvectors contain similar information The first eigen vector for Extrusion is 0.6541, -0.3385, 0.0359 and for Additive it is -0.6630, -0.3214, -0.0684 (not shown). The highest absolute value within these eigenvectors is for the response Tear, the second highest is for Gloss, and the value for Opacity is small. This implies that the Tear means have the largest differences between the two factor levels of either Extrusion or Additive, the Gloss means have the next largest differences, and the Opacity means have small differences.

MINITAB User's Guide 2

3-55

General MANOVA

Use general MANOVA to perform multivariate analysis of variance (MANOVA) with balanced and unbalanced designs or if you have covariates. This procedure takes advantage of the data covariance structure to simultaneously test the equality of means from different responses.

Calculations are done using a regression approach. A "full rank" design matrix is formed from the factors and covariates and each response variable is regressed on the columns of the design matrix.

Factors may be crossed or nested, but they cannot be declared as random; it is possible to work around this restriction by specifying the error term to test model terms (see *Specifying terms to test* on page 3-58). Covariates may be crossed with each other or with factors, or nested within factors. You can analyze up to 50 response variables with up to 9 factors and 50 covariates at one time.

Data

Set up your worksheet in the same manner as with balanced MANOVA: one column for each response variable, one column for each factor, and one column for each covariate, so that there is one row of the worksheet for each observation. The factor columns may be numeric, text, or date/time. If you wish to change the order in which text categories are processed from their default alphabetical order, you can define your own order. See *Ordering Text Categories* in the *Manipulating Data* chapter in *MINITAB User's Guide 1*.

Although models can be unbalanced in general MANOVA, they must be "full rank." That is, there must be enough data to estimate all the terms in your model. For example, suppose you have a two-factor crossed model with one empty cell. Then you can fit the model with terms A B, but not A B A*B. MINITAB will tell you if your model is not full rank. In most cases, eliminating some of the high order interactions in your model (assuming, of course, they are not important) can solve non-full rank problems.

Nesting does not need to be balanced. If factor B is nested within factor A, there can be unequal levels of B within each level of A. In addition, the subscripts used to identify the B levels can differ within each level of A.

If any response, factor, or covariate column contains missing data, that entire observation (row) is excluded from all computations.





General MANOVA

Analysis of Variance

To perform an analysis using general MANOVA

1 Choose Stat > ANOVA > General MANOVA.

General MANOVA			×
	R <u>e</u> sponses:		
	Mo <u>d</u> el:		×
		Covariates	Options
Select	G <u>r</u> aphs	Res <u>u</u> lts	<u>S</u> torage
Help		<u>0</u> K	Cancel

- 2 In Responses, enter up to 50 numeric columns containing the response variables.
- 3 In **Model**, type the model terms you want to fit. See *Overview of Balanced ANOVA and GLM* on page 3-18.
- 4 If you like, use one or more of the options described below, then click OK.

Options

Covariates subdialog box

include up to 50 covariates in the model

Options subdialog box

enter a column containing weights to perform weighted regression—see Weighted regression on page 2-6

Graphs subdialog box

- draw five different residual plots for regular, standardized, or deleted residuals—see Choosing a residual type on page 2-5. Available residual plots include a
 - histogram
 - normal probability plot
 - plot of residuals versus the fitted values (\hat{Y})
 - plot of residuals versus data order. The row number for each data point is shown on the x-axis—for example, 1 2 3 4... n
 - separate plot for the residuals versus each specified column

For a discussion of the residual plots, see *Residual plots* on page 2-6.

General MANOVA

Results subdialog box

- display different MANOVA output. You can request the display of the hypothesis matrix H, the error matrix E, and a matrix of partial correlations (see MANOVA tests on page 3-58), the eigenvalues and eigenvalues for the matrix E⁻¹ H, and univariate analysis of variance for each response.
- display a table of means corresponding to specified terms from the model. For example, if you specify A B D A*B*D, four tables of means will be printed, one for each main effect, A, B, D, and one for the three-way interaction, A*B*D.
- perform 4 multivariate tests for model terms that you specify. See Specifying terms to test on page 3-58. Default tests are performed for all model terms.

Storage subdialog box

- store model coefficients and fits in separate columns for each response.
- regular, standardized, and deleted residuals separately for each response—see Choosing a residual type on page 2-5.
- store leverages, Cook's distances, and DFITS, for identifying outliers—see Identifying outliers on page 2-9.
- store the design matrix. The design matrix multiplied by the coefficients will yield the fitted values. See Design matrix used by General Linear Model on page 3-42.

Specifying terms to test

In the Results subdialog box, you can specify model terms in **Custom multivariate test for the following terms** and designate the error term in **Error**. MINITAB will perform four multivariate tests (see below) for those terms. This option is most useful when you have factors that you consider as random factors. Model terms that are random or that are interactions with random terms may need a different error term than general MANOVA supplies. You can determine the appropriate error term by entering one response variable with General Linear Model (page 3-36), choose to display the expected mean squares, and determine which error term was used for each model terms (see *Expected mean squares* on page 3-27).

If you specify an error term, it must be a single term that is in the model. This error term is used for all requested tests. If you have different error terms for certain model terms, enter each separately and exercise the general MANOVA dialog for each one. If you do not specify an error term, MINITAB uses MSE.

MANOVA tests

3-58

The MANOVA tests with general MANOVA are similar to those performed for balanced MANOVA. See *MANOVA tests* on page 3-53 for details.

Test for Equal Variances

However, with general MANOVA, there are two SSCP matrices associated with each term in the model, the sequential SSCP matrix and the adjusted SSCP matrix. These matrices are analogous to the sequential SS and adjusted SS in univariate General Linear Model (see page 3-36). In fact, the univariate SS's are along the diagonal of the corresponding SSCP matrix. If you do not specify an error term in **Error** when you enter terms in **Custom multivariate tests for the following terms**, then the adjusted SSCP matrix is used for H and the SSCP matrix associated with MSE is used for E. If you do specify an error term, the sequential SSCP matrices associated with H and E are used. Using sequential SSCP matrices guarantees that H and E are statistically independent. See Help for details on these tests.

You can also perform Hotelling's T^2 test to compare the mean vectors of two groups (see *Hotelling's* T^2 *Test* on page 3-53). Refer to *Example of balanced MANOVA* on page 3-53 for an example of MANOVA. The dialog operation of general MANOVA is similar to that of balanced MANOVA.

Test for Equal Variances

Use the test for equal variances to perform hypothesis tests for equality or homogeneity of variance using Bartlett's and Levene's tests. An F-test replaces Bartlett's test when you have just two levels.

Many statistical procedures, including analysis of variance, assume that although different samples may come from populations with different means, they have the same variance. The effect of unequal variances upon inferences depends upon whether your model includes fixed or random effects, disparities in sample sizes, and the choice of multiple comparison procedure. The ANOVA F-test is only slightly affected by inequality of variance if the model only contains fixed factors and has equal or nearly equal sample sizes. F-tests involving random effects may be substantially affected, however [18]. Use the variance test procedure to test the validity of the equal variance assumption.

Data

Set up your worksheet with one column for the response variable and one column for each factor, so that there is one row for each observation. Your response data must be in one column. You may have up to 9 factors. Factor columns may be numeric, text, or date/time, and may contain any value. If there are many cells (factors and levels), the print in the output chart can get very small.

Rows where the response column contains missing data (*) are automatically omitted from the calculations. When one or more factor columns contain missing data, MINITAB displays the chart and Bartlett's test results, but not the Levene's test results.

Data limitations include: (1) if none of the cells have multiple observations, nothing is calculated. In addition, there must be at least one nonzero standard deviation; (2) the

MINITAB User's Guide 2

3-59



D

Chapter 3

3-60

Test for Equal Variances

F-test for 2 levels requires both cells to have multiple observations; (3) Bartlett's test requires two or more cells to have multiple observations; (4) Levene's test requires two or more cells to have multiple observations, but one cell must have three or more.

To perform a test for equal variances

1 Choose Stat > ANOVA > Test for Equal Variances.

Test for Equal Variances	×
	Response:
	Eactors:
	Confidence level: 95.0
	<u>T</u> itle:
Select	<u>S</u> torage
Help	<u>Q</u> K Cancel

- 2 In **Response**, enter the column containing the response.
- 3 In Factors, enter up to nine columns containing the factor levels.
- 4 If you like, use one or more of the options described below, then click OK.

Options

Test for Equal Variances dialog box

- specify a confidence level for the confidence interval (the default is 95%)
- replace the default graph title with your own title

Storage subdialog box

 store standard deviations, variances, and/or upper and lower confidence limits for σ by factor levels

Bartlett's versus Levene's tests

MINITAB calculates and displays a test statistic and p-value for both Bartlett's test and Levene's test where the null hypothesis is of equal variances versus the alternative of not all variances being equal. (When there are only two levels, an F-test is performed in place of Bartlett's test.) Use Bartlett's test when the data come from a normal distribution and Levene's test when the data come from a continuous, but not

Test for Equal Variances

Analysis of Variance

necessarily normal, distribution. Bartlett's test is not robust to departures from normality.

The computational method for Levene's Test is a modification of Levene's procedure [11] developed by [2]. This method considers the distances of the observations from their sample median rather than their sample mean. Using the sample median rather than the sample mean makes the test more robust for smaller samples.

See Help for the computational form of these tests.

Example of performing a test for equal variances

You study conditions conducive to potato rot by injecting potatoes with bacteria that cause rotting and subjecting them to different temperature and oxygen regimes. Before performing analysis of variance, you check the equal variance assumption using the test for equal variances.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Test for Equal Variances.
- 3 In **Response**, enter *Rot*.
- 4 In Factors, enter Temp Oxygen. Click OK.

Session	Test for	Equal	Variances
---------	----------	-------	-----------

vindow	Response	Rot	
output	Factors ConfLv1	Temp Oxygen 95.0000	١
	00 EV 1	50.0000	

Bonferroni confidence intervals for standard deviations

Lower	Sigma	Upper	Ν	Factor	Levels
2.26029	5.29150	81.890	3	10	2
1.28146	3.00000	46.427	3	10	6
2.80104	6.55744	101.481	3	10	10
1.54013	3.60555	55.799	3	16	2
1.50012	3.51188	54.349	3	16	6
3.55677	8.32666	128.862	3	16	10

Bartlett's Test (normal distribution)

Test Statistic: 2.712 P-Value : 0.744

Levene's Test (any continuous distribution)

Test Statistic: 0.372 P-Value : 0.858

MINITAB User's Guide 2

W

Interval Plot for Mean

Graph window output



Interpreting the results

The test for equal variances generates a plot that displays Bonferroni 95% confidence intervals for the response standard deviation at each level. Bartlett's and Levene's test results are displayed in both the Session window and in the graph. Note that the 95% confidence level applies to the the family of intervals and the asymmetry of the intervals is due to the skewness of the chi-square distribution.

For the potato rot example, the p-values of 0.744 and 0.858 are greater than reasonable choices of α , so you fail to reject the null hypothesis of the variances being equal. That is, these data do not provide enough evidence to claim that the populations have unequal variances

Interval Plot for Mean

Use Interval Plot to produce a plot of group means and standard error bars or confidence intervals about the means. This plot illustrates both a measure of central tendency and variability of the data.

Data

3-62

The response (Y variable) data must be stacked in one numeric column. You must also have a column that contains the group identifiers. The grouping column can be numeric, text, or date/time. If you wish to change the order in which text levels are processed, you can define your own order. See *Ordering Text Categories* in the *Manipulating Data* chapter in *MINITAB User's Guide 1*.

Interval Plot for Mean

Analysis of Variance

Special cases include one observation in a group or a standard deviation of 0 (such as when all observations are the same). In the first case, the mean is plotted, but not the interval bar. In the second case, you see a symbol for the mean *and* a horizontal interval bar.

MINITAB automatically omits rows with missing responses or factor levels from the calculations.

To display an interval plot for the mean

1 Choose Stat > ANOVA > Interval Plot.

Interval Plot for Mean		×
	Y variable:	
	<u>G</u> roup variable:	
	Type of interval plot	Display mean as
	Standard error	Symbol
	Multiple: 1.0	○ <u>B</u> ar
	C <u>C</u> onfidence interval	Side: two-sided 🔻
	Level: 95.0	- ,
		Pool error across groups
	X <u>a</u> xis label:	
	Y axis lab <u>e</u> l:	
1	<u>T</u> itle:	
Select	Sy <u>m</u> bol	Bar Interval Line
Help	[<u>O</u> K Cancel

- 2 In Y variable, enter the column containing the response data.
- 3 In Group variable, enter the column containing the grouping variable or subscripts.
- 4 If you like, use any of the options listed below, then click OK.

Options

Interval Plot for Mean dialog box

- determine the type of interval displayed on the plot. You can display
 - the default plot which uses standard error bars. That is, the error bars are $(s)/\sqrt{n}$ away from the mean. You also can specify a multiplier for the standard error bars. For example, specifying the multiplier allows you to display error bars that are two times the standard error away from the mean.
 - display error bars that show a normal distribution confidence interval for the mean (rather than using the standard error). You can change the confidence level from the default 95%.
- display a symbol at the mean position or a bar that extends from the x-axis (or a specified base) to the mean.

3-64



Interval Plot for Mean

- display error bars (or confidence intervals) above the mean (upper one-sided), below the mean (lower one-sided), or both above and below the mean (two-sided).
- pool the standard error across all subgroups instead of calculating the standard error for each subgroup separately.
- replace the default x- and y-axis labels with your own labels.
- replace the default graph title with your own title.

Symbol subdialog box

set the type, color, and size of symbols at each subgroup mean.

Bar subdialog box

set the fill type, foreground color, background color, edge size, and base (y-value to which bars extend to from the mean) of the bars.

Interval Line subdialog box

specify the type, color, and size of the error bar lines at each subgroup mean.

Example of an interval plot

Six varieties of alfalfa were grown on plots within four different fields. You are interested in comparing yields of the different varieties. After harvest, you wish to examine means with their standard errors using an error bar plot.

- 1 Open the worksheet ALFALFA.MTW.
- 2 Choose Stat > ANOVA > Interval Plot.
- 3 In Y variable, enter Yield.
- 4 In Group variable, enter Variety. Click OK.



Main Effects Plot

Analysis of Variance

Interpreting the results

The error bar plot plots the means of each alfalfa variety at the symbols with lines extending one standard error above and below the means. The variability between varieties appears to be large relative to the variability within varieties, as there is some distance between some of the error bars for the different varieties.

Main Effects Plot

Use Main Effects Plot to plot data means when you have multiple factors. The points in the plot are the means of the response variable at the various levels of each factor, with a reference line drawn at the grand mean of the response data. Use the main effects plot for comparing magnitudes of main effects.

Use the main effects plot described on page19-53 to generate main effects plots specifically for two-level factorial designs.

Data

Set up your worksheet with one column for the response variable and one column for each factor, so that each row in the response and factor columns represents one observation. It is not required that your data be balanced.

The factor columns may be numeric, text, or date/time and may contain any values. If you wish to change the order in which text levels are processed, you can define your own order. See *Ordering Text Categories* in the *Manipulating Data* chapter in *MINITAB* User's Guide 1. You may have up to 9 factors.

Missing values are automatically omitted from calculations.

To perform a main effects plot

1 Choose Stat > ANOVA > Main Effects Plot.

Main Effects Plot	×
	Responses:
	Eactors:
	A N
Select	O <u>p</u> tions
Help	<u>O</u> K Cancel

UGUIDE 2

2 In Responses, enter the column(s) containing the response data.



CONTENTS INDEX MEET MTB

HOW TO USE

Main Effects Plot

- 3 In **Factors**, enter the columns containing the factor levels. You can enter up to 9 factors.
- 4 If you like, use any of the options listed below, then click OK.

Options

Graph window

output

Options subdialog box

- specify the y-value(s) to use for the minimum and/or the maximum of the graph scale.
- you can replace the default graph title with your own title.

Example of a main effects plot

You grow six varieties of alfalfa on plots within four different fields and you weigh the yield of the cuttings. You are interested in comparing yields from the different varieties and consider the fields to be blocks. You want to preview the data and examine yield by variety and field using the main effects plot.

- 1 Open the worksheet ALFALFA.MTW.
- 2 Choose Stat > ANOVA > Main Effects Plot.
- 3 In Responses, enter Yield.
- 4 In Factors, enter Variety Field. Click OK.



Interpreting the results

The main effects plot displays the response means for each factor level in sorted order if the factors are numeric or date/time or in alphabetical order if text, unless value ordering has been assigned (see *Ordering Text Categories* in the *Manipulating Data* chapter in *MINITAB User's Guide 1*). A horizontal line is drawn at the grand mean. The effects are the differences between the means and the reference line. In the example, the variety effects upon yield are large compared to the effects of field (the blocking variable).

Interactions Plot

Analysis of Variance

Interactions Plot

Interactions Plot creates a single interaction plot for two factors, or a matrix of interaction plots for three to nine factors. An interactions plot is a plot of means for each level of a factor with the level of a second factor held constant. Interactions plots are useful for judging the presence of interaction.

Interaction is present when the response at a factor level depends upon the level(s) of other factors. Parallel lines in an interactions plot indicate no interaction. The greater the departure of the lines from the parallel state, the higher the degree of interaction. To use interactions plot, data must be available from all combinations of levels.

Use the Interactions Plot in Chapter 19 to generate interaction plots specifically for 2-level factorial designs, such as those generated by Fractional Factorial Design, Central Composite Design, and Box-Behnken Design.

Data

Set up your worksheet with one column for the response variable and one column for each factor, so that each row in the response and factor columns represents one observation. Your data is not required to be balanced.

The factor columns may be numeric, text, or date/time and may contain any values. If you wish to change the order in which text levels are processed, you can define your own order. See *Ordering Text Categories* in the *Manipulating Data* chapter in *MINITAB* User's Guide 1. You may have from 2 through 9 factors.

Missing data are automatically omitted from calculations.

To display an interactions plot

MEET MTB

INDEX

1 Choose Stat > ANOVA > Interactions Plot.

nteractions Plot	×
	Responses: Factors:
	<u>D</u> isplay full interaction plot matrix
Select	Options
Help	<u>O</u> K Cancel

- 2 In **Responses**, enter the column(s) containing the response data.
- 3 In Factors, enter from 2 to 9 columns containing the factor levels. If you have two factors, the x-variable will be the second factor that you enter.

UGUIDE 2

4 If you like, use any of the options listed below, then click OK.

UGUIDE 1

MINITAB User's Guide 2



HOW TO USE

Ð

Interactions Plot

Options

Main Effects Plot dialog box

 display the full interaction matrix for more than two factors, rather than the default upper right portion of the matrix.

Options subdialog box

- specify the y-value(s) to use for the minimum of the graph scale. You can enter one value to be used for all plots or one value for each response.
- specify the y-value(s) to use for the maximum of the graph scale. You can enter one value to be used for all plots or one value for each response.
- replace the default graph title with your own title.

Example of an interaction plot with two factors

You conduct an experiment to test the effect of temperature and glass type upon the light output of an oscilloscope (example and data from [14], page 252). There are three glass types and three temperatures, 100, 125, and 150 degrees Fahrenheit. You choose interactions plot to visually assess interaction in the data. You enter the quantitative variable second because you want this variable as the x variable in the plot.

- 1 Open the worksheet EXH_AOV.MTW.
- 2 Choose Stat > ANOVA > Interactions Plot.
- 3 In Responses, enter LightOutput.
- 4 In Factors, enter GlassType Temperature. Click OK.



Graph window output

MINITAB User's Guide 2

3-68

Interactions Plot

Analysis of Variance

Interpreting the results

This interaction plot shows the mean light output versus the temperature for each of the three glass types. The legend shows which symbols are assigned to the glass types. The means of the factor levels are plotted in sorted order if numeric or date/time or in alphabetical order if text, unless value ordering has been assigned (see Ordering Text Categories in the Manipulating Data chapter in MINITAB User's Guide 1).

This plot shows apparent interaction because the lines are not parallel, implying that the effect of temperature upon light output depends upon the glass type. You test this using GLM on page 3-43.

Example of an interaction plot with more than two factors

Plywood is made by cutting a thin layer of wood from logs as they are spun on their axis. Considerable force is required to turn a log hard enough so that a sharp blade can cut off a layer. Chucks are inserted into the ends of the log to apply the torque necessary to turn the log. You conduct an experiment to study factors that affect torque. These factors are diameter of the logs, penetration distance of the chuck into the log, and the temperature of the log. You wish to preview the data to check for the presence of interaction.

- Open the worksheet PLYWOOD.MTW.
- 2 Choose Stat > ANOVA > Interactions Plot.
- 3 In Responses, enter Torque.
- 4 In Factors, enter Diameter-Temp. Click OK.

Graph window output



Interpreting the results

MEET MTB

INDEX

UGUIDE 1

An interaction plot with three or more factors show separate two-way interaction plots for all two-factor combinations. In this example, the plot in the middle of the top row shows the mean torque versus the penetration levels for both levels of diameter, 4.5 and 7.5, averaged over all levels of temperature. There are analogous interactions plots for diameter by temperature (upper right) and penetration by temperature (second row).

UGUIDE 2

HOW TO USE

References

For this example, the diameter by penetration and the diameter by temperature plots show nonparallel lines, indicating interaction. The presence of penetration by temperature interaction is not so easy to judge. This interaction might best be judged in conjunction with a model-fitting procedure, such as GLM.

References

- [1] R.E. Bechhofer and C.W. Dunnett (1988). "Percentage points of multivariate Student t distributions," *Selected Tables in Mathematical Studies*, Vol.11. American Mathematical Society, Providence, R.I.
- [2] M.B. Brown and A.B. Forsythe (1974). *Journal of the American Statistical Association*, 69, 364–367.
- [3] H.L. Harter (1970). Order Statistics and Their Uses in Testing and Estimation, Vol.1. U.S. Government Printing Office, Washington, D.C.
- [4] A.J. Hayter (1984). "A proof of the conjecture that the Tukey-Kramer multiple comparisons procedure is conservative," *Annals of Statistics*, 12, pp.61–75.
- [5] D.L. Heck (1960). "Charts of Some Upper Percentage Points of the Distribution of the Largest Characteristic Root," *The Annals of Statistics*, pp.625–642.
- [6] C.R. Hicks (1982). *Fundamental Concepts in the Design of Experiments*, Third Edition, CBC College Publishing.
- [7] Y. Hochberg and A.C. Tamhane (1987). *Multiple Comparison Procedures*. John Wiley & Sons, New York.
- [8] J.C. Hsu (1984). "Constrained Two-Sided Simultaneous Confidence Intervals for Multiple Comparisons with the Best," Annals of Statistics, 12, pp.1136–1144.
- [9] J.C. Hsu (1996). *Multiple Comparisons, Theory and methods*, Chapman & Hall, New York.
- [10] R. Johnson and D. Wichern (1992). *Applied Multivariate Statistical Methods*, Third Edition, Prentice Hall.
- [11] H. Levene (1960). Contributions to Probability and Statistics, pp.278–292. Stanford University Press, CA.
- [12] T.M. Little (1981). "Interpretation and Presentation of Result," *HortScience*, 19, pp.637-640.
- [13] G.A. Milliken and D.E. Johnson (1984). *Analysis of Messy Data. Volume I: Designed Experiments*, Van Nostrand Reinhold.

References

Analysis of Variance

- [14] D.C. Montgomery (1991). Design and Analysis of Experiments, Third Edition, John Wiley & Sons.
- [15] D. Morrison (1967). *Multivariate Statistical Methods*, McGraw-Hill.
- [16] L.S. Nelson (1974). "Factors for the Analysis of Means," *Journal of Quality Technology*, 6, pp.175–181.
- [17] P.R. Nelson (1983). "A Comparison of Sample Sizes for the Analysis of Means and the Analysis of Variance," *Journal of Quality Technology*, 15, pp.33–39.
- [18] J. Neter, W. Wasserman and M.H. Kutner (1985). *Applied Linear Statistical Models*, Second Edition, Irwin, Inc.
- [19] R.A. Olshen (1973). "The conditional level of the F-test," *Journal of the American Statistical Association*, 68, pp.692–698.
- [20] E.R. Ott (1983). "Analysis of Means—A Graphical Procedure," *Journal of Quality Technology*, 15, pp.10–18.
- [21] E.R. Ott and E.G. Schilling (1990). *Process Quality Control—Troubleshooting and Interpretation of Data*, 2nd Edition, McGraw-Hill.
- [22] P.R. Ramig (1983). "Applications of the Analysis of Means," *Journal of Quality Technology*, 15, pp.19–25.
- [23] E.G. Schilling (1973). "A Systematic Approach to the Analysis of Means," *Journal of Quality Technology*, 5, pp.93–108, 147–159.
- [24] S.R. Searle, G. Casella, and C.E. McCulloch (1992). Variance Components, John Wiley & Sons.
- [25] N.R. Ullman (1989). "The Analysis of Means (ANOM) for Signal and Noise," Journal of Quality Technology, 21, pp.111–127.
- [26] E. Uusipaikka (1985). "Exact simultaneous confidence intervals for multiple comparisons among three or four mean values," *Journal of the American Statistical Association*, 80, pp.196–201.
- [27] B.J. Winer (1971). *Statistical Principals in Experimental Design*, Second Edition, McGraw-Hill.

Acknowledgment

We are grateful for assistance in the design and implementation of multiple comparisons from Jason C. Hsu, Department of Statistics, Ohio State University and for the guidance of James L. Rosenberger, Statistics Department, The Pennsylvania State University, in developing the Balanced ANOVA, Analysis of Covariance, and General Linear Models procedures.





D

