



THE
POWER
TO KNOW.

SAS/STAT[®] 14.1 User's Guide Introduction to Analysis of Variance Procedures

This document is an individual chapter from *SAS/STAT® 14.1 User's Guide*.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2015. *SAS/STAT® 14.1 User's Guide*. Cary, NC: SAS Institute Inc.

SAS/STAT® 14.1 User's Guide

Copyright © 2015, SAS Institute Inc., Cary, NC, USA

All Rights Reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

U.S. Government License Rights; Restricted Rights: The Software and its documentation is commercial computer software developed at private expense and is provided with RESTRICTED RIGHTS to the United States Government. Use, duplication, or disclosure of the Software by the United States Government is subject to the license terms of this Agreement pursuant to, as applicable, FAR 12.212, DFAR 227.7202-1(a), DFAR 227.7202-3(a), and DFAR 227.7202-4, and, to the extent required under U.S. federal law, the minimum restricted rights as set out in FAR 52.227-19 (DEC 2007). If FAR 52.227-19 is applicable, this provision serves as notice under clause (c) thereof and no other notice is required to be affixed to the Software or documentation. The Government's rights in Software and documentation shall be only those set forth in this Agreement.

SAS Institute Inc., SAS Campus Drive, Cary, NC 27513-2414

July 2015

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

Chapter 5

Introduction to Analysis of Variance Procedures

Contents

Overview: Analysis of Variance Procedures	105
Procedures That Perform Sum of Squares Analysis of Variance	106
Procedures That Perform General Analysis of Variance	107
Statistical Details for Analysis of Variance	108
From Sums of Squares to Linear Hypotheses	108
Tests of Effects Based on Expected Mean Squares	109
Analysis of Variance for Fixed-Effect Models	110
PROC GLM for General Linear Models	110
PROC ANOVA for Balanced Designs	111
Comparing Group Means	111
PROC TTEST for Comparing Two Groups	112
Analysis of Variance for Categorical Data and Generalized Linear Models	112
Nonparametric Analysis of Variance	113
Constructing Analysis of Variance Designs	113
For More Information	113
References	114

Overview: Analysis of Variance Procedures

The statistical term “analysis of variance” is used in a variety of circumstances in statistical theory and applications. In the *narrowest* sense, and the original sense of the phrase, it signifies a decomposition of a variance into contributing components. This was the sense used by R. A. Fisher when he defined the term to mean the expression of genetic variance as a sum of variance components due to environment, heredity, and so forth:

$$\sigma^2 = \sigma_1^2 + \sigma_2^2 + \cdots + \sigma_p^2$$

In this sense of the term, the SAS/STAT procedures that fit variance component models, such as the GLIMMIX, HPMIXED, MIXED, NESTED, and VARCOMP procedures, are “true” analysis of variance procedures.

Analysis of variance methodology in a slightly broader sense—and the sense most frequently understood today—applies the idea of an additive decomposition of variance to an additive decomposition of *sums of*

squares, whose expected values are functionally related to components of variation. A collection of sums of squares that measure and can be used for inference about meaningful features of a model is called a *sum of squares analysis of variance*, whether or not such a collection is an additive decomposition. In a linear model, the decomposition of sums of squares can be expressed in terms of projections onto orthogonal subspaces spanned by the columns of the design matrix \mathbf{X} . This is the general approach followed in the section “[Analysis of Variance](#)” on page 57 in Chapter 3, “[Introduction to Statistical Modeling with SAS/STAT Software](#).” Depending on the statistical question at hand, the projections can be formulated based on estimable functions, with different types of estimable functions giving rise to different types of sums of squares. Note that not all sum of squares analyses necessarily correspond to additive decompositions. For example, the Type III sums of squares often test hypotheses about the model that are more meaningful than those corresponding to the Type I sums of squares. But while the Type I sums of squares additively decompose the sum of squares due to all model contributions, the Type III sums of squares do not necessarily add up to any useful quantity. The four types of estimable functions in SAS/STAT software, their interpretation, and their construction are discussed in Chapter 15, “[The Four Types of Estimable Functions](#).” The application of sum of squares analyses is not necessarily limited to models with classification effects (factors). The methodology also applies to linear regression models that contain only continuous regressor variables.

An even broader sense of the term “analysis of variance” pertains to statistical models that contain classification effects (factors), and in particular, to models that contain *only* classification effects. Any statistical approach that measures features of such a model and can be used for inference is called a *general analysis of variance*. Thus the procedures for general analysis of variance in SAS/STAT are considered to be those that can fit statistical models containing factors, whether the data are experimental or observational. Some procedures for general analysis of variance have a statistical estimation principle that gives rise to a sum of squares analysis as discussed previously; others express a factor’s contribution to the model fit in some other form. Note that this view of analysis of variance includes, for example, maximum likelihood estimation in generalized linear models with the GENMOD procedure, restricted maximum likelihood estimation in linear mixed models with the MIXED procedure, the estimation of variance components with the VARCOMP procedure, the comparison of means of groups with the TTEST procedure, and the nonparametric analysis of rank scores with the NPAR1WAY procedure, and so on.

In summary, analysis of variance in the contemporary sense of statistical modeling and analysis is more aptly described as *analysis of variation*, the study of the influences on the variation of a phenomenon. This can take, for example, the following forms:

- an analysis of variance table based on sums of squares followed by more specific inquiries into the relationship among factors and their levels
- a deviance decomposition in a generalized linear model
- a series of Type III tests followed by comparisons of least squares means in a mixed model

Procedures That Perform Sum of Squares Analysis of Variance

The flagship procedure in SAS/STAT software for linear modeling with sum of squares analysis techniques is the GLM procedure. It handles most standard analysis of variance problems. The following list provides descriptions of PROC GLM and other procedures that are used for more specialized situations:

ANOVA	performs analysis of variance, multivariate analysis of variance, and repeated measures analysis of variance for <i>balanced</i> designs. PROC ANOVA also performs multiple comparison tests on arithmetic means.
GLM	performs analysis of variance, regression, analysis of covariance, repeated measures analysis, and multivariate analysis of variance. PROC GLM produces several diagnostic measures, performs tests for random effects, provides contrasts and estimates for customized hypothesis tests, provides tests for means adjusted for covariates, and performs multiple-comparison tests on both arithmetic and adjusted means.
LATTICE	computes the analysis of variance and analysis of simple covariance for data from an experiment with a lattice design. PROC LATTICE analyzes balanced square lattices, partially balanced square lattices, and some rectangular lattices.
MIXED	performs mixed model analysis of variance and repeated measures analysis of variance via covariance structure modeling. When you choose one of the method-of-moment estimation techniques, the MIXED procedure produces an analysis of variance table with sums of squares, mean squares, and expected mean squares. PROC MIXED constructs statistical tests and intervals, enables customized contrasts and estimates, and computes empirical Bayes predictions.
NESTED	performs analysis of variance and analysis of covariance for purely nested random models.
ORTHOREG	performs regression by using the Gentleman-Givens computational method. For ill-conditioned data, PROC ORTHOREG can produce more accurate parameter estimates than other procedures, such as PROC GLM. See Chapter 84, “ The ORTHOREG Procedure ,” for more information.
VARCOMP	estimates variance components for random or mixed models. If you choose the METHOD=TYPE1 or METHOD=GRR option, the VARCOMP procedure produces an analysis of variance table with sums of squares that correspond to the random effects in your models.
TRANSREG	fits univariate and multivariate linear models, optionally with spline and other nonlinear transformations. Models include ordinary regression and ANOVA, multiple and multivariate regression, metric and nonmetric conjoint analysis, metric and nonmetric vector and ideal point preference mapping, redundancy analysis, canonical correlation, and response surface regression. See Chapter 117, “ The TRANSREG Procedure ,” for more information.

Procedures That Perform General Analysis of Variance

Many procedures in SAS/STAT enable you to incorporate classification effects into your model and to perform statistical inferences for experimental factors and their interactions. These procedures do not necessarily rely on sums of squares decompositions to perform these inferences. Examples of such procedures are the CATMOD, GENMOD, GLIMMIX, LOGISTIC, NPAR1WAY, and TTEST procedures. In fact, any one of the more than two dozen SAS/STAT modeling procedures that include a CLASS statement can be said to perform analysis of variance in this general sense. For more information about individual procedures, refer to their corresponding chapters in this documentation.

The following section discusses procedures in SAS/STAT that compute analysis of variance in models with classification factors in the narrow sense—that is, they produce analysis of variance tables and form F tests based on sums of squares, mean squares, and expected mean squares.

The subsequent sections discuss procedures that perform statistical inference in models with classification effects in the broader sense.

The following section also presents an overview of some of the fundamental features of analysis of variance. Subsequent sections describe how this analysis is performed with procedures in SAS/STAT software. For more detail, see the chapters for the individual procedures. Additional sources are described in the section [References](#).

Statistical Details for Analysis of Variance

From Sums of Squares to Linear Hypotheses

Analysis of variance (ANOVA) is a technique for analyzing data in which one or more *response* (or *dependent* or simply Y) variables are measured under various conditions identified by one or more classification variables. The combinations of levels for the classification variables form the cells of the design for the data. This design can be the result of a controlled experiment or the result of an observational study in which you observe factors and factor level combinations in an uncontrolled environment. For example, an experiment might measure weight change (the dependent variable) for men and women who participated in three different weight-loss programs. The six cells of the design are formed by the six combinations of gender (men, women) and program (A, B, C).

In an analysis of variance, the variation in the response is separated into variation attributable to differences between the classification variables and variation attributable to random error. An analysis of variance constructs tests to determine the significance of the classification effects. A typical goal in such an analysis is to compare means of the response variable for various combinations of the classification variables.

The least squares principle is central to computing sums of squares in analysis of variance models. Suppose that you are fitting the linear model $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$ and that the error terms satisfy the usual assumptions (uncorrelated, zero mean, homogeneous variance). Further, suppose that \mathbf{X} is partitioned according to several model effects, $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2 \ \cdots \ \mathbf{X}_k]$. If $\hat{\boldsymbol{\beta}}$ denotes the ordinary least squares solution for this model, then the sum of squares attributable to the overall model can be written as

$$\text{SSM} = \hat{\boldsymbol{\beta}}' \mathbf{X}' \mathbf{Y} = \mathbf{Y}' \mathbf{H} \mathbf{Y}$$

where \mathbf{H} is the “hat” matrix $\mathbf{H} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$. (This model sum of squares is not yet corrected for the presence of an explicit or implied intercept. This adjustment would consist of subtracting $n\bar{Y}^2$ from SSM.) Because of the properties of the hat matrix \mathbf{H} , you can write $\mathbf{X}' = \mathbf{X}'\mathbf{H}$ and $\mathbf{H}\mathbf{X} = \mathbf{X}$. The (uncorrected) model sum of squares thus can also be written as

$$\text{SSM} = \hat{\boldsymbol{\beta}}' (\mathbf{X}'\mathbf{X}) \hat{\boldsymbol{\beta}}$$

This step is significant, because it demonstrates that sums of squares can be identified with quadratic functions in the least squares coefficients. The generalization of this idea is to do the following:

- consider hypotheses of interest in an analysis of variance model
- express the hypotheses in terms of linear estimable functions of the parameters
- compute the sums of squares associated with the estimable function
- construct statistical tests based on the sums of squares

Decomposing a model sum of squares into sequential, additive components, testing the significance of experimental factors, comparing factor levels, and performing other statistical inferences fall within this generalization. Suppose that $\mathbf{L}\boldsymbol{\beta}$ is an estimable function (see the section “[Estimable Functions](#)” on page 59 in Chapter 3, “[Introduction to Statistical Modeling with SAS/STAT Software](#),” and Chapter 15, “[The Four Types of Estimable Functions](#),” for details). The sum of squares associated with the hypothesis $H: \mathbf{L}\boldsymbol{\beta} = \mathbf{0}$ is

$$SS(H) = SS(\mathbf{L}\boldsymbol{\beta} = \mathbf{0}) = \hat{\boldsymbol{\beta}}' \mathbf{L}' (\mathbf{L}(\mathbf{X}'\mathbf{X})^{-1} \mathbf{L}')^{-1} \mathbf{L} \hat{\boldsymbol{\beta}}$$

One application would be to form sums of squares associated with the different components of \mathbf{X} . For example, you can form a matrix \mathbf{L}_2 matrix such that $\mathbf{L}_2\boldsymbol{\beta} = \mathbf{0}$ tests the effect of adding the columns for \mathbf{X}_2 to an empty model or to test the effect of adding \mathbf{X}_2 to a model that already contains \mathbf{X}_1 .

These sums of squares can also be expressed as the difference between two residual sums of squares, since $\mathbf{L}\boldsymbol{\beta} = \mathbf{0}$ can be thought of as a (linear) restriction on the parameter estimates in the model:

$$SS(H) = SSR(\text{constrained model}) - SSR(\text{full model})$$

If, in addition to the usual assumptions mentioned previously, the model errors are assumed to be normally distributed, then $SS(H)$ follows a distribution that is proportional to a chi-square distribution. This fact, and the independence of $SS(H)$ from the residual sum of squares, enables you to construct F tests based on sums of squares in least squares models.

The extension of sum of squares analysis of variance to general analysis of variance for classification effects depends on the fact that the distributional properties of quadratic forms in normal random variables are well understood. It is not necessary to first formulate a sum of squares to arrive at an exact or even approximate F test. The generalization of the expression for $SS(H)$ is to form test statistics based on quadratic forms

$$\hat{\boldsymbol{\beta}}' \mathbf{L}' \text{Var}[\mathbf{L}\hat{\boldsymbol{\beta}}]^{-1} \mathbf{L} \hat{\boldsymbol{\beta}}$$

that follow a chi-square distribution if $\hat{\boldsymbol{\beta}}$ is normally distributed.

Tests of Effects Based on Expected Mean Squares

Statistical tests in analysis of variance models can be constructed by comparing independent mean squares. To test a particular null hypothesis, you compute the ratio of two mean squares that have the same expected value under that hypothesis; if the ratio is much larger than 1, then that constitutes significant evidence against the null. In particular, in an analysis of variance model with fixed effects only, the expected value of

each mean square has two components: quadratic functions of fixed parameters and random variation. For example, for a fixed effect called A, the expected value of its mean square is

$$E[MS(A)] = Q(\boldsymbol{\beta}) + \sigma^2$$

where σ^2 is the common variance of the ϵ_i .

Under the null hypothesis of no A effect, the fixed portion $Q(\boldsymbol{\beta})$ of the expected mean square is zero. This mean square is then compared to another mean square—say, $MS(E)$ —that is independent of the first and has the expected value σ^2 . The ratio of the two mean squares

$$F = \frac{MS(A)}{MS(E)}$$

has an F distribution under the null hypothesis.

When the null hypothesis is false, the numerator term has a larger expected value, but the expected value of the denominator remains the same. Thus, large F values lead to rejection of the null hypothesis. The probability of getting an F value at least as large as the one observed given that the null hypothesis is true is called the *significance probability value* (or the p -value). A p -value of less than 0.05, for example, indicates that data with *no* A effect will yield F values as large as the one observed less than 5% of the time. This is usually considered moderate evidence that there *is* a real A effect. Smaller p -values constitute even stronger evidence. Larger p -values indicate that the effect of interest is less than random noise. In this case, you can conclude either that there is no effect at all or that you do not have enough data to detect the differences being tested.

The actual pattern in expected mean squares of terms related to fixed quantities ($Q(\boldsymbol{\beta})$) and functions of variance components depends on which terms in your model are fixed effects and which terms are random effects. This has bearing on how F statistics can be constructed. In some instances, exact tests are not available, such as when a linear combination of expected mean squares is necessary to form a proper denominator for an F test and a Satterthwaite approximation is used to determine the degrees of freedom of the approximation. The GLM and MIXED procedures can generate tables of expected mean squares and compute degrees of freedom by Satterthwaite's method. The MIXED and GLIMMIX procedures can apply Satterthwaite approximations and other degrees-of-freedom computations more widely than in analysis of variance models. See the section “Fixed, Random, and Mixed Models” on page 31 in Chapter 3, “Introduction to Statistical Modeling with SAS/STAT Software,” for a discussion of fixed versus random effects in statistical models.

Analysis of Variance for Fixed-Effect Models

PROC GLM for General Linear Models

The GLM procedure is the flagship tool for classical analysis of variance in SAS/STAT software. It performs analysis of variance by using least squares regression to fit general linear models. Among the statistical methods available in PROC GLM are regression, analysis of variance, analysis of covariance, multivariate analysis of variance, repeated measures analysis, and partial correlation analysis.

While PROC GLM can handle most common analysis of variance problems, other procedures are more efficient or have more features than PROC GLM for certain specialized analyses, or they can handle

specialized models that PROC GLM cannot. Much of the rest of this chapter is concerned with comparing PROC GLM to other procedures.

PROC ANOVA for Balanced Designs

When you design an experiment, you choose how many experimental units to assign to each combination of levels (or cells) in the classification. In order to achieve good statistical properties and simplify the computations, you typically attempt to assign the same number of units to every cell in the design. Such designs are called *balanced designs*.

In SAS/STAT software, you can use the ANOVA procedure to perform analysis of variance for balanced data. The ANOVA procedure performs computations for analysis of variance that assume the balanced nature of the data. These computations are simpler and more efficient than the corresponding general computations performed by PROC GLM. Note that PROC ANOVA can be applied to certain designs that are not balanced in the strict sense of equal numbers of observations for all cells. These additional designs include all one-way models, regardless of how unbalanced the cell counts are, as well as Latin squares, which do not have data in all cells. In general, however, the ANOVA procedure is recommended only for balanced data. **If you use ANOVA to analyze a design that is not balanced, you must assume responsibility for the validity of the output.** You are responsible for recognizing incorrect results, which might include negative values reported for the sums of squares. If you are not certain that your data fit into a balanced design, then you probably need the framework of general linear models in the GLM procedure.

Comparing Group Means

The F test for a classification factor that has more than two levels tells you whether the level effects are significantly different from each other, but it does not tell you which levels differ from which other levels.

If the level comparisons are expressed through differences of the arithmetic cell means, you can use the MEANS statement in the GLM and ANOVA procedures for comparison. If arithmetic means are not appropriate for comparison, for example, because your data are unbalanced or means need to be adjusted for other model effects, then you can use the LSMEANS statement in the GLIMMIX, GLM, and MIXED procedures for level comparisons.

If you have specific comparisons in mind, you can use the CONTRAST statement in these procedures to make these comparisons. However, if you make many comparisons that use some given significance level (0.05, for example), you are more likely to make a type 1 error (incorrectly rejecting a hypothesis that the means are equal) simply because you have more chances to make the error.

Multiple-comparison methods give you more detailed information about the differences among the means and enable you to control error rates for a multitude of comparisons. A variety of multiple-comparison methods are available with the MEANS statement in both the ANOVA and GLM procedures, as well as the LSMEANS statement in the GLIMMIX, GLM, and MIXED procedures. These are described in detail in the section “[Multiple Comparisons](#)” on page 3615 in Chapter 46, “[The GLM Procedure](#),” and in Chapter 45, “[The GLIMMIX Procedure](#),” and Chapter 77, “[The MIXED Procedure](#).”

PROC TTEST for Comparing Two Groups

If you want to perform an analysis of variance and have only one classification variable with two levels, you can use PROC TTEST. In this special case, the results generated by PROC TTEST are equivalent to the results generated by PROC ANOVA or PROC GLM.

You can use PROC TTEST with balanced or unbalanced groups. In addition to the test assuming equal variances, PROC TTEST also performs a Satterthwaite test assuming unequal variances.

The TTEST procedure also performs equivalence tests, computes confidence limits, and supports both normal and lognormal data. If you have an AB/BA crossover design with no carryover effects, then you can use the TTEST procedure to analyze the treatment and period effects.

The PROC NPAR1WAY procedure performs nonparametric analogues to *t* tests. See Chapter 16, “[Introduction to Nonparametric Analysis](#),” for an overview and Chapter 83, “[The NPAR1WAY Procedure](#),” for details on PROC NPAR1WAY.

Analysis of Variance for Categorical Data and Generalized Linear Models

A *categorical variable* is defined as one that can assume only a limited number of values. For example, a person’s gender is a categorical variable that can assume one of two values. Variables with levels that simply name a group are said to be measured on a *nominal scale*. Categorical variables can also be measured using an *ordinal scale*, which means that the levels of the variable are ordered in some way. For example, responses to an opinion poll are usually measured on an ordinal scale, with levels ranging from “strongly disagree” to “no opinion” to “strongly agree.”

For two categorical variables, one measured on an ordinal scale and one measured on a nominal scale, you can assign scores to the levels of the ordinal variable and test whether the mean scores for the different levels of the nominal variable are significantly different. This process is analogous to performing an analysis of variance on continuous data, which can be performed by PROC CATMOD. If there are *n* nominal variables, rather than 1, then PROC CATMOD can perform an *n*-way analysis of variance of the mean scores.

For two categorical variables measured on a nominal scale, you can test whether the distribution of the first variable is significantly different for the levels of the second variable. This process is an analysis of variance of proportions, rather than means, and can be performed by PROC CATMOD. The corresponding *n*-way analysis of variance can also be performed by PROC CATMOD.

See Chapter 8, “[Introduction to Categorical Data Analysis Procedures](#),” and Chapter 32, “[The CATMOD Procedure](#),” for more information.

The GENMOD procedure uses maximum likelihood estimation to fit generalized linear models. This family includes models for categorical data such as logistic, probit, and complementary log-log regression for binomial data and Poisson regression for count data, as well as continuous models such as ordinary linear regression, gamma, and inverse Gaussian regression models. PROC GENMOD performs analysis of variance through likelihood ratio and Wald tests of fixed effects in generalized linear models, and provides contrasts and estimates for customized hypothesis tests. It performs analysis of repeated measures data with generalized estimating equation (GEE) methods.

See Chapter 8, “Introduction to Categorical Data Analysis Procedures,” and Chapter 44, “The GENMOD Procedure,” for more information.

Nonparametric Analysis of Variance

Analysis of variance is sensitive to the distribution of the error term. If the error term is not normally distributed, the statistics based on normality can be misleading. The traditional test statistics are called *parametric tests* because they depend on the specification of a certain probability distribution except for a set of free parameters. Parametric tests are said to depend on distributional assumptions. Nonparametric methods perform the tests without making any strict distributional assumptions. Even if the data are distributed normally, nonparametric methods are often almost as powerful as parametric methods.

Most nonparametric methods are based on taking the ranks of a variable and analyzing these ranks (or transformations of them) instead of the original values. The NPAR1WAY procedure performs a nonparametric one-way analysis of variance. Other nonparametric tests can be performed by taking ranks of the data (using the RANK procedure) and using a regular parametric procedure (such as GLM or ANOVA) to perform the analysis. Some of these techniques are outlined in the description of PROC RANK in *SAS Language Reference: Concepts* and in Conover and Iman (1981).

Constructing Analysis of Variance Designs

Analysis of variance is most often used for data from designed experiments. You can use the PLAN procedure to construct designs for many experiments. For example, PROC PLAN constructs designs for completely randomized experiments, randomized blocks, Latin squares, factorial experiments, certain balanced incomplete block designs, and balanced crossover designs.

Randomization, or randomly assigning experimental units to cells in a design and to treatments within a cell, is another important aspect of experimental design. For either a new or an existing design, you can use PROC PLAN to randomize the experimental plan.

Additional features for design of experiments are available in SAS/QC software. The FACTEX and OPTEX procedures can construct a wide variety of designs, including factorials, fractional factorials, and D-optimal or A-optimal designs. These procedures, as well as the ADX Interface, provide features for randomizing and replicating designs; saving the design in an output data set; and interactively changing the design by changing its size, use of blocking, or the search strategies used. For more information, see the *SAS/QC User's Guide*.

For More Information

Analysis of variance was pioneered by Fisher (1925). For a general introduction to analysis of variance, see an intermediate statistical methods textbook such as Steel and Torrie (1980); Snedecor and Cochran (1980); Milliken and Johnson (1984); Mendenhall (1968); John (1971); Ott (1977); Kirk (1968). A classic source is Scheffé (1959). Freund, Littell, and Spector (1991) bring together a treatment of these statistical methods and SAS/STAT software procedures. Schlotzhauer and Littell (1997) cover how to perform *t* tests and one-way

analysis of variance with SAS/STAT procedures. Texts on linear models include Searle (1971); Graybill (1976); Hocking (1985). Kennedy and Gentle (1980) survey the computing aspects. Other references appear in the reference section.

References

- Conover, W. J., and Iman, R. L. (1981). "Rank Transformations as a Bridge between Parametric and Nonparametric Statistics." *American Statistician* 35:124–129.
- Fisher, R. A. (1925). *Statistical Methods for Research Workers*. Edinburgh: Oliver & Boyd.
- Freund, R. J., Littell, R. C., and Spector, P. C. (1991). *SAS System for Linear Models*. Cary, NC: SAS Institute Inc.
- Graybill, F. A. (1976). *Theory and Applications of the Linear Model*. North Scituate, MA: Duxbury Press.
- Hocking, R. R. (1985). *The Analysis of Linear Models*. Monterey, CA: Brooks/Cole.
- John, P. W. M. (1971). *Statistical Design and Analysis of Experiments*. New York: Macmillan.
- Kennedy, W. J., Jr., and Gentle, J. E. (1980). *Statistical Computing*. New York: Marcel Dekker.
- Kirk, R. E. (1968). *Experimental Design: Procedures for the Behavioral Sciences*. Monterey, CA: Brooks/Cole.
- Mendenhall, W. (1968). *Introduction to Linear Models and the Design and Analysis of Experiments*. Belmont, CA: Duxbury Press.
- Milliken, G. A., and Johnson, D. E. (1984). *Designed Experiments*. Vol. 1 of Analysis of Messy Data. Belmont, CA: Lifetime Learning Publications.
- Ott, L. (1977). *Introduction to Statistical Methods and Data Analysis*. 2nd ed. Belmont, CA: Duxbury Press.
- Scheffé, H. (1959). *The Analysis of Variance*. New York: John Wiley & Sons.
- Schlotzhauer, S. D., and Littell, R. C. (1997). *SAS System for Elementary Statistical Analysis*. 2nd ed. Cary, NC: SAS Institute Inc.
- Searle, S. R. (1971). *Linear Models*. New York: John Wiley & Sons.
- Snedecor, G. W., and Cochran, W. G. (1980). *Statistical Methods*. 7th ed. Ames: Iowa State University Press.
- Steel, R. G. D., and Torrie, J. H. (1980). *Principles and Procedures of Statistics*. 2nd ed. New York: McGraw-Hill.

Index

- analysis of variance
 - Introduction to ANOVA Procedures, 105
- estimable function
 - Introduction to ANOVA Procedures, 106, 109
- Introduction to ANOVA
 - SAS/STAT procedures, 105
- Introduction to ANOVA Procedures
 - analysis of covariance, 107
 - analysis of ranks, 106, 113
 - analysis of variance, 107
 - balanced design, 111
 - classification effect, 106–108
 - constructing designs, 113
 - controlled experiment, 108
 - covariance structure modeling, 107
 - definition, 105
 - design matrix, 106
 - empirical Bayes predictions, 107
 - estimable function, 106, 109
 - exact test, 110
 - expected mean squares, 110
 - experimental data, 106, 108
 - F-test based on sum of squares, 109
 - fixed effect, 110
 - general analysis of variance model, 106
 - general ANOVA procedures, 107
 - group comparisons, 111
 - hat matrix, 108
 - hypothesis sum of squares, 109
 - lattice design, 107
 - least squares, 108
 - linear model, 106
 - mean squares, 109
 - mean squares, expected, 110
 - method of moments, 107
 - model sum of squares, 108
 - multiple comparisons, 111
 - multivariate analysis of variance, 107
 - nested model, 107
 - nonlinear transformation, 107
 - nonparametric analysis, 106, 113
 - observational data, 106, 108
 - p-value, 110
 - projection, 106
 - random effect, 110
 - repeated measures, 107
 - residual sum of squares, 109
 - Satterthwaite approximation, 110
 - spline transformation, 107
 - sum of squares decomposition, 105, 107
 - Type I sum of squares, 106
 - Type III sum of squares, 106
 - variance components, 107
- least squares
 - Introduction to ANOVA Procedures, 108
- matrix
 - design (Introduction to ANOVA Procedures), 106
 - hat (Introduction to ANOVA Procedures), 108
 - projection (Introduction to ANOVA Procedures), 106
- mixed model
 - linear (Introduction to ANOVA Procedures), 107
- sum of squares
 - decomposition (Introduction to ANOVA Procedures), 105, 107
 - F-test (Introduction to ANOVA Procedures), 109
 - for linear hypothesis (Introduction to ANOVA Procedures), 109
 - model (Introduction to ANOVA Procedures), 108
 - residual (Introduction to ANOVA Procedures), 109
 - Type I (Introduction to ANOVA Procedures), 106
 - Type III (Introduction to ANOVA Procedures), 106