

User's Guide III

Standard Deviations, Variances, Normality Tests,
Survival Analysis, Correlations, Regression, Design of
Experiments, and Tools, Helps, and Aids

PASS
Power Analysis
and
Sample Size
System

Published by
NCSS
Dr. Jerry L. Hintze
Kaysville, Utah

PASS User's Guide III

Copyright © 2008
Dr. Jerry L. Hintze
Kaysville, Utah 84037

All Rights Reserved

Direct inquiries to:

NCSS
329 North 1000 East
Kaysville, Utah 84037
Phone (801) 546-0445
Fax (801) 546-3907
Email: support@ncss.com

PASS is a trademark of Dr. Jerry L. Hintze.

Warning:

This software and manual are both protected by U.S. Copyright Law (Title 17 United States Code). Unauthorized reproduction and/or sales may result in imprisonment of up to one year and fines of up to \$10,000 (17 USC 506). Copyright infringers may also be subject to civil liability.

About This Manual

Congratulations on your purchase of the *PASS* package! *PASS* offers:

- Easy parameter entry.
- A comprehensive list of power analysis routines that are accurate and verified, yet are quick and easy to learn and use.
- Straightforward procedures for creating paper printouts and file copies of both the numerical and graphical reports.

Our goal is that with the help of these user's guides, you will be up and running on *PASS* quickly. After reading the quick start manual (at the front of User's Guide I) you will only need to refer to the chapters corresponding to the procedures you want to use. The discussion of each procedure includes one or more tutorials that will take you step-by-step through the tasks necessary to run the procedure.

I believe you will find that these user's guides provides a quick, easy, efficient, and effective way for first-time *PASS* users to get up and running.

I look forward to any suggestions you have to improve the usefulness of this manual and/or the *PASS* system. Meanwhile, good computing!

Jerry Hintze, Author

PASS License Agreement

Important: The enclosed Power Analysis and Sample Size software program (PASS) is licensed by NCSS to customers for their use only on the terms set forth below. Your purchase and use of the PASS system indicates your acceptance of these terms.

1. **LICENSE.** NCSS hereby agrees to grant you a non-exclusive license to use the accompanying PASS program subject to the terms and restrictions set forth in this License Agreement.
2. **COPYRIGHT.** PASS and its documentation are copyrighted. You may not copy or otherwise reproduce any part of PASS or its documentation, except that you may load PASS into a computer as an essential step in executing it on the computer and make backup copies for your use on the same computer.
3. **BACKUP POLICY.** PASS may be backed up by you for your use on the same machine for which PASS was purchased.
4. **RESTRICTIONS ON USE AND TRANSFER.** The original and any backup copies of PASS and its documentation are to be used only in connection with a single user. This user may load PASS onto several machines for his/her convenience (such as a desktop and laptop computer), but only for use by the licensee. You may physically transfer PASS from one computer to another, provided that PASS is used in connection with only one user. You may not distribute copies of PASS or its documentation to others. You may transfer this license together with the original and all backup copies of PASS and its documentation, provided that the transferee agrees to be bound by the terms of this License Agreement. PASS licenses may not be transferred more frequently than once in twelve months. Neither PASS nor its documentation may be modified or translated without written permission from NCSS.
You may not use, copy, modify, or transfer PASS, or any copy, modification, or merged portion, in whole or in part, except as expressly provided for in this license.
5. **NO WARRANTY OF PERFORMANCE.** NCSS does not and cannot warrant the performance or results that may be obtained by using PASS. Accordingly, PASS and its documentation are licensed "as is" without warranty as to their performance, merchantability, or fitness for any particular purpose. The entire risk as to the results and performance of PASS is assumed by you. Should PASS prove defective, you (and not NCSS nor its dealer) assume the entire cost of all necessary servicing, repair, or correction.
6. **LIMITED WARRANTY ON CD.** To the original licensee only, NCSS warrants the medium on which PASS is recorded to be free from defects in materials and faulty workmanship under normal use and service for a period of ninety days from the date PASS is delivered. If, during this ninety-day period, a defect in a CD should occur, the CD may be returned to NCSS at its address, or to the dealer from which PASS was purchased, and NCSS will replace the CD without charge to you, provided that you have sent a copy of your receipt for PASS. Your sole and exclusive remedy in the event of a defect is expressly limited to the replacement of the CD as provided above.
Any implied warranties of merchantability and fitness for a particular purpose are limited in duration to a period of ninety (90) days from the date of delivery. If the failure of a CD has resulted from accident, abuse, or misapplication of the CD, NCSS shall have no responsibility to replace the CD under the terms of this limited warranty. This limited warranty gives you specific legal rights, and you may also have other rights which vary from state to state.
7. **LIMITATION OF LIABILITY.** Neither NCSS nor anyone else who has been involved in the creation, production, or delivery of PASS shall be liable for any direct, incidental, or consequential damages, such as, but not limited to, loss of anticipated profits or benefits, resulting from the use of PASS or arising out of any breach of any warranty. Some states do not allow the exclusion or limitation of direct, incidental, or consequential damages, so the above limitation may not apply to you.
8. **TERM.** The license is effective until terminated. You may terminate it at any time by destroying PASS and documentation together with all copies, modifications, and merged portions in any form. It will also terminate if you fail to comply with any term or condition of this License Agreement. You agree upon such termination to destroy PASS and documentation together with all copies, modifications, and merged portions in any form.
9. **YOUR USE OF PASS ACKNOWLEDGES** that you have read this customer license agreement and agree to its terms. You further agree that the license agreement is the complete and exclusive statement of the agreement between us and supersedes any proposal or prior agreement, oral or written, and any other communications between us relating to the subject matter of this agreement.

Dr. Jerry L. Hintze & NCSS, Kaysville, Utah

Preface

PASS (Power Analysis and Sample Size) is an advanced, easy-to-use statistical analysis software package. The system was designed and written by Dr. Jerry L. Hintze over the last Seventeen years. Dr. Hintze drew upon his experience both in teaching statistics at the university level and in various types of statistical consulting.

The present version, written for 32-bit versions of Microsoft Windows (Vista, XP, NT, ME, 2000, 98, etc.) computer systems, is the result of several iterations. Experience over the years with several different types of users has helped the program evolve into its present form.

NCSS maintains a website at www.ncss.com where we make the latest edition of **PASS** available for free downloading. The software is password protected, so only users with valid serial numbers may use this downloaded edition. We hope that you will download the latest edition routinely and thus avoid any bugs that have been corrected since you purchased your copy.

We believe **PASS** to be an accurate, exciting, easy-to-use program. If you find any portion which you feel needs to be changed, please let us know. Also, we openly welcome suggestions for additions and enhancements.

Verification

All calculations used in this program have been extensively tested and verified. First, they have been verified against the original journal article or textbook that contained the formulas. Second, they have been verified against second and third sources when these exist.

User's Guide III

Table of Contents

Standard Deviations

- 640 Confidence Intervals for One Standard Deviation using Standard Deviation
- 641 Confidence Intervals for One Standard Deviation with Tolerance Probability
- 642 Confidence Intervals for One Standard Deviation using Relative Error

Variances

One Variance

- 650 Inequality Tests for One Variance
- 651 Confidence Intervals for One Variance using Variance
- 652 Confidence Intervals for One Variance with Tolerance Probability
- 653 Confidence Intervals for One Variance using Relative Error

Two Variances

- 655 Inequality Tests for Two Variances
- 656 Confidence Intervals for the Ratio of Two Variances using Variances
- 657 Confidence Intervals for the Ratio of Two Variances using Relative Error

Normality Tests

- 670 Normality Tests (Simulation)

Survival Analysis

- 700 Logrank Tests (Freedman)
- 705 Logrank Tests (Lachin and Foulkes)
- 706 Logrank Tests for Non-Inferiority
- 710 Group-Sequential Logrank Tests
- 715 Logrank Tests (Lakatos)

Correlations

- 800 Inequality Tests for One Correlation
- 801 Confidence Intervals for One Correlation
- 805 Inequality Tests for Two Correlations
- 810 Inequality Tests for Intraclass Correlation

- 811 Kappa Test for Agreement Between Two Raters
- 815 Inequality Tests for One Coefficient Alpha
- 820 Inequality Tests for Two Coefficient Alphas

Regression

- 850 Cox Regression
- 855 Linear Regression
- 856 Confidence Intervals for Linear Regression Slope
- 860 Logistic Regression
- 865 Multiple Regression
- 870 Poisson Regression

Design of Experiments

- 880 Randomization Lists
- 881 Two-Level Designs
- 882 Fractional Factorial Designs
- 883 Balanced Incomplete Block Designs
- 884 Latin Square Designs
- 885 Response Surface Designs
- 886 Screening Designs
- 887 Taguchi Designs
- 888 D-Optimal Designs
- 889 Design Generator

Tools, Helps, and Aids

- 900 Chi-Square Effect Size Estimator
- 905 Standard Deviation Estimator
- 910 Odds Ratio and Proportions Estimator
- 915 Probability Calculator
- 920 Data Simulator
- 925 The Spreadsheet
- 930 Macros

References and Index

- References
- Index

User's Guide I

Table of Contents

Quick Start

1	Installation
2	Running PASS
3	The PASS Home Window
4	The Procedure Window
5	The Output Window
6	The Map (Quick Launch) Window
7	Introduction to Power Analysis
8	Proportions
9	Means

Quick Start Index

Proportions

One Proportion

100	Inequality Tests
105	Non-Inferiority & Superiority Tests
110	Equivalence Tests
115	Confidence Intervals
120	Single-Stage Phase II Clinical Trials
125	Two-Stage Phase II Clinical Trials
130	Three-Stage Phase II Clinical Trials
135	Post-Marketing Surveillance

Two Correlated Proportions

150	Inequality Tests (McNemar Test)
155	Inequality Tests (Matched Case-Control Design)
160	Non-Inferiority Tests
165	Equivalence Tests

Two Independent Proportions

200	Inequality Tests
201	Inequality Tests (Repeated Measures Design)
205	Inequality Tests (Offset Null Hypothesis)
210	Non-Inferiority & Superiority Tests
215	Equivalence Tests
216	Confidence Intervals
220	Group-Sequential Tests
225	Inequality Tests (Stratified Design – Cochran-Mantel-Haenszel Test)

Two Independent Proportions in a Cluster-Randomized Design

230	Inequality Tests
235	Non-Inferiority & Superiority Tests
240	Equivalence Tests

Many Proportions (Contingency Tables)

250	Chi-Square Tests
255	Cochran-Armitage Test for Trend in Proportions

ROC Curves

260	Inequality Tests for One ROC Curve
265	Inequality Tests for Two ROC Curves

References and Index

References
Index

User's Guide II

Table of Contents

Means

One Mean

- 400 Inequality Tests (One-Sample or Paired T-Test)
- 405 Inequality Tests (Exponential Data)
- 410 Inequality Tests (Simulation)
- 415 Non-Inferiority & Superiority Tests
- 420 Confidence Intervals
- 421 Confidence Intervals with Tolerance Probability

Two Independent Means

- 430 Inequality Tests using Differences (Two-Sample T-Test)
- 431 Inequality Tests (Repeated Measures Design)
- 435 Inequality Tests (Exponential Data)
- 440 Inequality Tests (Simulation)
- 445 Inequality Tests using Ratios (Two-Sample T-Test)
- 450 Non-Inferiority & Superiority Tests using Differences
- 455 Non-Inferiority & Superiority Tests using Ratios
- 460 Equivalence Tests using Differences
- 465 Equivalence Tests (Simulation)
- 470 Equivalence Tests using Ratios
- 471 Confidence Intervals for the Difference
- 472 Confidence Intervals for the Difference with Tolerance Probability
- 475 Group-Sequential Tests
- 480 Inequality Tests (Cluster-Randomized Design)

Two Correlated (Paired) Means

- 490 Inequality Tests (Simulation)
- 495 Equivalence Tests (Simulation)
- 496 Confidence Intervals
- 497 Confidence Intervals with Tolerance Probability

Two Independent Means in a 2x2 Cross-Over Design

- 500 Inequality Tests using Differences
- 505 Inequality Tests using Ratios
- 510 Non-Inferiority & Superiority Tests using Differences

- 515 Non-Inferiority & Superiority Tests using Ratios
- 520 Equivalence Tests using Differences
- 525 Equivalence Tests using Ratios

Two Independent Means in a Higher-Order Cross-Over Design

- 530 Non-Inferiority & Superiority Tests using Differences
- 535 Non-Inferiority & Superiority Tests using Ratios
- 540 Equivalence Tests using Differences
- 545 Equivalence Tests using Ratios

Many Means (ANOVA)

- 550 One-Way Analysis of Variance
- 555 One-Way Analysis of Variance (Simulation)
- 560 Fixed Effects Analysis of Variance
- 565 Randomized Block Analysis of Variance
- 570 Repeated Measures Analysis of Variance

Mixed Models

- 571 Mixed Models

Multiple Comparisons

- 575 Multiple Comparisons
- 580 Pair-Wise Multiple Comparisons (Simulation)
- 585 Multiple Comparisons of Treatments vs. a Control (Simulation)
- 590 Multiple Contrasts (Simulation)

Multivariate Means

- 600 Hotelling's T2
- 605 Multivariate Analysis of Variance (MANOVA)

Microarrays

- 610 One-Sample or Paired T-Test
- 615 Two-Sample T-Test

References and Index

- References
- Index

Chapter 640

Confidence Intervals for One Standard Deviation Using Standard Deviation

Introduction

This routine calculates the sample size necessary to achieve a specified interval width or distance from the standard deviation to the confidence limit at a stated confidence level for a confidence interval about the standard deviation when the underlying data distribution is normal.

Caution: This procedure assumes that the standard deviation of the future sample will be the same as the standard deviation that is specified. If the standard deviation to be used in the procedure is estimated from a previous sample or represents the population standard deviation, the Confidence Intervals for One Standard Deviation with Tolerance Probability procedure should be considered. That procedure controls the probability that the width or distance from the standard deviation to the confidence limit will be less than or equal to the value specified. The Confidence Intervals for One Standard Deviation using Relative Error controls the width or distance from the standard deviation to the limit by controlling the distance as a percent of the true standard deviation.

Technical Details

For a single standard deviation from a normal distribution with unknown mean, a two-sided, $100(1 - \alpha)\%$ confidence interval is calculated by

$$\left[s \left\{ \frac{n-1}{\chi^2_{1-\alpha/2, n-1}} \right\}^{1/2}, s \left\{ \frac{n-1}{\chi^2_{\alpha/2, n-1}} \right\}^{1/2} \right]$$

640-2 Confidence Intervals for One Standard Deviation using Standard Deviation

A one-sided $100(1 - \alpha)\%$ upper confidence limit is calculated by

$$s \left\{ \frac{n-1}{\chi^2_{\alpha, n-1}} \right\}^{1/2}$$

Similarly, the one-sided $100(1 - \alpha)\%$ lower confidence limit is

$$s \left\{ \frac{n-1}{\chi^2_{1-\alpha, n-1}} \right\}^{1/2}$$

For two-sided intervals, the distance from the standard deviation to each of the limits is different. Thus, instead of specifying the distance to the limits we specify the width of the interval, W .

The basic equation for determining sample size for a two-sided interval when W has been specified is

$$W = s \left\{ \frac{n-1}{\chi^2_{\alpha/2, n-1}} \right\}^{1/2} - s \left\{ \frac{n-1}{\chi^2_{1-\alpha/2, n-1}} \right\}^{1/2}$$

For one-sided intervals, the distance from the standard deviation to limits, D , is specified.

The basic equation for determining sample size for a one-sided upper limit when D has been specified is

$$D = s \left\{ \frac{n-1}{\chi^2_{\alpha/2, n-1}} \right\}^{1/2} - s$$

The basic equation for determining sample size for a one-sided lower limit when D has been specified is

$$D = s - s \left\{ \frac{n-1}{\chi^2_{1-\alpha/2, n-1}} \right\}^{1/2}$$

These equations can be solved for any of the unknown quantities in terms of the others.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population standard deviation is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population standard deviation is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as *0.90, 0.95* or *0.90 to 0.99 by 0.01*.

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

Precision

Confidence Interval Width (Two-Sided)

This is the distance from the lower confidence limit to the upper confidence limit. The distance from the standard deviation to the lower and upper limits is not equal.

You can enter a single value or a list of values. The value(s) must be greater than zero.

640-4 Confidence Intervals for One Standard Deviation using Standard Deviation

Distance from SD to Limit (One-Sided)

This is the distance from the standard deviation to the lower or upper limit of the confidence interval, depending on whether the Interval Type is set to Lower Limit or Upper Limit.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Standard Deviation

S (Standard Deviation)

Enter an estimate of the standard deviation (must be positive). The sample size and width calculations assume that the value entered here is the standard deviation estimate that is obtained from the sample. If the sample standard deviation is different from the one specified here, the width may be narrower or wider than specified.

For controlling the probability that the width is less than the value specified, see the procedure 'Confidence Intervals for One Standard Deviation with Tolerance Probability'.

For confidence intervals with widths that are specified in terms of a percentage of relative error, see the procedure 'Confidence Intervals for One Standard Deviation using Relative Error'.

One common method for estimating the standard deviation is the range divided by 4, 5, or 6.

You can enter a range of values such as *1 2 3* or *1 to 10 by 1*.

Press the Standard Deviation Estimator button to load the Standard Deviation Estimator window.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to construct a two-sided 95% confidence interval for the standard deviation such that the width of the interval is no wider than 20 units. The confidence level is set at 0.95, but 0.99 is included for comparative purposes. The standard deviation estimate, based on the range of data values, is 14. Instead of examining only the interval width of 20, a series of widths from 16 to 24 will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Standard Deviation using Standard Deviation** procedure window by clicking on **Confidence Intervals**, then **Standard Deviations**, then **One Standard Deviation using Standard Deviation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95 0.99
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	16 to 24 by 1
S (Standard Deviation)	14

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals

Confidence Level	Sample Size (N)	Target Width	Actual Width	Standard Deviation (S)	Lower Limit	Upper Limit
0.950	40	16.000	15.806	34.000	27.851	43.657
0.990	67	16.000	15.873	34.000	27.715	43.588
0.950	36	17.000	16.774	34.000	27.577	44.351
0.990	60	17.000	16.870	34.000	27.420	44.289
0.950	32	18.000	17.944	34.000	27.258	45.202
0.990	54	18.000	17.891	34.000	27.128	45.019
0.950	30	19.000	18.629	34.000	27.078	45.707
0.990	49	19.000	18.900	34.000	26.850	45.750
0.950	27	20.000	19.819	34.000	26.776	46.595
0.990	45	20.000	19.842	34.000	26.599	46.441
0.950	25	21.000	20.751	34.000	26.548	47.299
0.990	41	21.000	20.939	34.000	26.317	47.256
0.950	23	22.000	21.827	34.000	26.295	48.122
0.990	38	22.000	21.892	34.000	26.080	47.972

640-6 Confidence Intervals for One Standard Deviation using Standard Deviation

0.950	22	23.000	22.430	34.000	26.158	48.588
0.990	35	23.000	22.986	34.000	25.818	48.804
0.950	20	24.000	23.803	34.000	25.857	49.659
0.990	33	24.000	23.813	34.000	25.627	49.440

References

Hahn, G. J. and Meeker, W.Q. 1991. Statistical Intervals. John Wiley & Sons. New York.

Report Definitions

Confidence level is the proportion of confidence intervals (constructed with this same confidence level, sample size, etc.) that would contain the population standard deviation.

N is the size of the sample drawn from the population.

Width is distance from the lower limit to the upper limit.

Target Width is the value of the width that is entered into the procedure.

Actual Width is the value of the width that is obtained from the procedure.

Standard Deviation (S) is the assumed sample standard deviation.

Lower Limit is the lower limit of the confidence interval.

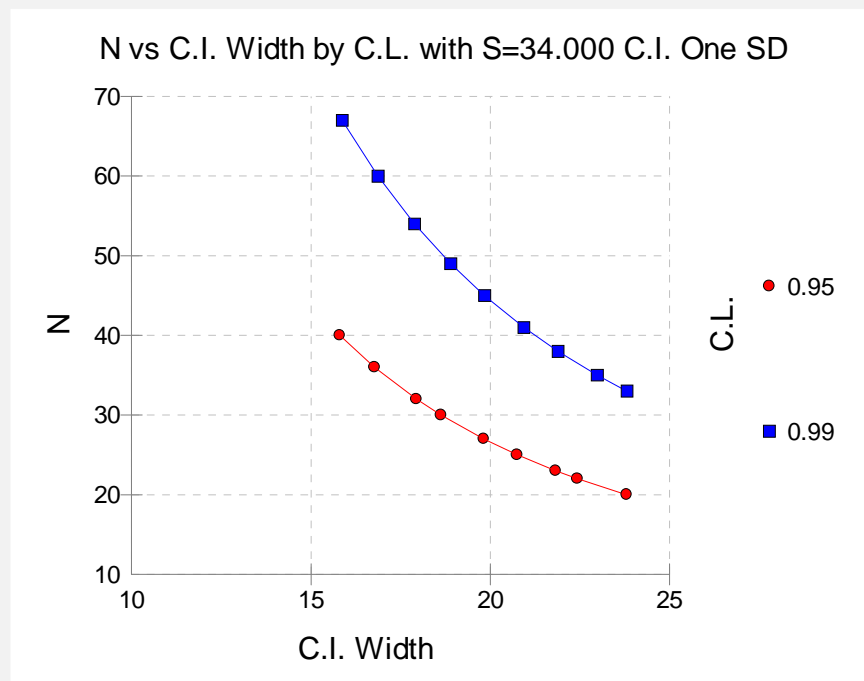
Upper Limit is the upper limit of the confidence interval.

Summary Statements

A sample size of 40 produces a two-sided 95% confidence interval with a width equal to 15.806 when the standard deviation is 34.000.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the confidence interval width for the two confidence levels.

Example 2 – Validation using Hahn and Meeker

Hahn and Meeker (1991) page 56 give an example of a calculation for a confidence interval on the standard deviation when the confidence level is 95%, the standard deviation is 1.31, and the interval width is 2.9795. The necessary sample size is 5.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Standard Deviation using Standard Deviation** procedure window by clicking on **Confidence Intervals**, then **Standard Deviations**, then **One Standard Deviation using Standard Deviation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	2.9795
S (Standard Deviation)	1.31

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Confidence Level	Sample Size (N)	Target Width	Actual Width	Standard Deviation (S)	Lower Limit	Upper Limit
0.950	5	2.980	2.979	1.310	0.785	3.764

PASS also calculated the necessary sample size to be 5.

640-8 Confidence Intervals for One Standard Deviation using Standard Deviation

Chapter 641

Confidence Intervals for One Standard Deviation with Tolerance Probability

Introduction

This procedure calculates the sample size necessary to achieve a specified width (or in the case of one-sided intervals, the distance from the standard deviation to the confidence limit) with a given tolerance probability at a stated confidence level for a confidence interval about a single standard deviation when the underlying data distribution is normal.

Technical Details

For a single standard deviation from a normal distribution with unknown mean, a two-sided, $100(1 - \alpha)\%$ confidence interval is calculated by

$$\left[s \left\{ \frac{n-1}{\chi^2_{1-\alpha/2, n-1}} \right\}^{1/2}, s \left\{ \frac{n-1}{\chi^2_{\alpha/2, n-1}} \right\}^{1/2} \right]$$

A one-sided $100(1 - \alpha)\%$ upper confidence limit is calculated by

$$s \left\{ \frac{n-1}{\chi^2_{\alpha, n-1}} \right\}^{1/2}$$

641-2 Confidence Intervals for One Standard Deviation with Tolerance Probability

Similarly, the one-sided $100(1 - \alpha)\%$ lower confidence limit is

$$s \left\{ \frac{n-1}{\chi^2_{1-\alpha, n-1}} \right\}^{1/2}$$

For two-sided intervals, the distance from the standard deviation to each of the limits is different. Thus, instead of specifying the distance to the limits we specify the width of the interval, W .

The basic equation for determining sample size for a two-sided interval when W has been specified is

$$W = s \left\{ \frac{n-1}{\chi^2_{\alpha/2, n-1}} \right\}^{1/2} - s \left\{ \frac{n-1}{\chi^2_{1-\alpha/2, n-1}} \right\}^{1/2}$$

For one-sided intervals, the distance from the standard deviation to limits, D , is specified.

The basic equation for determining sample size for a one-sided upper limit when D has been specified is

$$D = s \left\{ \frac{n-1}{\chi^2_{\alpha/2, n-1}} \right\}^{1/2} - s$$

The basic equation for determining sample size for a one-sided lower limit when D has been specified is

$$D = s - s \left\{ \frac{n-1}{\chi^2_{1-\alpha/2, n-1}} \right\}^{1/2}$$

These equations can be solved for any of the unknown quantities in terms of the others.

There is an additional subtlety that arises when the standard deviation is to be chosen for estimating sample size. The sample sizes determined from the formulas above produce confidence intervals with the specified widths only when the future sample has a sample standard deviation that is no greater than the value specified.

As an example, suppose that 15 individuals are sampled in a pilot study, and a standard deviation estimate of 3.5 is obtained from the sample. The purpose of a later study is to estimate the standard deviation with a confidence interval with width no greater than 1.3 units. Suppose further that the sample size needed is calculated to be 105 using the formula above with 3.5 as the estimate for the standard deviation. The sample of size 105 is then obtained from the population, but the standard deviation of the 105 individuals turns out to be 3.9 rather than 3.5. The confidence interval is computed and the width of the interval is greater than 1.3 units.

This example illustrates the need for an adjustment to adjust the sample size such that the width or distance from the standard deviation to the confidence limits will be below the specified value with known probability.

Such an adjustment for situations where a previous sample is used to estimate the standard deviation is derived for the case of confidence intervals for a mean by Harris, Horvitz, and Mood

(1948) and discussed in Zar (1984) and Hahn and Meeker (1991). The adjustment is made by replacing s with

$$s = \sigma \sqrt{F_{1-\gamma; n-1, m-1}}$$

where $1 - \gamma$ is the probability that the width or distance from the standard deviation to the confidence limit will be below the specified value, and m is the sample size in the previous sample that was used to estimate the standard deviation.

The corresponding adjustment when no previous sample is available is discussed in Kupper and Hafner (1989) and Hahn and Meeker (1991). The adjustment in this case is

$$s = \sigma \left(\frac{\chi^2_{1-\gamma, n-1}}{n-1} \right)^{1/2}$$

where, again, $1 - \gamma$ is the probability that the width or distance from the standard deviation to the confidence limit will be below the specified value.

Each of these adjustments accounts for the variability in a future estimate of the standard deviation. In the first adjustment formula (Harris, Horvitz, and Mood, 1948), the distribution of the standard deviation is based on the estimate from a previous sample. In the second adjustment formula, the distribution of the standard deviation is based on a specified value that is assumed to be the population standard deviation.

For this procedure, both adjustments are adapted from the case of a one-sample confidence interval for a single mean to the case of a one-sample confidence interval for a single standard deviation.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population standard deviation is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence and Tolerance

Confidence Level (1 – Alpha)

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population standard deviation is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as 0.90, 0.95 or 0.90 to 0.99 by 0.01.

Tolerance Probability

This is the probability that a future interval with sample size N and the specified confidence level will have a width or distance from the standard deviation to the limit that is less than or equal to the width or distance specified.

If a tolerance probability is not used, as in the 'Confidence Intervals for One Standard Deviation using Standard Deviation' procedure, the sample size is calculated for the expected width or distance from the SD to the limit, which assumes that the future standard deviation will also be the one specified.

Using a tolerance probability implies that the standard deviation of the future sample will not be known in advance, and therefore, an adjustment is made to the sample size formula to account for the variability in the standard deviation. Use of a tolerance probability is similar to using an upper bound for the standard deviation in the 'Confidence Intervals for One Standard Deviation using Standard Deviation' procedure.

Values between 0 and 1 can be entered. The choice of the tolerance probability depends upon how important it is that the width distance from the interval limits to the standard deviation is at most the value specified.

You can enter a range of values such as 0.70 0.80 0.90 or 0.70 to 0.95 by 0.05.

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

Precision

Confidence Interval Width (Two-Sided)

This is the distance from the lower confidence limit to the upper confidence limit. The distance from the standard deviation to the lower and upper limits is not equal.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Distance from SD to Limit (One-Sided)

This is the distance from the standard deviation to the lower or upper limit of the confidence interval, depending on whether the Interval Type is set to Lower Limit or Upper Limit.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Standard Deviation

Standard Deviation Source

This procedure permits two sources for estimates of the standard deviation:

- **S is a Population Standard Deviation**

This option should be selected if there is no previous sample that can be used to obtain an estimate of the standard deviation. In this case, the algorithm assumes that future sample obtained will be from a population with standard deviation S.

- **S from a Previous Sample**

This option should be selected if the estimate of the standard deviation is obtained from a previous random sample from the same distribution as the one to be sampled. The sample size of the previous sample must also be entered under 'Sample Size of Previous Sample'.

Standard Deviation – S is a Population Standard Deviation

S (Standard Deviation)

Enter an estimate of the standard deviation (must be positive). In this case, the algorithm assumes that future samples obtained will be from a population with standard deviation S.

One common method for estimating the standard deviation is the range divided by 4, 5, or 6.

You can enter a range of values such as *1 2 3* or *1 to 10 by 1*.

Press the Standard Deviation Estimator button to load the Standard Deviation Estimator window.

Standard Deviation – S from a Previous Sample

S (SD Estimated from a Previous Sample)

Enter an estimate of the standard deviation from a previous (or pilot) study. This value must be positive.

A range of values may be entered.

Press the Standard Deviation Estimator button to load the Standard Deviation Estimator window.

Sample Size of Previous Sample

Enter the sample size that was used to estimate the standard deviation entered in S (SD Estimated from a Previous Sample).

This value is entered only when 'Standard Deviation Source:' is set to 'S from a Previous Sample'.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

A researcher would like to estimate the standard deviation of a population with 95% confidence. It is very important that the interval width is less than 5 grams. Data available from a previous study are used to provide an estimate of the standard deviation. The estimate of the standard deviation is 25.4 grams, from a sample of size 12.

The goal is to determine the sample size necessary to obtain a two-sided confidence interval such that the width of the interval is less than 5 grams. Tolerance probabilities of 0.70 to 0.95 will be examined.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Standard Deviation with Tolerance Probability** procedure window by clicking on **Confidence Intervals**, then **Standard Deviations**, then **One Standard Deviation with Tolerance Probability**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
Tolerance Probability	0.70 to 0.95 by 0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width	5
Standard Deviation Source	S from a Previous Sample
S	25.4
Sample Size of Previous Sample.....	12

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals

Confidence Level	Sample Size (N)	Target Width	Actual Width	Standard Deviation (S)	Tolerance Probability
0.950	273	5.000	4.997	25.400	0.700
0.950	293	5.000	4.998	25.400	0.750
0.950	318	5.000	4.997	25.400	0.800
0.950	351	5.000	4.993	25.400	0.850
0.950	398	5.000	4.995	25.400	0.900
0.950	484	5.000	4.997	25.400	0.950

Sample size for estimate of S from previous sample = 12.

References

Hahn, G. J. and Meeker, W.Q. 1991. Statistical Intervals. John Wiley & Sons. New York.
 Zar, J. H. 1984. Biostatistical Analysis. Second Edition. Prentice-Hall. Englewood Cliffs, New Jersey.
 Harris, M., Horvitz, D. J., and Mood, A. M. 1948. 'On the Determination of Sample Sizes in Designing Experiments', Journal of the American Statistical Association, Volume 43, No. 243, pp. 391-402.

Report Definitions

Confidence Level is the proportion of confidence intervals (constructed with this same confidence level, sample size, etc.) that would contain the population standard deviation.

N is the size of the sample drawn from the population.

Width is the distance from the lower limit to the upper limit.

Target Width is the value of the width that is entered into the procedure.

Actual Width is the value of the width that is obtained from the procedure.

Standard Deviation (S) is the estimated standard deviation based on a previous sample.

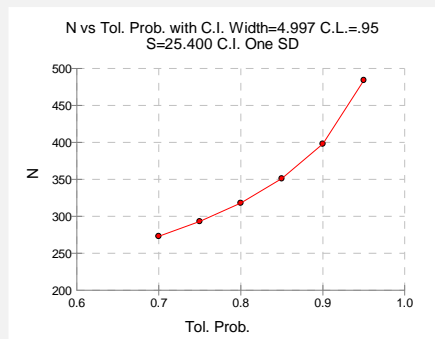
Tolerance Probability is the probability that a future interval with sample size N and corresponding confidence level will have a width that is less than or equal to the specified width.

Summary Statements

The probability is 0.700 that a sample size of 273 will produce a two-sided 95% confidence interval with a width that is less than or equal to 4.997 if the population standard deviation is estimated to be 25.400 by a previous sample of size 12.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the tolerance probability.

Example 2 – Validation using Simulation

We could not find a published result for confidence intervals for a standard deviation with tolerance probability. This procedure is validated using a simulation. A simulation was run with a confidence level of 95%, sample size of 100, a specified confidence interval width of 6, and assumed population standard deviation of 20.

The number of simulations was 100,000. The proportion of intervals with width less than 6 (tolerance probability) was 0.80168.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Standard Deviation with Tolerance Probability** procedure window by clicking on **Confidence Intervals**, then **Standard Deviations**, then **One Standard Deviation with Tolerance Probability**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Tolerance Probability
Confidence Level	0.95
Tolerance Probability	<i>Ignored since this is the Find setting</i>
N (Sample Size)	100
Interval Type	Two-Sided
Confidence Interval Width	6
Standard Deviation Source	S is a Population Standard Deviation
S	20

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals					
Confidence Level	Sample Size (N)	Target Width	Actual Width	Standard Deviation (S)	Tolerance Probability
0.950	100	6.000	6.000	20.000	0.802

PASS calculated the tolerance probability to be 0.802, which is well within the simulation error of the simulation estimate of 0.80168.

Chapter 642

Confidence Intervals for One Standard Deviation Using Relative Error

Introduction

This routine calculates the necessary sample size such that a sample standard deviation estimate will achieve a specified relative distance from the true population standard deviation at a stated confidence level when the underlying data distribution is normal.

Caution: This procedure controls the relative width of the interval as a proportion of the true population standard deviation. For controlling the absolute width of the interval see the procedures Confidence Intervals for One Standard Deviation using Standard Deviation and Confidence Intervals for One Standard Deviation with Tolerance Probability.

Technical Details

Following the results of Desu and Raghavarao (1990) and Greenwood and Sandomire (1950), let s be the standard deviation estimate based on a sample from a normal distribution with unknown μ and unknown σ . Let r be the proportion of σ such that s is within $r\sigma$ of σ with desired confidence. If

$$p_1 = \Pr(s > \sigma + r\sigma) = \Pr(s > \sigma(1 + r))$$

and

$$p_2 = \Pr(s < \sigma - r\sigma) = \Pr(s < \sigma(1 - r))$$

642-2 Confidence Intervals for One Standard Deviation using Relative Error

The confidence level for estimating σ within proportion r is $1 - p_1 - p_2$.

Since $\frac{(n-1)s^2}{\sigma^2} \sim \chi_{n-1}^2$, it is useful to rewrite p_1 and p_2 as

$$p_1 = \Pr\left(\frac{(n-1)s^2}{\sigma^2} > (n-1)(1+r)^2\right)$$

and

$$p_2 = \Pr\left(\frac{(n-1)s^2}{\sigma^2} < (n-1)(1-r)^2\right)$$

Using the chi-square distribution, these equations can be solved for any of the unknown quantities (n , r , $p_1 + p_2$) in terms of the others.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and the standard deviation estimates are obtained from each sample, the proportion of those estimates that are within $r\sigma$ of σ is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and the standard deviation estimates are obtained from each sample, the proportion of those estimates that are within $r\sigma$ of σ is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as 0.90, 0.95 or 0.90 to 0.99 by 0.01.

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

In each case the limits are based on the relative error of the true population standard deviation.

Precision

Relative Error

This is the distance from the true standard deviation as a proportion of the true standard deviation.

You can enter a single value or a list of values. The value(s) must be between 0 and 100000.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to be 95% confident that estimated standard deviation is within 10% of the true population standard deviation. In addition to 10% relative error, 5%, 15%, 20% and 25% will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Standard Deviation using Relative Error** procedure window by clicking on **Confidence Intervals**, then **Standard Deviations**, then **One Standard Deviation using Relative Error**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Relative Error	0.05 to 0.25 by 0.05

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Relative Error Confidence Intervals

Target Confidence Level	Actual Confidence Level	Sample Size (N)	Relative Error
0.950	0.950	769	0.050
0.950	0.950	193	0.100
0.950	0.951	87	0.150
0.950	0.951	49	0.200
0.950	0.952	32	0.250

References

Greenwood, J. A. and Sandomire, M. M. 1950. 'Sample Size Required for Estimating the Standard Deviation as a Per Cent of its True Value', Journal of the American Statistical Association, Vol. 45, No. 250, pp. 257-260.
Desu, M. M. and Raghavarao, D. 1990. Sample Size Methodology. Academic Press. New York.

Report Definitions

Confidence Level is the proportion of standard deviation estimates that will be within the relative error of the true standard deviation.

Target Confidence Level is the value of the confidence level that is entered into the procedure.

Actual Confidence Level is the value of the confidence level that is obtained from the procedure.

Sample Size (N) is the size of the sample drawn from the population.

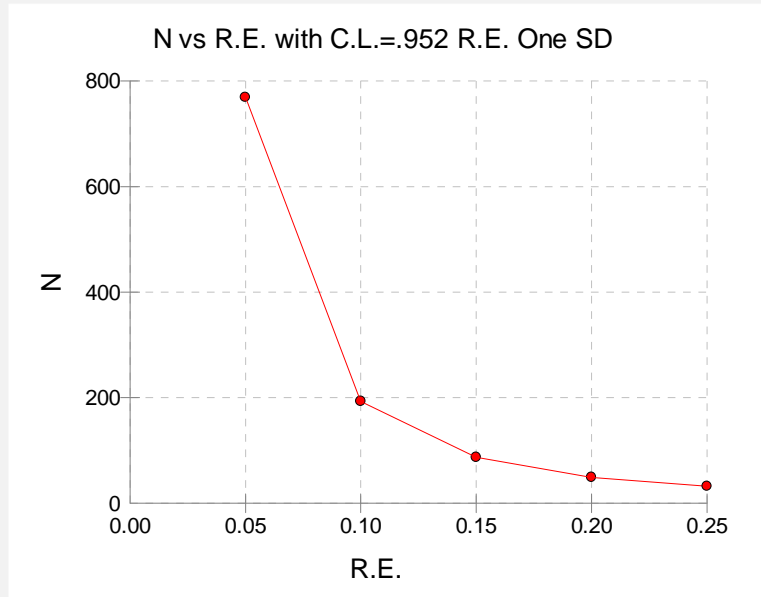
Relative Error is the distance from the true standard deviation as a proportion of the true standard deviation.

Summary Statements

With a sample size of 769, the probability is 0.950 (95% confidence) that the estimate of the standard deviation will be within 5% of the true population standard deviation.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the relative error.

Example 2 – Validation using Greenwood and Sandomire

Greenwood and Sandomire (1950) page 259 give an example of a calculation in which the desired confidence level is 80% and the relative error is 10%. The necessary sample size is 84.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Standard Deviation using Relative Error** procedure window by clicking on **Confidence Intervals**, then **Standard Deviations**, then **One Standard Deviation using Relative Error**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.80
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Relative Error	0.1

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Target Confidence Level	Actual Confidence Level	Sample Size (N)	Relative Error
0.800	0.802	84	0.100

PASS also calculated the necessary sample size to be 84.

Chapter 650

Inequality Tests for One Variance

Introduction

Occasionally, researchers are interested in the estimation of the variance (or standard deviation) rather than the mean. This module calculates the sample size and performs power analysis for hypothesis tests concerning a single variance.

Technical Details

Assuming that a variable X is normally distributed with mean μ and variance σ^2 , the sample variance is distributed as a Chi-square random variable with $N - 1$ degrees of freedom, where N is the sample size. That is,

$$X^2 = \frac{(N-1)s^2}{\sigma^2}$$

is distributed as a Chi-square random variable. The sample statistic, s^2 , is calculated as follows

$$s^2 = \frac{\sum_{i=1}^N (X_i - \bar{X})^2}{N-1}.$$

The power or sample size of a hypothesis test about the variance can be calculated using the appropriate one of the following three formulas from Ostle (1988) page 130.

Case 1: $H_0: \sigma_1^2 = \sigma_0^2$ versus $H_a: \sigma_1^2 \neq \sigma_0^2$

$$\beta = P\left(\frac{\sigma_0^2}{\sigma_1^2} \chi_{\alpha/2, N-1}^2 < \chi^2 < \frac{\sigma_0^2}{\sigma_1^2} \chi_{1-\alpha/2, N-1}^2\right)$$

Case 2: $H_0: \sigma_1^2 = \sigma_0^2$ versus $H_a: \sigma_1^2 > \sigma_0^2$

$$\beta = P\left(\chi^2 < \frac{\sigma_0^2}{\sigma_1^2} \chi_{1-\alpha, N-1}^2\right)$$

650-2 Inequality Tests for One Variance

Case 3: $H_0: \sigma_1^2 = \sigma_0^2$ versus $H_a: \sigma_1^2 < \sigma_0^2$

$$\beta = P\left(\chi^2 > \frac{\sigma_0^2}{\sigma_1^2} \chi_{\alpha, N-1}^2\right)$$

Procedure Options

This section describes the options that are unique to this procedure. To find out more about using the other tabs, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

If your only interest is in determining the appropriate sample size for a confidence interval, set power or beta to 0.5.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when a true null hypothesis is rejected.

Values must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

This is the number of observations in the study.

Effect Size

Scale

Specify whether *V0* and *V1* are variances or standard deviations.

V0 (Baseline Variance)

Enter one or more value(s) of the baseline variance. This variance will be compared to the alternative variance. It must be greater than zero.

Actually, only the ratio of the two variances (or standard deviations) is used, so you can enter a one here and enter the ratio value in the *V1* box.

If Scale is *Standard Deviation* this value is treated as a standard deviation rather than a variance.

V1 (Alternative Variance)

Enter one or more value(s) of the alternative variance. This variance will be compared to the baseline variance. It must be greater than zero.

Actually, only the ratio of the two variances (or standard deviations) is used, so you can enter a one for *V0* and enter a ratio value here.

If Scale is *Standard Deviation* this value is treated as a standard deviation rather than a variance.

Test

Alternative Hypothesis

This option specifies the alternative hypothesis. This implicitly specifies the direction of the hypothesis test. The null hypothesis is always $H_0: \sigma_0^2 = \sigma_1^2$.

Note that the alternative hypothesis enters into power calculations by specifying the rejection region of the hypothesis test. Its accuracy is critical.

Possible selections are:

- **Ha: V0 <> V1**

This selection yields a *two-tailed* test. Use this option when you are testing whether the variances are different but you do not want to specify beforehand which variance is larger.

- **Ha: V0 > V1**

The options yields a *one-tailed* test. Use it when you are only interested in the case in which *V1* is less than *V0*.

- **Ha: V0 < V1**

This option yields a *one-tailed* test. Use it when you are only interested in the case in which *V1* is greater than *V0*.

650-4 Inequality Tests for One Variance

Known Mean

The degrees of freedom of the Chi-square test is $N - 1$ if the mean is calculated from the data (this is usually the case) or it is N if the mean is known. Check this box if the mean is known. This will cause an increase of the sample size by one.

Iterations Tab

This tab sets the option used in the iterative procedure.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating the Power

A machine used to perform a particular analysis is to be replaced with a new type of machine if the new machine reduces the variation in the output. The current machine has been tested repeatedly and found to have an output variance of 42.5. The new machine will be cost effective if it can reduce the variance by 30% to 29.75. If the significance level is set to 0.05, calculate the power for sample sizes of 10, 50, 90, 130, 170, 210, and 250.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Variance** procedure window by clicking on **Variances**, then **One Variance**, then **Inequality Tests**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	10 50 90 130 170 210 250
Scale	Variance
V0 (Baseline Variance)	42.5
V1 (Alternative Variance)	29.75
Alternative Hypothesis	Ha: V0 > V1
Known Mean	Not checked

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when $H_0: V_0 = V_1$ versus $H_a: V_0 > V_1$

Power	N	V0	V1	Alpha	Beta
0.144479	10	42.5000	29.7500	0.050000	0.855521
0.505556	50	42.5000	29.7500	0.050000	0.494444
0.747747	90	42.5000	29.7500	0.050000	0.252253
0.881740	130	42.5000	29.7500	0.050000	0.118260
0.947851	170	42.5000	29.7500	0.050000	0.052149
0.978059	210	42.5000	29.7500	0.050000	0.021941
0.991108	250	42.5000	29.7500	0.050000	0.008892

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N is the size of the sample drawn from the population.

V0 is the value of the population variance under the null hypothesis.

V1 is the value of the population variance under the alternative hypothesis.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

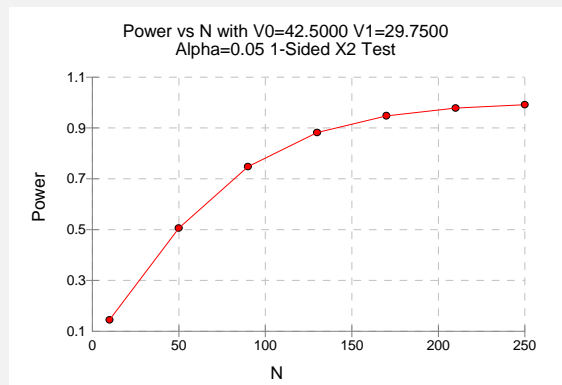
Beta is the probability of accepting a false null hypothesis. It should be small.

Summary Statements

A sample size of 10 achieves 14% power to detect a difference of 12.7500 between the null hypothesis variance of 42.5000 and the alternative hypothesis variance of 29.7500 using a one-sided, Chi-square hypothesis test with a significance level (alpha) of 0.050000.

This report shows the calculated power for each scenario.

Plots Section



This plot shows the power versus the sample size. We see that a sample size of about 150 is necessary to achieve a power of 0.90.

Example 2 – Calculating Sample Size

Continuing with the previous example, the analyst wants to find the necessary sample sizes to achieve a power of 0.9, for two significance levels, 0.01 and 0.05, and for several variance values. To make interpreting the output easier, the analyst decides to switch from the absolute scale to a ratio scale. To accomplish this, the baseline variance is set at 1.0 and the alternative variances of 0.2, 0.3, 0.4, 0.5, 0.6, and 0.7 are tried.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Variance** procedure window by clicking on **Variances**, then **One Variance**, then **Inequality Tests**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.01 0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Scale	Variance
V0 (Baseline Variance)	1.0
V1 (Alternative Variance)	0.2 to 0.7 by 0.1
Alternative Hypothesis	Ha: V0 > V1
Known Mean	Not checked

Output

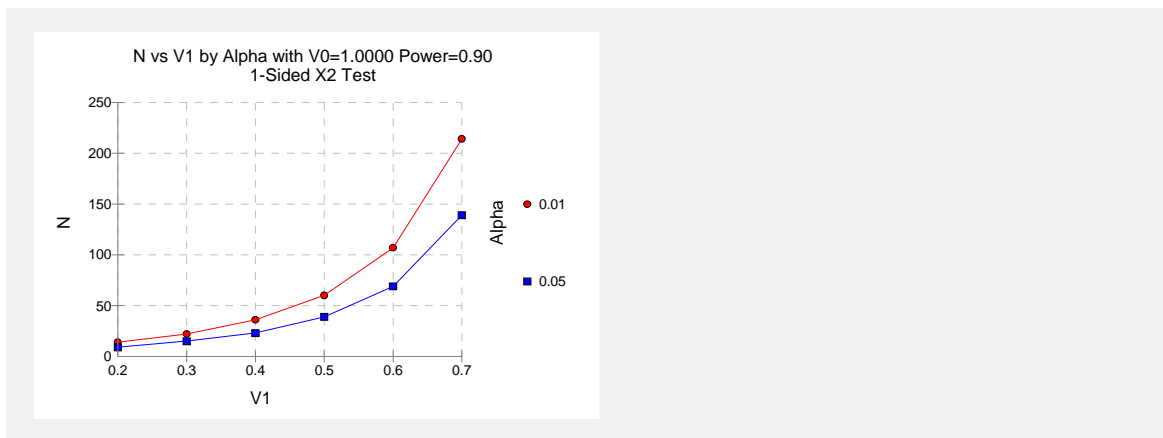
Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H0: V0 = V1 versus Ha: V0>V1					
Power	N	V0	V1	Alpha	Beta
0.917341	14	1.0000	0.2000	0.010000	0.082659
0.909022	9	1.0000	0.2000	0.050000	0.090978
0.900907	22	1.0000	0.3000	0.010000	0.099093
0.919352	15	1.0000	0.3000	0.050000	0.080648
0.903675	36	1.0000	0.4000	0.010000	0.096325
0.900670	23	1.0000	0.4000	0.050000	0.099330
0.901634	60	1.0000	0.5000	0.010000	0.098366
0.904235	39	1.0000	0.5000	0.050000	0.095765
0.901264	107	1.0000	0.6000	0.010000	0.098736
0.900775	69	1.0000	0.6000	0.050000	0.099225
0.900601	214	1.0000	0.7000	0.010000	0.099399
0.901171	139	1.0000	0.7000	0.050000	0.098829

This report shows the necessary sample size for each scenario.

Plots Section



This plot shows the necessary sample size for various values of $V1$. Note that as $V1$ gets farther from zero, the required sample size increases.

Example 3 – Validation using Zar

Zar (1994) page 117 presents an example with $V0 = 1.5$, $V1 = 2.6898$, $N = 40$, $Alpha = 0.05$, and $Power = 0.84$. We will run this example through *PASS*.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the *PASS* Home window, load the **Inequality Tests for One Variance** procedure window by clicking on **Variances**, then **One Variance**, then **Inequality Tests**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	40
Scale	Variance
V0 (Baseline Variance)	1.5
V1 (Alternative Variance)	2.6898
Alternative Hypothesis	Ha: V0 < V1
Known Mean	Not checked

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H0: $V_0 = V_1$ versus Ha: $V_0 < V_1$					
Power	N	V0	V1	Alpha	Beta
0.835167	40	1.5000	2.6898	0.050000	0.164833

PASS calculated the power at 0.835167 which matches Zar's result of 0.84 within rounding.

Example 4 – Validation using Davies

Davies (1971) page 40 presents an example of determining N when (in the standard deviation scale) $V_0 = 0.04$, $V_1 = 0.10$, $\text{Alpha} = 0.05$, and $\text{Power} = 0.99$. Davies calculates N to be 13. We will run this example through *PASS*.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the *PASS* Home window, load the **Inequality Tests for One Variance** procedure window by clicking on **Variances**, then **One Variance**, then **Inequality Tests**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.99
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Scale	Standard Deviation
V0 (Baseline Variance)	0.04
V1 (Alternative Variance)	0.10
Alternative Hypothesis	Ha: $V_0 < V_1$
Known Mean	Not checked

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H0: $S_0 = S_1$ versus Ha: $S_0 < S_1$					
Power	N	S0	S1	Alpha	Beta
0.992381	13	0.0400	0.1000	0.050000	0.007619

PASS calculated an N of 13 which matches Davies' result.

Chapter 651

Confidence Intervals for One Variance using Variance

Introduction

This routine calculates the sample size necessary to achieve a specified interval width or distance from the variance to the confidence limit at a stated confidence level for a confidence interval about the variance when the underlying data distribution is normal.

Caution: This procedure assumes that the variance of the future sample will be the same as the variance that is specified. If the variance to be used in the procedure is estimated from a previous sample or represents the population variance, the Confidence Intervals for One Variance with Tolerance Probability procedure should be considered. That procedure controls the probability that the width or the distance from the variance to the confidence limits will be less than or equal to the value specified. The Confidence Intervals for One Variance using Relative Error controls the width or distance from the variance to the limits by controlling the width or distance as a percent of the true variance.

Technical Details

For a single variance from a normal distribution with unknown mean, a two-sided, $100(1 - \alpha)\%$ confidence interval is calculated by

$$\left[\frac{(n-1)s^2}{\chi^2_{1-\alpha/2, n-1}}, \frac{(n-1)s^2}{\chi^2_{\alpha/2, n-1}} \right]$$

A one-sided $100(1 - \alpha)\%$ upper confidence limit is calculated by

$$\frac{(n-1)s^2}{\chi^2_{\alpha, n-1}}$$

651-2 Confidence Intervals for One Variance using Variance

Similarly, the one-sided $100(1 - \alpha)\%$ lower confidence limit is

$$\frac{(n-1)s^2}{\chi^2_{1-\alpha, n-1}}$$

For two-sided intervals, the distance from the variance to each of the limits is different. Thus, instead of specifying the distance to the limits we specify the width of the interval, W .

The basic equation for determining sample size for a two-sided interval when W has been specified is

$$W = \frac{(n-1)s^2}{\chi^2_{\alpha/2, n-1}} - \frac{(n-1)s^2}{\chi^2_{1-\alpha/2, n-1}}$$

For one-sided intervals, the distance from the variance to limits, D , is specified.

The basic equation for determining sample size for a one-sided upper limit when D has been specified is

$$D = \frac{(n-1)s^2}{\chi^2_{\alpha, n-1}} - s^2$$

The basic equation for determining sample size for a one-sided lower limit when D has been specified is

$$D = s^2 - \frac{(n-1)s^2}{\chi^2_{1-\alpha, n-1}}$$

These equations can be solved for any of the unknown quantities in terms of the others.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population variance is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population variance is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as 0.90, 0.95 or 0.90 to 0.99 by 0.01.

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

Precision

Confidence Interval Width (Two-Sided)

This is the distance from the lower confidence limit to the upper confidence limit. The distance from the variance to the lower and upper limits is not equal.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Distance from Var to Limit (One-Sided)

This is the distance from the variance to the lower or upper limit of the confidence interval, depending on whether the Interval Type is set to Lower Limit or Upper Limit.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Variance

Variance

Enter an estimate of the variance (must be positive). The sample size and width calculations assume that the value entered here is the variance estimate that is obtained from the sample. If the sample variance is different from the one specified here, the width may be narrower or wider than specified.

For controlling the probability that the width is less than the value specified, see the procedure 'Confidence Intervals for One Variance with Tolerance Probability'.

For confidence intervals with widths that are specified in terms of a percentage of relative error, see the procedure 'Confidence Intervals for One Variance using Relative Error'.

651-4 Confidence Intervals for One Variance using Variance

You can enter a range of values such as *1 2 3* or *1 to 10 by 1*.

Press the Standard Deviation Estimator button to load the Standard Deviation Estimator window.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to construct a two-sided 95% confidence interval for the variance such that the width of the interval is no wider than 40 units. The confidence level is set at 0.95, but 0.99 is included for comparative purposes. The variance estimate, based on the range of data values, is 24. Instead of examining only the interval width of 40, a series of widths from 30 to 50 will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Variance using Variance** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **One Variance using Variance**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95 0.99
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	30 to 50 by 5
Variance	24

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals

Confidence Level	Sample Size (N)	Target Width	Actual Width	Variance	Lower Limit	Upper Limit
0.950	28	30.000	29.463	24.000	15.002	44.465
0.990	46	30.000	29.663	24.000	14.761	44.424
0.950	22	35.000	34.808	24.000	14.206	49.013
0.990	36	35.000	34.924	24.000	13.936	48.860
0.950	19	40.000	38.783	24.000	13.703	52.486
0.990	30	40.000	39.745	24.000	13.299	53.044
0.950	16	45.000	44.392	24.000	13.096	57.488
0.990	26	45.000	44.251	24.000	12.786	57.036
0.950	14	50.000	49.678	24.000	12.613	62.291
0.990	23	50.000	48.754	24.000	12.338	61.092

References

Zar, Jerrold H. 1984. Biostatistical Analysis. Second Edition. Prentice-Hall. Englewood Cliffs, New Jersey.

Report Definitions

Confidence level is the proportion of confidence intervals (constructed with this same confidence level, sample size, etc.) that would contain the population variance.

N is the size of the sample drawn from the population.

Width is the distance from the lower limit to the upper limit.

Target Width is the value of the width that is entered into the procedure.

Actual Width is the value of the width that is obtained from the procedure.

Variance is the assumed sample variance.

Lower Limit is the lower limit of the confidence interval.

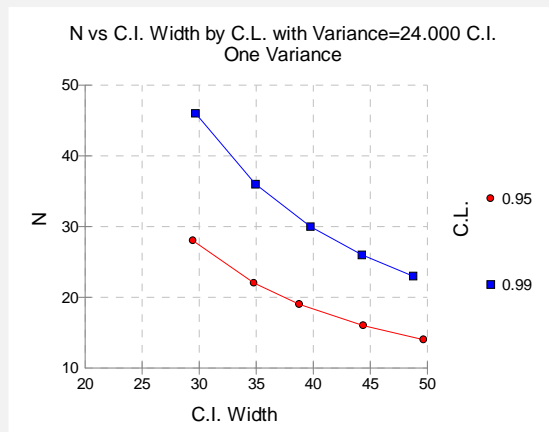
Upper Limit is the upper limit of the confidence interval.

Summary Statements

A sample size of 28 produces a two-sided 95% confidence interval with a width equal to 29.463 when the sample variance is 24.000.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size vs. the confidence interval width for the two confidence levels.

Example 2 – Validation using Zar

Zar (1984) page 115 give an example of a calculation for a confidence interval on the variance when the confidence level is 95%, the variance is 18.0388, and the interval width is 23.91244. The necessary sample size is 25.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Variance using Variance** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **One Variance using Variance**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	23.91244
Variance	18.0388

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Confidence Level	Sample Size (N)	Target Width	Actual Width	Variance	Lower Limit	Upper Limit
0.950	25	23.912	23.912	18.039	10.998	34.911

PASS also calculated the necessary sample size to be 25.

Chapter 652

Confidence Intervals for One Variance with Tolerance Probability

Introduction

This procedure calculates the sample size necessary to achieve a specified width (or in the case of one-sided intervals, the distance from the variance to the confidence limit) with a given tolerance probability at a stated confidence level for a confidence interval about a single variance when the underlying data distribution is normal.

Technical Details

For a single variance from a normal distribution with unknown mean, a two-sided, $100(1 - \alpha)\%$ confidence interval is calculated by

$$\left[\frac{(n-1)s^2}{\chi^2_{1-\alpha/2, n-1}}, \frac{(n-1)s^2}{\chi^2_{\alpha/2, n-1}} \right]$$

A one-sided $100(1 - \alpha)\%$ upper confidence limit is calculated by

$$\frac{(n-1)s^2}{\chi^2_{\alpha, n-1}}$$

Similarly, the one-sided $100(1 - \alpha)\%$ lower confidence limit is

$$\frac{(n-1)s^2}{\chi^2_{1-\alpha, n-1}}$$

652-2 Confidence Intervals for One Variance with Tolerance Probability

For two-sided intervals, the distance from the variance to each of the limits is different. Thus, instead of specifying the distance to the limits we specify the width of the interval, W .

The basic equation for determining sample size for a two-sided interval when W has been specified is

$$W = \frac{(n-1)s^2}{\chi_{\alpha/2, n-1}^2} - \frac{(n-1)s^2}{\chi_{1-\alpha/2, n-1}^2}$$

For one-sided intervals, the distance from the variance to limits, D , is specified.

The basic equation for determining sample size for a one-sided upper limit when D has been specified is

$$D = \frac{(n-1)s^2}{\chi_{\alpha, n-1}^2} - s^2$$

The basic equation for determining sample size for a one-sided lower limit when D has been specified is

$$D = s^2 - \frac{(n-1)s^2}{\chi_{1-\alpha, n-1}^2}$$

These equations can be solved for any of the unknown quantities in terms of the others.

There is an additional subtlety that arises when the variance is to be chosen for estimating sample size. The sample sizes determined from the formulas above produce confidence intervals with the specified widths only when the future sample has a sample variance that is no greater than the value specified.

As an example, suppose that 15 individuals are sampled in a pilot study, and a variance estimate of 6.5 is obtained from the sample. The purpose of a later study is to estimate the variance with a confidence interval with width no greater than 3 units. Suppose further that the sample size needed is calculated to be 105 using the formula above with 6.5 as the estimate for the variance. The sample of size 105 is then obtained from the population, but the variance of the 105 individuals turns out to be 7.2 rather than 6.5. The confidence interval is computed and the width of the interval is greater than 3 units.

This example illustrates the need for an adjustment to adjust the sample size such that the width or distance from the variance to the confidence limits will be below the specified value with known probability.

Such an adjustment for situations where a previous sample is used to estimate the variance is derived for the case of confidence intervals for a mean by Harris, Horvitz, and Mood (1948) and discussed in Zar (1984) and Hahn and Meeker (1991). The adjustment is made by replacing s^2 with

$$s^2 = \sigma^2 F_{1-\gamma; n-1, m-1}$$

where $1 - \gamma$ is the probability that the width or distance from the variance to the confidence limit will be below the specified value, and m is the sample size in the previous sample that was used to estimate the variance.

The corresponding adjustment when no previous sample is available is discussed in Kupper and Hafner (1989) and Hahn and Meeker (1991). The adjustment in this case is

$$s^2 = \sigma^2 \frac{\chi_{1-\gamma, n-1}^2}{n-1}$$

where, again, $1 - \gamma$ is the probability that the width or distance from the variance to the confidence limit will be below the specified value.

Each of these adjustments accounts for the variability in a future estimate of the variance. In the first adjustment formula (Harris, Horvitz, and Mood, 1948), the distribution of the variance is based on the estimate from a previous sample. In the second adjustment formula, the distribution of the variance is based on a specified value that is assumed to be the population variance.

For this procedure, both adjustments are adapted from the case of a one-sample confidence interval for a single mean to the case of a one-sample confidence interval for a single variance.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population variance is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence and Tolerance

Confidence Level (1 – Alpha)

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population variance is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as *0.90, 0.95* or *0.90 to 0.99 by 0.01*.

652-4 Confidence Intervals for One Variance with Tolerance Probability

Tolerance Probability

This is the probability that a future interval with sample size N and the specified confidence level will have a width or distance from the variance to the limit that is less than or equal to the width or distance specified.

If a tolerance probability is not used, as in the 'Confidence Intervals for One Variance using Variance' procedure, the sample size is calculated for the expected width or distance from the variance to the limit, which assumes that the future variance will also be the one specified.

Using a tolerance probability implies that the variance of the future sample will not be known in advance, and therefore, an adjustment is made to the sample size formula to account for the variability in the variance. Use of a tolerance probability is similar to using an upper bound for the variance in the 'Confidence Intervals for One Variance using Variance' procedure.

Values between 0 and 1 can be entered. The choice of the tolerance probability depends upon how important it is that the width or distance from the interval limits to the variance is at most the value specified.

You can enter a range of values such as *0.70 0.80 0.90* or *0.70 to 0.95 by 0.05*.

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

Precision

Confidence Interval Width (Two-Sided)

This is the distance from the lower confidence limit to the upper confidence limit. The distance from the variance to the lower and upper limits is not equal.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Distance from Var to Limit (One-Sided)

This is the distance from the variance to the lower or upper limit of the confidence interval, depending on whether the Interval Type is set to Lower Limit or Upper Limit.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Variance

Variance Source

This procedure permits two sources for estimates of the variance:

- **V is a Population Variance**

This option should be selected if there is no previous sample that can be used to obtain an estimate of the variance. In this case, the algorithm assumes that the future sample obtained will be from a population with variance V .

- **V from a Previous Sample**

This option should be selected if the estimate of the variance is obtained from a previous random sample from the same distribution as the one to be sampled. The sample size of the previous sample must also be entered under 'Sample Size of Previous Sample'.

Variance – V is a Population Variance

V (Variance)

Enter an estimate of the variance (must be positive). In this case, the algorithm assumes that future samples obtained will be from a population with variance V .

You can enter a range of values such as *1 2 3* or *1 to 10 by 1*.

Press the Standard Deviation Estimator button to load the Standard Deviation Estimator window.

Variance – V from a Previous Sample

V (Var Estimated from a Previous Sample)

Enter an estimate of the variance from a previous (or pilot) study. This value must be positive.

A range of values may be entered.

Press the Standard Deviation Estimator button to load the Standard Deviation Estimator window.

Sample Size of Previous Sample

Enter the sample size that was used to estimate the variance entered in V (Var Estimated from a Previous Sample).

This value is entered only when 'Variance Source:' is set to 'V from a Previous Sample'.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

A researcher would like to estimate the variance of a population with 95% confidence. It is very important that the interval width is less than 10 grams. Data available from a previous study are used to provide an estimate of the variance. The estimate of the variance is 35.4 grams, from a sample of size 16.

The goal is to determine the sample size necessary to obtain a two-sided confidence interval such that the width of the interval is less than 10 grams. Tolerance probabilities of 0.70 to 0.95 will be examined.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Variance with Tolerance Probability** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **One Variance with Tolerance Probability**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
Tolerance Probability	0.70 to 0.95 by 0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width	10
Variance Source	V from a Previous Sample
V	35.4
Sample Size of Previous Sample	16

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals

Confidence Level	Sample Size (N)	Target Width	Actual Width	Variance	Tolerance Probability
0.950	641	10.000	9.995	35.400	0.700
0.950	722	10.000	9.998	35.400	0.750
0.950	827	10.000	9.998	35.400	0.800
0.950	973	10.000	9.995	35.400	0.850
0.950	1200	10.000	9.996	35.400	0.900
0.950	1658	10.000	9.999	35.400	0.950

Sample size for estimate of the variance from previous sample = 16.

References

Hahn, G. J. and Meeker, W.Q. 1991. Statistical Intervals. John Wiley & Sons. New York.
 Zar, J. H. 1984. Biostatistical Analysis. Second Edition. Prentice-Hall. Englewood Cliffs, New Jersey.
 Harris, M., Horvitz, D. J., and Mood, A. M. 1948. 'On the Determination of Sample Sizes in Designing Experiments', Journal of the American Statistical Association, Volume 43, No. 243, pp. 391-402.

Report Definitions

Confidence Level is the proportion of confidence intervals (constructed with this same confidence level, sample size, etc.) that would contain the population variance.

N is the size of the sample drawn from the population.

Width is the distance from the lower limit to the upper limit.

Target Width is the value of the width that is entered into the procedure.

Actual Width is the value of the width that is obtained from the procedure.

Variance is the estimated variance based on a previous sample.

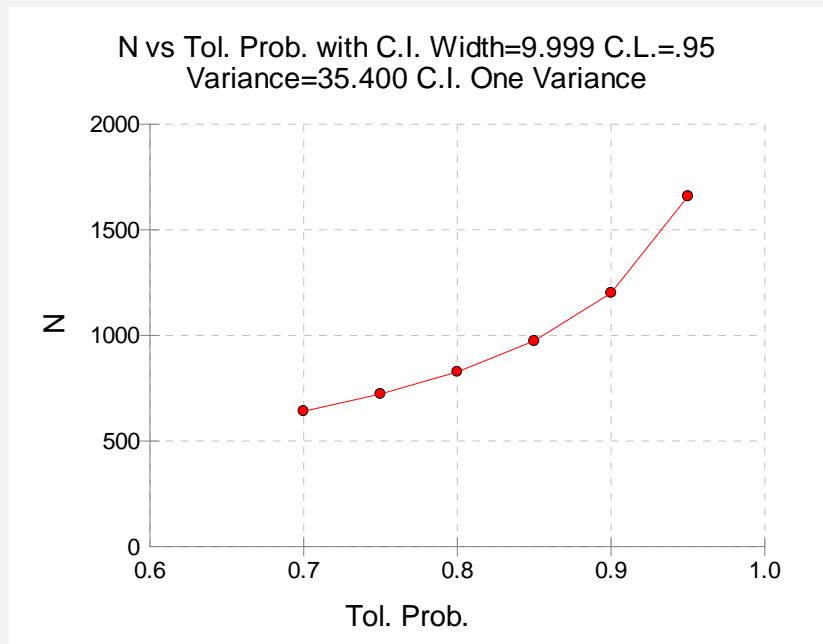
Tolerance Probability is the probability that a future interval with sample size N and corresponding confidence level will have a width that is less than or equal to the specified width.

Summary Statements

The probability is 0.700 that a sample size of 641 will produce a two-sided 95% confidence interval with a width that is less than or equal to 9.995 if the population variance is estimated to be 35.400 by a previous sample of size 16.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the tolerance probability.

Example 2 – Validation using Simulation

We could not find a published result for confidence intervals for a variance with tolerance probability. This procedure is validated using a simulation. A simulation was run with a confidence level of 95%, sample size of 57, a specified confidence interval width of 0.989, and assumed population variance of 1.

The number of simulations was 100,000. The proportion of intervals with width less than 0.989 (tolerance probability) was 0.90009.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Variance with Tolerance Probability** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **One Variance with Tolerance Probability**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Tolerance Probability
Confidence Level	0.95
Tolerance Probability	<i>Ignored since this is the Find setting</i>
N (Sample Size)	57
Interval Type	Two-Sided
Confidence Interval Width	6
Variance Source	V is a Population Variance
V	1

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals					
Confidence Level	Sample Size (N)	Target Width	Actual Width	Variance	Tolerance Probability
0.950	57	0.989	0.989	1.000	0.900

PASS calculated the tolerance probability to be 0.900, which is well within the simulation error of the simulation estimate of 0.90009.

Chapter 653

Confidence Intervals for One Variance using Relative Error

Introduction

This routine calculates the necessary sample size such that a sample variance estimate will achieve a specified relative distance from the true population variance at a stated confidence level when the underlying data distribution is normal.

Caution: This procedure controls the relative width of the interval as a proportion of the true population variance. For controlling the absolute width of the interval see the procedures Confidence Intervals for One Variance using Variance and Confidence Intervals for One Variance with Tolerance Probability.

Technical Details

Following the results of Desu and Raghavarao (1990) and Greenwood and Sandomire (1950), let s^2 be the variance estimate based on a sample from a normal distribution with unknown μ and unknown σ^2 . Let r be the proportion of σ^2 such that s^2 is within $r\sigma^2$ of σ^2 with desired confidence. If

$$p_1 = \Pr(s^2 > \sigma^2 + r\sigma^2) = \Pr(s^2 > \sigma^2(1 + r))$$

and

$$p_2 = \Pr(s^2 < \sigma^2 - r\sigma^2) = \Pr(s^2 < \sigma^2(1 - r))$$

The confidence level for estimating σ^2 within proportion r is $1 - p_1 - p_2$.

653-2 Confidence Intervals for One Variance using Relative Error

Since $\frac{(n-1)s^2}{\sigma^2} \sim \chi_{n-1}^2$, it is useful to rewrite p_1 and p_2 as

$$p_1 = \Pr\left(\frac{(n-1)s^2}{\sigma^2} > (n-1)(1+r)\right)$$

and

$$p_2 = \Pr\left(\frac{(n-1)s^2}{\sigma^2} < (n-1)(1-r)\right)$$

Using the chi-square distribution, these equations can be solved for any of the unknown quantities (n , r , $p_1 + p_2$) in terms of the others.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and the variance estimates are obtained from each sample, the proportion of those estimates that are within $r\sigma^2$ of σ^2 is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and the variance estimates are obtained from each sample, the proportion of those estimates that are within $r\sigma^2$ of σ^2 is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as 0.90, 0.95 or 0.90 to 0.99 by 0.01.

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

In each case the limits are based on the relative error of the true population variance.

Precision

Relative Error

This is the distance from the true variance as a proportion of the true variance.

You can enter a single value or a list of values. The value(s) must be between 0 and 100000.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to be 95% confident that estimated variance is within 10% of the true population variance. In addition to 10% relative error, 5%, 15%, 20% and 25% will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Variance using Relative Error** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **One Variance using Relative Error**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Relative Error	0.05 to 0.25 by 0.05

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Relative Error Confidence Intervals

Target Confidence Level	Actual Confidence Level	Sample Size (N)	Relative Error
0.950	0.950	3074	0.050
0.950	0.950	769	0.100
0.950	0.950	342	0.150
0.950	0.950	192	0.200
0.950	0.950	123	0.250

References

Desu, M. M. and Raghavarao, D. 1990. Sample Size Methodology. Academic Press. New York.
Greenwood, J. A. and Sandomire, M. M. 1950. 'Sample Size Required for Estimating the Standard Deviation as a Per Cent of its True Value', Journal of the American Statistical Association, Vol. 45, No. 250, pp. 257-260.

Report Definitions

Confidence Level is the proportion of variance estimates that will be within the relative error of the true variance.

Target Confidence Level is the value of the confidence level that is entered into the procedure.

Actual Confidence Level is the value of the confidence level that is obtained from the procedure.

Sample Size (N) is the size of the sample drawn from the population.

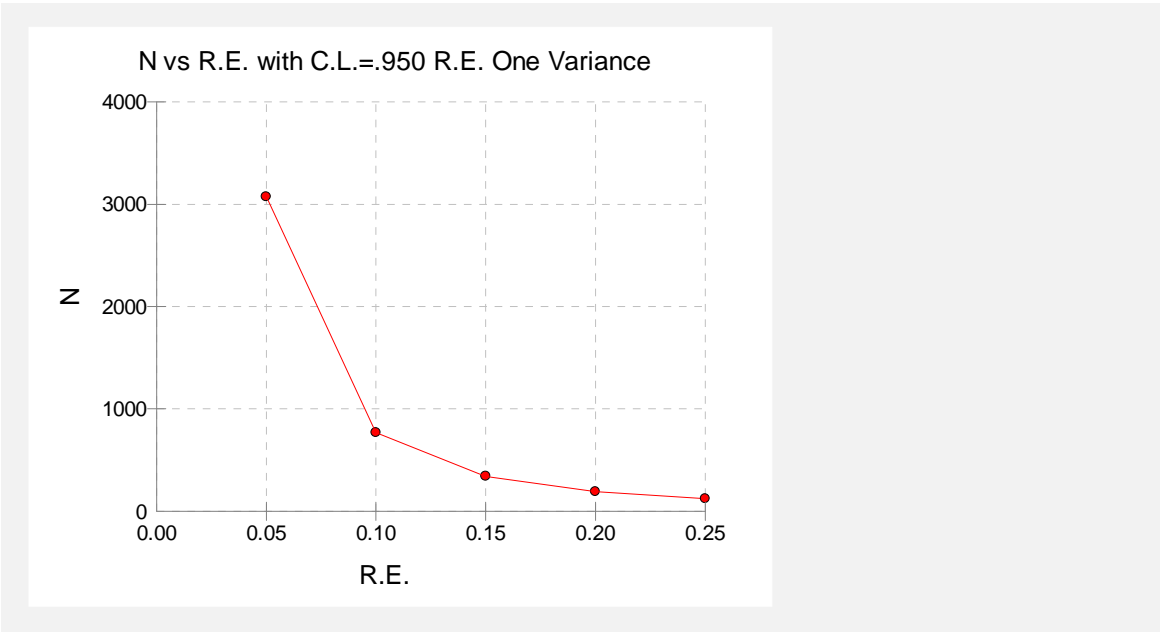
Relative Error is the distance from the true variance as a proportion of the true variance.

Summary Statements

With a sample size of 3074, the probability is 0.950 (95% confidence) that the estimate of the variance will be within 5% of the true population variance.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the relative error.

Example 2 – Validation using Desu and Raghavarao

Desu and Raghavarao (1990) page 6 give an example of a calculation in which the desired confidence level is 95% and the relative error is 20%. This calculation is based on a large sample approximation formula. The necessary sample size is 194.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Variance using Relative Error** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **One Variance using Relative Error**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Relative Error	0.2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Target Confidence Level	Actual Confidence Level	Sample Size (N)	Relative Error
0.950	0.950	192	0.200

PASS calculated the necessary sample size to be 192, but did not use the approximation formulas.

Using direct calculation with the non-approximate chi-square formulas with a sample size of 194, the confidence level is $1 - (0.0300164564 + 0.0188126383) = 0.951171$.

Using direct calculation with the non-approximate chi-square formulas with a sample size of 192, the confidence level is $1 - (0.0306634808 + 0.0193315700) = 0.950005$, which is closer to the prescribed confidence level. Thus, 192 is the correct value.

Chapter 655

Inequality Tests for Two Variances

Introduction

Occasionally, researchers are interested in comparing the variances (or standard deviations) of two groups rather than their means. This module calculates the sample sizes and performs power analyses for hypothesis tests concerning two variances.

Technical Details

Assuming that variables X_1 and X_2 are normally distributed variances σ_1^2 and σ_2^2 (the means are ignored), the distribution of the ratio of the sample variances follows the F distribution. That is,

$$F = \frac{s_1^2}{s_2^2}$$

is distributed as an F random variable with $N_1 - 1$ and $N_2 - 1$ degrees of freedom. The sample statistic, s_j^2 , is calculated as follows

$$s_j^2 = \frac{\sum_{i=1}^N (X_{ji} - \bar{X}_j)^2}{N_j - 1}.$$

The power or sample size of a hypothesis test about the variance can be calculated using the appropriate one of the following three formulas:

Case 1: $H_0: \sigma_1^2 = \sigma_2^2$ versus $H_a: \sigma_1^2 \neq \sigma_2^2$

$$\beta = P\left(\frac{\sigma_1^2}{\sigma_2^2} F_{\alpha/2, N_1-1, N_2-1} < F < \frac{\sigma_1^2}{\sigma_2^2} F_{1-\alpha/2, N_1-1, N_2-1}\right)$$

Case 2: $H_0: \sigma_1^2 = \sigma_2^2$ versus $H_a: \sigma_1^2 > \sigma_2^2$

$$\beta = P\left(F > \frac{\sigma_1^2}{\sigma_2^2} F_{\alpha, N_1-1, N_2-1}\right)$$

655-2 Inequality Tests for Two Variances

Case 3: $H_0: \sigma_1^2 = \sigma_2^2$ versus $H_a: \sigma_1^2 < \sigma_2^2$

$$\beta = P\left(F < \frac{\sigma_1^2}{\sigma_2^2} F_{1-\alpha, N1-1, N2-1}\right)$$

Procedure Options

This section describes the options that are unique to this procedure. These are located on the panels associated with the Data and Iterations tabs. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

If your only interest is in determining the appropriate sample size for a confidence interval, set power or beta to 0.5.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when a true null hypothesis is rejected.

Values must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N1 (Sample Size Group 1)

Enter a value (or range of values) for the sample size of this group. Note that these values are ignored when you are solving for $N1$. You may enter a range of values such as *10 to 100 by 10*.

N2 (Sample Size Group 2)

Enter a value (or range of values) for the sample size of group 2 or enter *Use R* to base $N2$ on the value of $N1$. You may enter a range of values such as *10 to 100 by 10*.

- **Use R**

When *Use R* is entered, $N2$ is calculated using the formula

$$N2 = [R N1]$$

where R is the Sample Allocation Ratio and $[Y]$ is the first integer greater than or equal to Y . For example, if you want $N1 = N2$, select *Use R* here and set $R = 1$.

R (Sample Allocation Ratio)

Enter a value (or range of values) for R , the allocation ratio between samples. This value is only used when $N2$ is set to *Use R*.

When used, $N2$ is calculated from $N1$ using the formula: $N2 = [R N1]$ where $[Y]$ is the next integer greater than or equal to Y . Note that setting $R = 1.0$ forces $N2 = N1$.

Effect Size

Scale

Specify whether $V1$ and $V2$ are variances or standard deviations.

V1 and V2

Enter one or more value(s) for the variances of the groups, σ_1^2 and σ_2^2 . All entries must be greater than zero.

Note that since the ratio of these variances is all that is used in the power equations, you can specify the problem in terms of the variance ratio instead of the two variances. To do this, enter 1.0 for $V2$ and enter the desired variance ratio in $V1$.

If Scale is *Standard Deviation* this value is the standard deviation rather than the variance.

Test

Alternative Hypothesis

This option specifies the alternative hypothesis. This implicitly specifies the direction of the hypothesis test. The null hypothesis is always $H_0: V_1 = V_2$.

Note that the alternative hypothesis enters into power calculations by specifying the rejection region of the hypothesis test. Its accuracy is critical.

655-4 Inequality Tests for Two Variances

Possible selections are:

- **Ha: $V1 \neq V2$**

This selection yields a *two-tailed* test. Use this option when you are testing whether the variances are different but you do not want to specify beforehand which variance is larger.

- **Ha: $V1 > V2$**

The options yields a *one-tailed* test. Use it when you are only interested in the case in which $V2$ is less than $V1$.

- **Ha: $V1 < V2$**

This option yields a *one-tailed* test. Use it when you are only interested in the case in which $V2$ is greater than $V1$.

Iterations Tab

This tab sets the option used in the iterative procedure.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating the Power

A machine used to perform a particular analysis is to be replaced with a new type of machine if the new machine reduces the variance in the output by 50%. If the significance level is set to 0.05, calculate the power for sample sizes of 5, 10, 20, 35, 50, 90, 130, and 200.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Variances** procedure window by clicking on **Variances**, then **Two Variances**, then **Inequality Tests**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

Option

Value

Data Tab

Find (Solve For) **Power and Beta**
Power *Ignored since this is the Find setting*
Alpha **0.01 0.05**
N1 (Sample Size Group 1) **5 10 20 35 50 90 130 200**
N2 (Sample Size Group 2) **Use R**
R (Sample Allocation Ratio) **1.0**

Data Tab (continued)

Scale **Variance**
 V1 (Variance of Group 1) **1.0**
 V2 (Variance of Group 2) **0.5**
 Alternative Hypothesis **Ha: V1 > V2**

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results**Numeric Results when H0: V1 = V2 versus Ha: V1>V2**

Power	N1	N2	V1	V2	Alpha	Beta
0.034378	5	5	1.0000	0.5000	0.010000	0.965622
0.079391	10	10	1.0000	0.5000	0.010000	0.920609
0.187122	20	20	1.0000	0.5000	0.010000	0.812878
0.362650	35	35	1.0000	0.5000	0.010000	0.637350
0.526210	50	50	1.0000	0.5000	0.010000	0.473790
0.821599	90	90	1.0000	0.5000	0.010000	0.178401
0.944261	130	130	1.0000	0.5000	0.010000	0.055739
0.994526	200	200	1.0000	0.5000	0.010000	0.005474
0.143437	5	5	1.0000	0.5000	0.050000	0.856563
0.250418	10	10	1.0000	0.5000	0.050000	0.749582
0.431043	20	20	1.0000	0.5000	0.050000	0.568957
0.636863	35	35	1.0000	0.5000	0.050000	0.363137
0.776507	50	50	1.0000	0.5000	0.050000	0.223493
0.946023	90	90	1.0000	0.5000	0.050000	0.053977
0.988497	130	130	1.0000	0.5000	0.050000	0.011503
0.999364	200	200	1.0000	0.5000	0.050000	0.000636

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N1 is the size of the sample drawn from the population 1.

N2 is the size of the sample drawn from the population 2.

V1 is the value of the population variance of group 1.

V2 is the value of the population variance of group 2.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

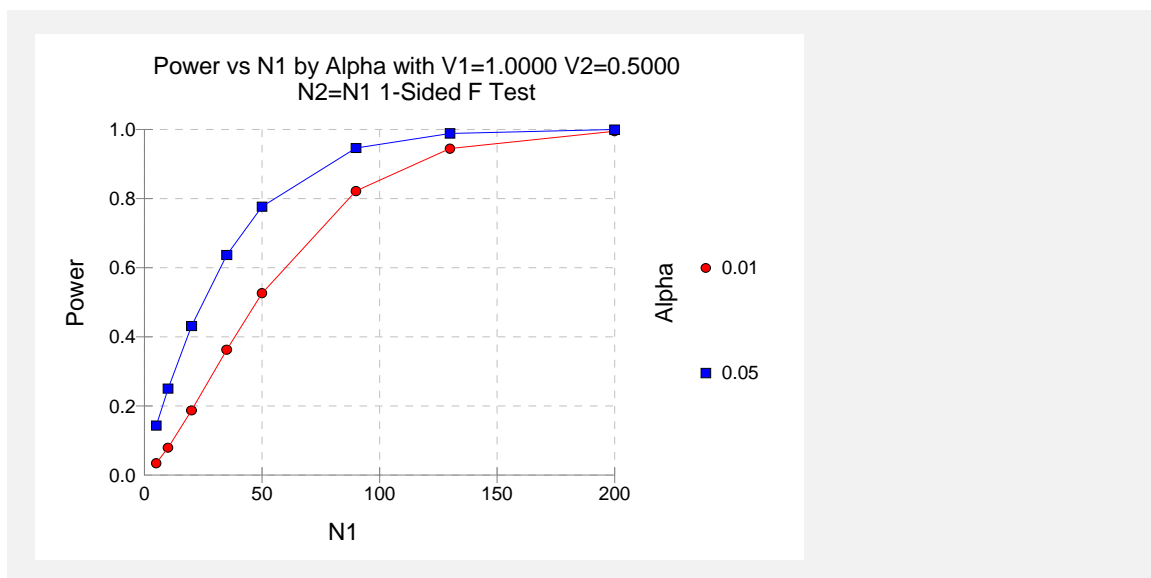
Beta is the probability of accepting a false null hypothesis. It should be small.

Summary Statements

Group sample sizes of 5 and 5 achieve 3% power to detect a ratio of 2.0000 between the group one variance of 1.0000 and the group two variance of 0.5000 using a one-sided F test with a significance level (alpha) of 0.010000.

This report shows the calculated power for each scenario.

Plots Section



This plot shows the power versus the sample size for the two significance levels. It is now easy to determine an appropriate sample size to meet both the alpha and beta objectives of the study.

Example 2 – Calculating Sample Size

Continuing with the previous example, the analyst wants to find the necessary sample sizes to achieve a power of 0.9 for two significance levels, 0.01 and 0.05, and for several variance ratio values of 0.2, 0.3, 0.4, 0.5, 0.6, and 0.7.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Variances** procedure window by clicking on **Variances**, then **Two Variances**, then **Inequality Tests**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N1
Power	0.90
Alpha	0.05
N1 (Sample Size Group 1)	<i>Ignored since this is the Find setting</i>
N2 (Sample Size Group 2)	Use R
R (Sample Allocation Ratio)	1.0
Scale	Variance
V1 (Variance of Group 1)	1.0
V2 (Variance of Group 2)	0.2 to 0.7 by 0.1
Alternative Hypothesis	Ha: V1 > V2

Output

Click the Run button to perform the calculations and generate the following output.

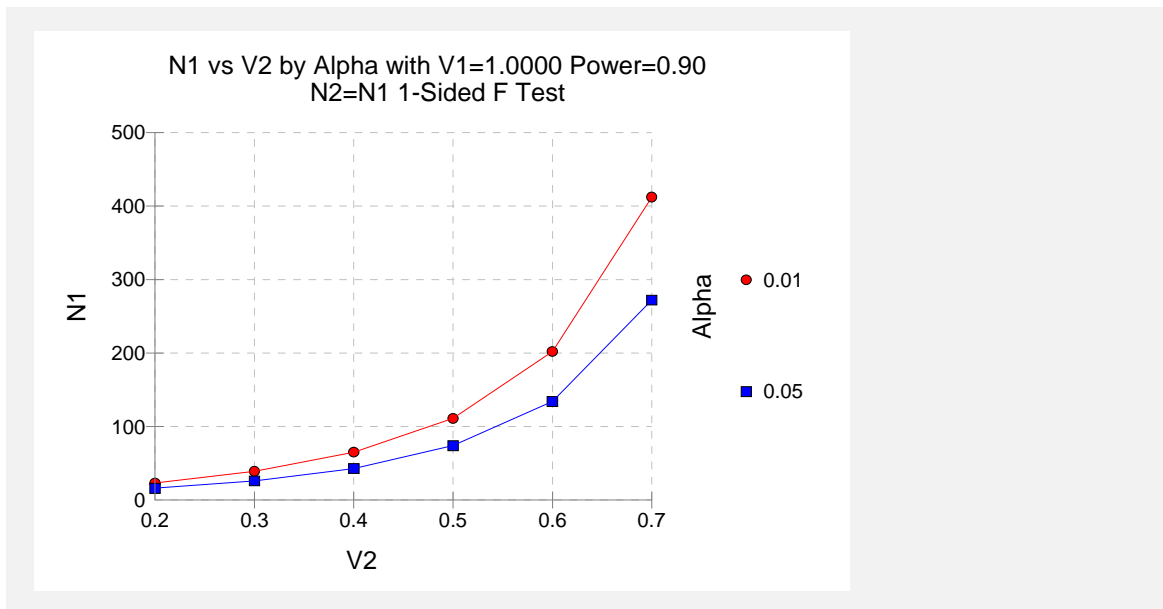
Numeric Results

Numeric Results when $H_0: V_1 = V_2$ versus $H_a: V_1 > V_2$

Power	N1	N2	V1	V2	Alpha	Beta
0.911109	23	23	1.0000	0.2000	0.010000	0.088891
0.916243	16	16	1.0000	0.2000	0.050000	0.083757
0.907774	39	39	1.0000	0.3000	0.010000	0.092226
0.905324	26	26	1.0000	0.3000	0.050000	0.094676
0.904075	65	65	1.0000	0.4000	0.010000	0.095925
0.902069	43	43	1.0000	0.4000	0.050000	0.097931
0.901318	111	111	1.0000	0.5000	0.010000	0.098682
0.902949	74	74	1.0000	0.5000	0.050000	0.097051
0.900454	202	202	1.0000	0.6000	0.010000	0.099546
0.901654	134	134	1.0000	0.6000	0.050000	0.098346
0.900417	412	412	1.0000	0.7000	0.010000	0.099583
0.900817	272	272	1.0000	0.7000	0.050000	0.099183

This report shows the necessary sample size for each scenario.

Plot Section



This plot shows the necessary sample size for various values of V_2 . Note that as V_2 nears V_1 , the sample size is increased.

Example 3 – Validation using Davies

Davies (1971) page 41 presents an example with $V1 = 4$, $V2 = 1$, $Alpha = 0.05$, and $Power = 0.99$ in which the sample sizes, $N1$ and $N2$, are calculated to be 36. We will run this example through *PASS*.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the *PASS* Home window, load the **Inequality Tests for Two Variances** procedure window by clicking on **Variances**, then **Two Variances**, then **Inequality Tests**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N1
Power	0.99
Alpha	0.05
N1 (Sample Size Group 1)	<i>Ignored since this is the Find setting</i>
N2 (Sample Size Group 2)	Use R
R (Sample Allocation Ratio)	1.0
Scale	Variance
V1 (Variance of Group 1)	4
V2 (Variance of Group 2)	1
Alternative Hypothesis	Ha: V1 > V2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H0: V1 = V2 versus Ha: V1>V2

Power	N1	N2	V1	V2	Alpha	Beta
0.991423	36	36	4.0000	1.0000	0.050000	0.008577

PASS calculated the $N1$ and $N2$ to be 36 which matches Davies' result.

Chapter 656

Confidence Intervals for the Ratio of Two Variances using Variances

Introduction

This routine calculates the group sample sizes necessary to achieve a specified interval width or distance from the variance ratio to the confidence limit at a stated confidence level for a confidence interval about the variance ratio when the underlying data distribution is normal.

Caution: This procedure assumes that the variances of the future samples will be the same as the variances that are specified. The Confidence Intervals for the Ratio of Two Variances using Relative Error controls the width or distance from the variance ratio to the limits by controlling the width or distance as a percent of the true variance ratio.

Technical Details

For a ratio of two variances from normal distributions, a two-sided, $100(1 - \alpha)\%$ confidence interval is calculated by

$$\left[\frac{s_1^2}{s_2^2} \frac{1}{F_{\alpha/2, n_1-1, n_2-1}}, \frac{s_1^2}{s_2^2} F_{\alpha/2, n_2-1, n_1-1} \right]$$

A one-sided $100(1 - \alpha)\%$ upper confidence limit is calculated by

$$\frac{s_1^2}{s_2^2} F_{\alpha, n_2-1, n_1-1}$$

656-2 Confidence Intervals for the Ratio of Two Variances using Variances

Similarly, the one-sided $100(1 - \alpha)\%$ lower confidence limit is

$$\frac{s_1^2}{s_2^2} \frac{1}{F_{\alpha, n_1-1, n_2-1}}$$

For two-sided intervals, the distance from the variance ratio to each of the limits is different. Thus, instead of specifying the distance to the limits we specify the width of the interval, W .

The basic equation for determining sample size for a two-sided interval when W has been specified is

$$W = \frac{s_1^2}{s_2^2} F_{\alpha/2, n_2-1, n_1-1} - \frac{s_1^2}{s_2^2} \frac{1}{F_{\alpha/2, n_1-1, n_2-1}}$$

For one-sided intervals, the distance from the variance ratio to limit, D , is specified.

The basic equation for determining sample size for a one-sided upper limit when D has been specified is

$$D = \frac{s_1^2}{s_2^2} F_{\alpha, n_2-1, n_1-1} - \frac{s_1^2}{s_2^2}$$

The basic equation for determining sample size for a one-sided lower limit when D has been specified is

$$D = \frac{s_1^2}{s_2^2} - \frac{s_1^2}{s_2^2} \frac{1}{F_{\alpha, n_1-1, n_2-1}}$$

These equations can be solved for any of the unknown quantities in terms of the others.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of random samples of size n_1 and n_2 are drawn from populations 1 and 2, respectively, and a confidence interval for the variance ratio is calculated for each pair of samples, the proportion of those intervals that will include the true variance ratio is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of random samples of size n_1 and n_2 are drawn from populations 1 and 2, respectively, and a confidence interval for the variance ratio is calculated for each pair of samples, the proportion of those intervals that will include the true variance ratio is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as *0.90, 0.95 or 0.90 to 0.99 by 0.01*.

Sample Size

N1 (Sample Size Group 1)

Enter a value (or range of values) for the sample size of this group. Note that these values are ignored when you are solving for $N1$. You may enter a range of values such as *10 to 100 by 10*.

N2 (Sample Size Group 2)

Enter a value (or range of values) for the sample size of group 2 or enter *Use R* to base $N2$ on the value of $N1$. You may enter a range of values such as *10 to 100 by 10*.

- **Use R**

When *Use R* is entered here, $N2$ is calculated using the formula

$$N2 = [R(N1)]$$

where R is the Sample Allocation Ratio and the operator $[Y]$ is the first integer greater than or equal to Y . For example, if you want $N1 = N2$, select *Use R* and set $R = 1$.

R (Sample Allocation Ratio)

Enter a value (or range of values) for R , the allocation ratio between samples. This value is only used when $N2$ is set to *Use R*.

When used, $N2$ is calculated from $N1$ using the formula: $N2 = [R(N1)]$ where $[Y]$ is the next integer greater than or equal to Y . Note that setting $R = 1.0$ forces $N2 = N1$.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

Precision

Confidence Interval Width (Two-Sided)

This is the distance from the lower confidence limit to the upper confidence limit. The distance from the variance ratio to the lower and upper limits is not equal.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Distance from Ratio to Limit (One-Sided)

This is the distance from the variance ratio to the lower or upper limit of the confidence interval, depending on whether the Interval Type is set to Lower Limit or Upper Limit.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Variances for Variance Ratio (V1/V2)

V1 (Variance Group 1)

Enter an estimate of the variance for group 1 (must be positive). The sample size and width calculations assume that the value entered here is the variance estimate that is obtained from the sample. If the sample variance is different from the one specified here, the width may be narrower or wider than specified.

For confidence intervals with widths that are specified in terms of a percentage of relative error, see the procedure 'Confidence Intervals for the Ratio of Two Variances using Relative Error'.

You can enter a range of values such as *1 2 3* or *1 to 10 by 1*.

V2 (Variance Group 2)

Enter an estimate of the variance for group 2 (must be positive). The sample size and width calculations assume that the value entered here is the variance estimate that is obtained from the sample. If the sample variance is different from the one specified here, the width may be narrower or wider than specified.

For confidence intervals with widths that are specified in terms of a percentage of relative error, see the procedure 'Confidence Intervals for the Ratio of Two Variances using Relative Error'.

You can enter a range of values such as *1 2 3* or *1 to 10 by 1*.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to construct a two-sided 95% confidence interval for the variance ratio such that the width of the interval is no wider than 0.5. The confidence level is set at 0.95, but 0.99 is included for comparative purposes. The variance estimates to be used are 5 for Group 1, and 10 for Group 2. Instead of examining only the interval width of 0.5, a series of widths from 0.2 to 0.8 will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for the Ratio of Two Variances using Variance** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **Ratio of Two Variances using Variances**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N1
Confidence Level	0.95 0.99
N1 (Sample Size Group 1).....	<i>Ignored since this is the Find setting</i>
N2 (Sample Size Group 2).....	Use R
R (Sample Allocation Ratio).....	1.0
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	0.2 to 0.8 by 0.1
V1	5
V2	10

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals for the Variance Ratio

Confidence Level	N1	N2	Allocation Ratio	Target Width	Actual Width	V1	V2	Lower Limit	Upper Limit
0.950	392	392	1.000	0.200	0.200	5.00	10.00	0.41	0.61
0.950	178	178	1.000	0.300	0.300	5.00	10.00	0.37	0.67
0.950	104	104	1.000	0.400	0.398	5.00	10.00	0.34	0.74
0.950	69	69	1.000	0.500	0.498	5.00	10.00	0.31	0.81
0.950	50	50	1.000	0.600	0.597	5.00	10.00	0.28	0.88
0.950	39	39	1.000	0.700	0.691	5.00	10.00	0.26	0.95
0.950	31	31	1.000	0.800	0.796	5.00	10.00	0.24	1.04
0.990	675	675	1.000	0.200	0.200	5.00	10.00	0.41	0.61
0.990	307	307	1.000	0.300	0.300	5.00	10.00	0.37	0.67
0.990	178	178	1.000	0.400	0.399	5.00	10.00	0.34	0.74
0.990	118	118	1.000	0.500	0.498	5.00	10.00	0.31	0.81
0.990	85	85	1.000	0.600	0.598	5.00	10.00	0.28	0.88
0.990	65	65	1.000	0.700	0.699	5.00	10.00	0.26	0.96
0.990	53	53	1.000	0.800	0.791	5.00	10.00	0.24	1.03

References

Ostle, B. and Malone, L.C. 1988. Statistics in Research. Iowa State University Press. Ames, Iowa.
 Zar, Jerrold H. 1984. Biostatistical Analysis (Second Edition). Prentice-Hall. Englewood Cliffs, New Jersey.

Report Definitions

Confidence level is the proportion of confidence intervals (constructed with this same confidence level, sample size, etc.) that would contain the true variance ratio.

N1 and N2 are the sample sizes drawn from the two populations.

Allocation Ratio is the ratio of the sample sizes, N2/N1.

Width is the distance from the lower limit to the upper limit.

Target Width is the value of the width that is entered into the procedure.

Actual Width is the value of the width that is obtained from the procedure.

V1 and V2 are the assumed sample variances upon which the width calculations are based.

Lower Limit is the lower limit of the confidence interval.

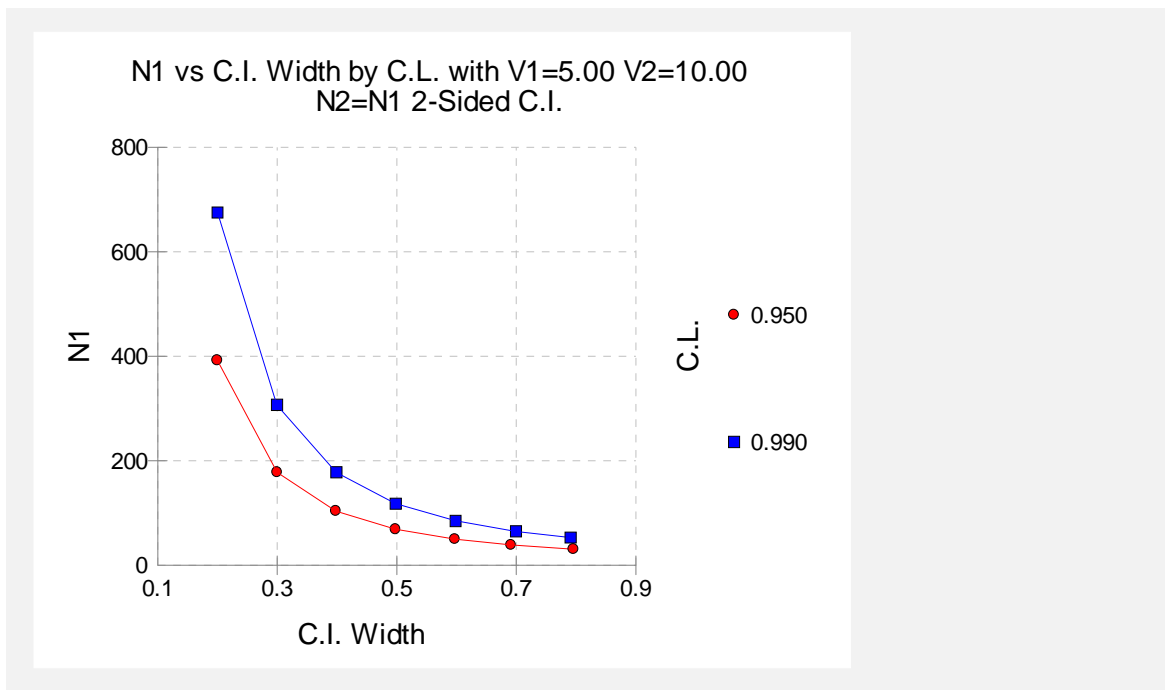
Upper Limit is the upper limit of the confidence interval.

Summary Statements

Group sample sizes of 392 and 392 produce a two-sided 95% confidence interval with a width that is equal to 0.200 when the estimated numerator variance is 5.00 and the estimated denominator variance is 10.00.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the group sample size versus the confidence interval width for the two confidence levels.

Example 2 – Validation using Sachs

Sachs (1984) page 261 gives an example of a calculation for a confidence interval for the variance ratio when the confidence level is 90%, the variances are 8 and 3, and the interval width is 4.56. The necessary sample size is 20 per group.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for the Ratio of Two Variances using Variance** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **Ratio of Two Variances using Variances**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N1
Confidence Level	0.90
N1 (Sample Size Group 1)	<i>Ignored since this is the Find setting</i>
N2 (Sample Size Group 2)	Use R
R (Sample Allocation Ratio)	1.0
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	4.56
V1	8
V2	3

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Confidence Level	N1	N2	Allocation Ratio	Target Width	Actual Width	V1	V2	Lower Limit	Upper Limit
0.900	20	20	1.000	4.560	4.552	8.00	3.00	1.23	5.78

PASS also calculated the necessary sample size to be 20 per group.

Chapter 657

Confidence Intervals for the Ratio of Two Variances using Relative Error

Introduction

This routine calculates the necessary sample size such that a variance ratio estimate from two independent samples will achieve a specified relative distance from the true variance ratio at a stated confidence level when the underlying data distributions are normal.

Caution: This procedure controls the relative width of the interval as a proportion of the true variance ratio. For controlling the absolute width of the interval see the procedure Confidence Intervals for the Ratio of Two Variances using Variances.

Technical Details

Following the results of Desu and Raghavarao (1990), let s_1^2 be the variance estimate based on a sample from a normal distribution with unknown μ_1 and unknown σ_1^2 . Let s_2^2 be the variance estimate based on a sample from a normal distribution with unknown μ_2 and unknown σ_2^2 . Let r be the proportion of σ_1^2/σ_2^2 such that s_1^2/s_2^2 is within $r\sigma_1^2/\sigma_2^2$ of σ_1^2/σ_2^2 with desired confidence $(1 - \alpha)$. That is, the desired condition to be satisfied is

$$\Pr\left(\left|\frac{s_1^2}{s_2^2} - \frac{\sigma_1^2}{\sigma_2^2}\right| \leq r \frac{\sigma_1^2}{\sigma_2^2}\right) \geq 1 - \alpha$$

which can also be written as

657-2 Confidence Intervals for the Ratio of Two Variances using Relative Error

$$\Pr \left(\left| \frac{\frac{s_1^2}{s_2^2} - \frac{\sigma_1^2}{\sigma_2^2}}{\frac{\sigma_1^2}{\sigma_2^2}} \right| \leq r \right) \geq 1 - \alpha$$

which simplifies to

$$\Pr(1 - r \leq F(df_1, df_2) \leq 1 + r) \geq 1 - \alpha$$

since

$$\frac{s_1^2}{s_2^2} \bigg/ \frac{\sigma_1^2}{\sigma_2^2} \sim F(df_1, df_2)$$

If $G_{df_1, df_2}(\cdot)$ is the distribution function of $F(df_1, df_2)$, then probability statement can be rewritten as

$$G_{df_1, df_2}(1 + r) - G_{df_1, df_2}(1 - r) \geq 1 - \alpha$$

This equation can be solved for any of the unknown quantities (df_1, df_2, r, α) in terms of the others.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of random samples of size n_1 and n_2 are drawn from populations 1 and 2, respectively, and a variance ratio is calculated for each pair of samples, the proportion of those estimates that are within $r\sigma_1^2/\sigma_2^2$ of σ_1^2/σ_2^2 is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of random samples of size n_1 and n_2 are drawn from populations 1 and 2, respectively, and a variance ratio is calculated

for each pair of samples, the proportion of those estimates that are within $r\sigma_1^2/\sigma_2^2$ of σ_1^2/σ_2^2 is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as 0.90, 0.95 or 0.90 to 0.99 by 0.01.

Sample Size

N1 (Sample Size Group 1)

Enter a value (or range of values) for the sample size of this group. Note that these values are ignored when you are solving for $N1$. You may enter a range of values such as 10 to 100 by 10.

N2 (Sample Size Group 2)

Enter a value (or range of values) for the sample size of group 2 or enter *Use R* to base $N2$ on the value of $N1$. You may enter a range of values such as 10 to 100 by 10.

- **Use R**

When *Use R* is entered here, $N2$ is calculated using the formula

$$N2 = [R(N1)]$$

where R is the Sample Allocation Ratio and the operator $[Y]$ is the first integer greater than or equal to Y . For example, if you want $N1 = N2$, select *Use R* and set $R = 1$.

R (Sample Allocation Ratio)

Enter a value (or range of values) for R , the allocation ratio between samples. This value is only used when $N2$ is set to *Use R*.

When used, $N2$ is calculated from $N1$ using the formula: $N2 = [R(N1)]$ where $[Y]$ is the next integer greater than or equal to Y . Note that setting $R = 1.0$ forces $N2 = N1$.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

In each case the limits are based on the relative error of the true population variance.

Precision

Relative Error

This is the distance from the true variance ratio as a proportion of the true variance ratio.

You can enter a single value or a list of values. The value(s) must be between 0 and 100000.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to be 95% confident that estimated variance ratio is within 10% of the true variance ratio. In addition to 10% relative error, 5%, 15%, 20% and 25% will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for the Ratio of Two Variances using Relative Error** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **Ratio of Two Variances using Relative Error**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N1 (Sample Size Group 1)	<i>Ignored since this is the Find setting</i>
N2 (Sample Size Group 2)	Use R
R (Sample Allocation Ratio)	1.0
Interval Type	Two-Sided
Relative Error	0.05 to 0.25 by 0.05

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Relative Error Confidence Intervals for the Variance Ratio

Target Confidence Level	Actual Confidence Level	N1	N2	Allocation Ratio	Relative Error
0.950	0.950	6154	6154	1.000	0.050
0.950	0.950	1544	1544	1.000	0.100
0.950	0.950	690	690	1.000	0.150
0.950	0.950	392	392	1.000	0.200
0.950	0.950	253	253	1.000	0.250

References

Desu, M. M. and Raghavarao, D. 1990. Sample Size Methodology. Academic Press. New York.

Report Definitions

Confidence Level is the proportion of variance ratio estimates that will be within the relative error of the true variance ratio.

Target Confidence Level is the value of the confidence level that is entered into the procedure.

Actual Confidence Level is the value of the confidence level that is obtained from the procedure.

N1 and N2 are the sample sizes drawn from the two populations.

Allocation Ratio is the ratio of the sample sizes, $N2/N1$.

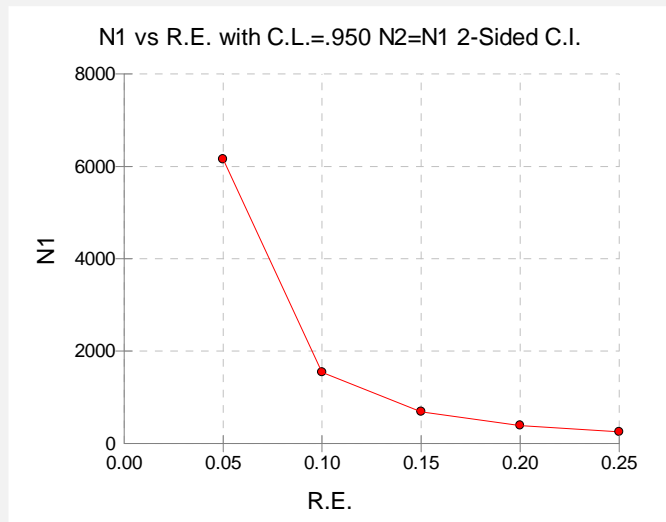
Relative Error is the distance from the true variance ratio as a proportion of the true variance ratio.

Summary Statements

With group sample sizes of 6154 and 6154, the probability is 0.950 (95% confidence) that the estimate of the variance ratio will be within 5% of the true variance ratio.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the relative error.

Example 2 – Validation using Direct Calculation

Suppose a study is planned in which the researcher wishes to be confident that estimated variance ratio is within 20% of the true variance ratio. For sample sizes of 200 per group, the resulting confidence level is $0.9003379878 - 0.0581710981 = 0.8421668897$.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for the Ratio of Two Variances using Relative Error** procedure window by clicking on **Confidence Intervals**, then **Variances**, then **Ratio of Two Variances using Relative Error**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.8421668897
N1 (Sample Size Group 1)	<i>Ignored since this is the Find setting</i>
N2 (Sample Size Group 2)	Use R
R (Sample Allocation Ratio)	1.0
Interval Type	Two-Sided
Relative Error	0.2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Target Confidence Level	Actual Confidence Level	N1	N2	Allocation Ratio	Relative Error
0.842	0.842	200	200	1.000	0.200

PASS calculated the necessary sample size to be 200 per group, which is the same as the direct calculation sample size.

Chapter 670

Normality Tests (Simulation)

Introduction

This procedure allows you to study the power and sample size of eight statistical tests of normality. Since there are no formulas that allow the calculation of power directly, simulation is used. This gives you the ability to compare the adequacy of each test under a wide variety of solutions.

The reason there are so many different normality tests is that there are many different forms of normality. Thode (2002) presents the following recommendations concerning which tests to use for each situation. Note that the details of each test will be presented later.

Normal vs. Long-Tailed Symmetric Alternative Distributions

The Shapiro-Wilk and the kurtosis tests have been found to be best for normality testing against long-tailed symmetric alternatives.

Normal vs. Short-Tailed Symmetric Alternative Distributions

The Shapiro-Wilk and the range tests have been found to be best for normality testing against short-tailed symmetric alternatives.

Normal vs. Asymmetric Alternative Distributions

The Shapiro-Wilk and the skewness tests have been found to be best for normality testing against asymmetric alternatives.

Technical Details

Computer simulation allows one to estimate the power and significance level that is actually achieved by a test procedure in situations that are not mathematically tractable. Computer simulation was once limited to mainframe computers. Currently, due to increased computer speeds, simulation studies can be completed on desktop and laptop computers in a reasonable period of time.

The steps to a simulation study are as follows.

1. Specify which the normality test is to be used. This includes specifying the significance level.

670-2 Normality Tests (Simulation)

2. Generate a random sample, X_1, X_2, \dots, X_n , from the distribution specified by the alternative hypothesis. Calculate the test statistic from the simulated data and determine if the null hypothesis is accepted or rejected. Each of these samples is used to calculate the power of the test. Note that if the alternative distribution is set to normal, the power is the significance level of the test.
3. Repeat step 2 several hundred times, tabulating the number of times the simulated data lead to a rejection of the null hypothesis. The power is the proportion of simulation samples in step 2 that lead to rejection.

Data Distributions

A wide variety of distributions may be studied. These distributions can vary in skewness, elongation, or other features such as bimodality. A detailed discussion of the distributions that may be used in the simulation is provided in the chapter 'Data Simulator'.

Test Statistics

This section describes the test statistics that are available of study in this procedure.

Anderson-Darling Test

This test, developed by Anderson and Darling (1954), is a popular normality test based on EDF statistics. In some situations, it has been found to be as powerful as the Shapiro-Wilk test. This test is available when n is greater than or equal to 8.

Kolmogorov-Smirnov

This test for normality is based on the maximum difference between the observed distribution and expected cumulative-normal distribution. Since it uses the sample mean and standard deviation to calculate the expected normal distribution, the Lilliefors' adjustment is used. The Lilliefors' adjusted critical values used are those given by Dallal (1986).

This test is available when n is greater than or equal to 3.

This test has been shown to be less powerful than the other tests in most situations. It is included because of its historical popularity.

Kurtosis

D'Agostino (1990) describes a normality test based on the kurtosis coefficient, b_2 . Recall that for the normal distribution, the theoretical value of b_2 is 3. Hence, a test can be developed to determine if the value of b_2 is significantly different from 3. If it is, the data are obviously nonnormal. The statistic, z_k , is, under the null hypothesis of normality, approximately normally distributed for sample sizes $n > 20$. This test is available when n is greater than or equal to 8.

The calculation of this test proceeds as follows:

$$z_k = \frac{\left(1 - \frac{2}{9A}\right) - \left(\frac{1 - \frac{2}{A}}{1 + G\sqrt{\frac{2}{A-4}}}\right)^{1/3}}{\sqrt{\frac{2}{9A}}}$$

where

$$b_2 = \frac{m_4}{m_2^2}$$

$$G = \frac{b_2 - \left(\frac{3n-3}{n+1}\right)}{\sqrt{\frac{24n(n-2)(n-3)}{(n+1)^2(n+3)(n+5)}}}$$

$$E = \frac{6(n^2 - 5n + 2)}{(n+7)(n+9)} \sqrt{\frac{6(n+3)(n+5)}{n(n-2)(n-3)}}$$

$$A = 6 + \frac{8}{E} \left(\frac{2}{E} + \sqrt{1 + \frac{4}{E^2}} \right)$$

Martinez-Iglewicz

This test for normality, developed by Martinez and Iglewicz (1981), is based on the median and a robust estimator of dispersion. They have shown that this test is very powerful for heavy-tailed symmetric distributions as well as a variety of other situations. A value of the test statistic that is close to one indicates that the distribution is normal. This test is recommended for exploratory data analysis by Hoaglin (1983). The formula for this test is:

$$I = \frac{\sum_{i=1}^n (x_i - \bar{x})^2}{(n-1)s_{bi}^2}$$

where s_{bi}^2 is a biweight estimator of scale.

This test is available when n is greater than or equal to 3.

Omnibus

D'Agostino (1990) describes a normality test that combines the tests for skewness and kurtosis. The statistic, K^2 , is approximately distributed as a chi-square with two degrees of freedom.

After calculating z_s and z_k , calculate K^2 as follows:

$$K^2 = z_s^2 + z_k^2$$

This test is available when n is greater than or equal to 8.

Range Test

The range test, u , was created to test for normality when the alternative distribution is actually the uniform distribution. It is calculated by dividing the range by the standard deviation. Tables of critical values of the range test are given in Pearson and Hartley (1976) for n equal 3 to 1000.

Shapiro-Wilk W Test

This test for normality, developed by Shapiro and Wilk (1965), has been found to be the most powerful test in most situations. It is the ratio of two estimates of the variance of a normal distribution based on a random sample of n observations. The numerator is proportional to the square of the best linear estimator of the standard deviation. The denominator is the sum of squares of the observations about the sample mean. W may be written as the square of the Pearson correlation coefficient between the ordered observations and a set of weights which are used to calculate the numerator. Since these weights are asymptotically proportional to the corresponding expected normal order statistics, W is roughly a measure of the straightness of the normal quantile-quantile plot. Hence, the closer W is to one, the more normal the sample.

The probability values for W are valid for samples in the range of 3 to 5000.

Skewness

D'Agostino (1990) describes a normality test based on the skewness coefficient, $\sqrt{b_1}$. Recall that because the normal distribution is symmetrical, $\sqrt{b_1}$ is equal to zero for normal data. Hence, a test can be developed to determine if the value of $\sqrt{b_1}$ is significantly different from zero. If it is, the data are obviously nonnormal. The statistic, z_s , is, under the null hypothesis of normality, approximately normally distributed. The computation of this statistic, which is restricted to sample sizes $n > 8$, is

$$z_s = d \ln \left(\frac{T}{a} + \sqrt{\left(\frac{T}{a} \right)^2 + 1} \right)$$

where

$$b_1 = \frac{m_3^2}{m_2^3}$$

$$T = \sqrt{b_1 \left(\frac{(n+1)(n+3)}{6(n-2)} \right)}$$

$$C = \frac{3(n^2 + 27n - 70)(n+1)(n+3)}{(n-2)(n+5)(n+7)(n+9)}$$

$$W^2 = -1 + \sqrt{2(C-1)}$$

$$a = \sqrt{\frac{2}{W^2 - 1}}$$

$$d = \frac{1}{\sqrt{\ln(W)}}$$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be calculated using the values of the other parameters. Under most conditions, you would select either *Power* or *N*.

Select *Power* when you want to estimate the power for a specific scenario.

Select *N* when you want to determine the sample size needed to achieve a given power and alpha error level. This option can be very computationally intensive, and may take considerable time to complete.

Search/Report Test

Specify the specific normality test to be used for searching (if Find = 'N') and reporting. Note that each test was developed for a specific situation and that no one test is best in all situations. Apparently, if you have no other information, you would choose the Shapiro-Wilk test.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when a true null hypothesis is rejected.

Values must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

This option specifies one or more values of the sample size, the number of individuals in the study. This value must be an integer greater than one. Note that you may enter a list of values using the syntax *50 100 150 200 250* or *50 to 250 by 50*.

Simulations

Simulations

This option specifies the number of iterations, M , used in the simulation. Larger numbers of iterations result in longer running time and more accurate results.

The precision of the simulated power estimates can be determined by recognizing that they follow the binomial distribution. Thus, confidence intervals may be constructed for power estimates. The following table gives an estimate of the precision that is achieved for various simulation sizes when the power is either 0.50 or 0.95. The table values are interpreted as follows: a 95% confidence interval of the true power is given by the power reported by the simulation plus and minus the 'Precision' amount given in the table.

Simulation Size M	Precision when Power = 0.50	Precision when Power = 0.95
100	0.100	0.044
500	0.045	0.019
1000	0.032	0.014
2000	0.022	0.010
5000	0.014	0.006
10000	0.010	0.004
50000	0.004	0.002
100000	0.003	0.001

Notice that a simulation size of 1000 gives a precision of plus or minus 0.014 when the true power is 0.95. Also note that as the simulation size is increased beyond 5000, there is only a small amount of additional precision achieved.

Effect Size

Data Distribution

This option specifies distribution that is being compared to normality. That is, this is the actual (true) distribution of the data from which the power is computed. Usually, the mean is specified by entering 'M1' for the mean parameter in the distribution expression and then entering values for the M1 parameter below. All of the distributions are parameterized so that the mean is entered first.

The parameters of each distribution are specified using numbers or letters. If letters are used, their values are specified in the boxes below. The values 'M0' and 'M1' are usually used to specify the value of the mean.

Following is a list of the distributions that are available and the syntax used to specify them:

Beta=A(M1,A,B,Minimum)
 Binomial=B(M1,N)
 Cauchy=C(M1,Scale)
 Constant=K(Value)
 Exponential=E(M1)
 F=F(M1,DF1)
 Gamma=G(M1,A)
 Multinomial=M(P1,P2,P3,...,Pk)
 Normal=N(M1,SD)
 Poisson=P(M1)
 Student's T=T(M1,D)
 Tukey's Lambda=L(M1,S,Skewness,Elongation)
 Uniform=U(M1,Minimum)
 Weibull=W(M1,B)

Details of writing mixture distributions, combined distributions, and compound distributions are found in the chapter on Data Simulation and will not be repeated here.

Effect Size – Distribution Parameters

M0 and M1

These values are substituted for the M0 and M1 in the distribution specifications given above. M0 is intended to be the value of the mean.

You can enter a list of values using syntax such as *0 1 2 3* or *0 to 3 by 1*.

Parameter Values (S, A, B, C)

Enter the numeric value(s) of parameter listed above. These values are substituted for the corresponding letter in the Data Distribution specifications.

You can enter a list of values using syntax such as *0 1 2 3* or *0 to 3 by 1*.

You can also change the letter that is used as the name of this parameter.

Iterations Tab

The Iterations tab contains limits on the number of iterations and various options about individual tests.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations before the search for the sample size, N, is aborted. When the maximum number of iterations is reached without convergence, the sample size is not reported. We recommend a value of at least 500.

Example 1 – Power at Various Sample Sizes

A researcher is planning an experiment to test whether a certain data distribution is reasonably close to normality. He will begin his research by generating data from the exponential distribution. The researcher is particularly interested in the Shapiro-Wilk normality test, but he wants to see the power of various other choices when the data actually come from an exponential distribution. Alpha is set to 0.05. He wants to compute the power at various sample sizes from 5 to 40. Since this is an exploratory analysis, he sets the number of simulation iterations to 1000.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Normality Tests (Simulation)** procedure window by clicking on **Means**, then **Normality Tests (Simulation)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power
Search/Report Test	Shapiro-Wilk
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	5 10 15 20 30 40
Simulations.....	1000
Data Distribution.....	E(M1)
M1	1
Reports Tab	
Show Numeric Report	Checked
Show Comparative Report.....	Checked
Show Confidence Interval for Power.....	Checked
Show Definitions	Checked
Show Plots	Checked
Number of Summary Statements.....	1

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for the Shapiro-Wilk Normality Test Actual Distribution: Expo(M1)

	Lower 95% C.L. of Power	Upper 95% C.L. of Power	N	Alpha	Beta	M1	No.
0.175	0.152	0.200	5	0.050	0.825	1.0	1
0.442	0.411	0.473	10	0.050	0.558	1.0	2
0.690	0.660	0.719	15	0.050	0.310	1.0	3
0.840	0.816	0.862	20	0.050	0.160	1.0	4
0.968	0.955	0.978	30	0.050	0.032	1.0	5
0.996	0.990	0.999	40	0.050	0.004	1.0	6

Notes

Simulations: 1000. Run Time: 12.28 seconds.

References

Henry C. Thode, Jr. 2002. Testing for Normality. Marcel Dekker. New York.
Devroye, Luc. 1986. Non-Uniform Random Variate Generation. Springer-Verlag. New York.

Report Definitions

Power is the probability of rejecting a false null hypothesis.
Lower/Upper 95% C.L. are the boundaries of an exact binomial confidence interval for power.
N is the size of the sample drawn from the population.
Alpha is the probability of rejecting when the distribution is normal.
Beta is the probability of concluding normality when the distribution is actually the one stated.
No. is the sequence number of this combination of parameter values.

Summary Statements

A sample size of 5 achieves 18% power to detect non-normality using the Shapiro-Wilk test when the significance level is 0.050 and the actual distribution is Expo(M1). These results are based on 1000 Monte Carlo samples from this distribution.

This report shows the estimated power with its confidence interval (using the binomial distribution) for each combination of parameter values. Note that because these are results of a simulation study, the computed power and alpha will vary from run to run. Thus, another report obtained using the same input parameters will be slightly different from the one above.

If the researcher wants 80% power, he will need an N of 20. To achieve 90% power, he will need an N of 30.

Comparative Results

Power of Various Normality Tests Actual Distribution: Expo(M1)

Sequence No.	N	Anderson Darling	Kolmo'v Smirnov	Kur- tosis	Martinez Iglewicz	Omni- bus	Range /SD	Shapiro Wilk	Skew- ness	Any Test
1	5		0.124		0.225		0.076	0.175		0.290
		Lower 95% C.L.	0.104		0.199		0.060	0.152	0.000	0.262
		Upper 95% C.L.	0.146		0.252		0.094	0.200	0.000	0.319

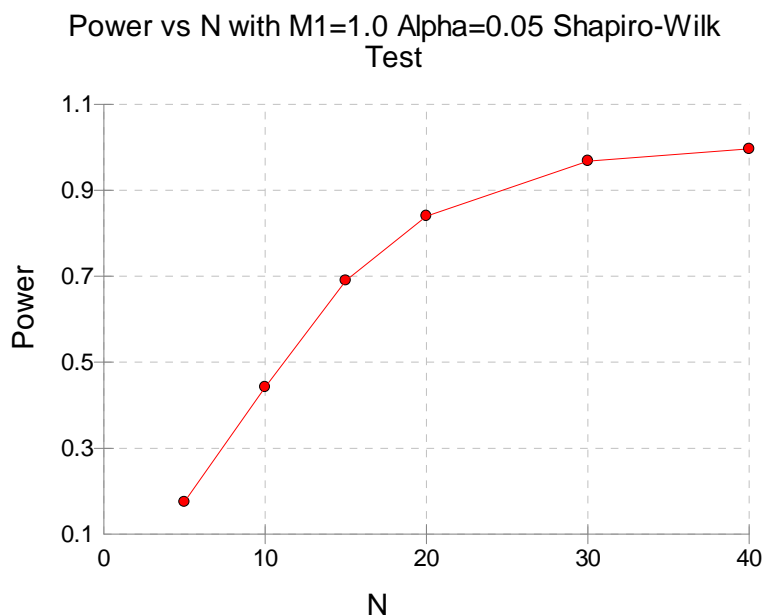
670-10 Normality Tests (Simulation)

2	10	0.401	0.297	0.215	0.311	0.315	0.062	0.442	0.366	0.549
	Lower 95% C.L.	0.370	0.269	0.190	0.282	0.286	0.048	0.411	0.336	0.518
	Upper 95% C.L.	0.432	0.326	0.242	0.341	0.345	0.079	0.473	0.397	0.580
3	15	0.627	0.448	0.290	0.456	0.474	0.068	0.690	0.562	0.750
	Lower 95% C.L.	0.596	0.417	0.262	0.425	0.443	0.053	0.660	0.531	0.722
	Upper 95% C.L.	0.657	0.479	0.319	0.487	0.505	0.085	0.719	0.593	0.777
4	20	0.784	0.597	0.349	0.598	0.607	0.074	0.840	0.713	0.871
	Lower 95% C.L.	0.757	0.566	0.319	0.567	0.576	0.059	0.816	0.684	0.849
	Upper 95% C.L.	0.809	0.628	0.379	0.629	0.637	0.092	0.862	0.741	0.891
5	30	0.937	0.775	0.485	0.735	0.777	0.131	0.968	0.895	0.975
	Lower 95% C.L.	0.920	0.748	0.454	0.706	0.750	0.111	0.955	0.874	0.963
	Upper 95% C.L.	0.951	0.801	0.516	0.762	0.802	0.154	0.978	0.913	0.984
6	40	0.987	0.886	0.588	0.839	0.894	0.142	0.996	0.958	0.998
	Lower 95% C.L.	0.978	0.865	0.557	0.815	0.873	0.121	0.990	0.944	0.993
	Upper 95% C.L.	0.993	0.905	0.619	0.861	0.912	0.165	0.999	0.970	1.000

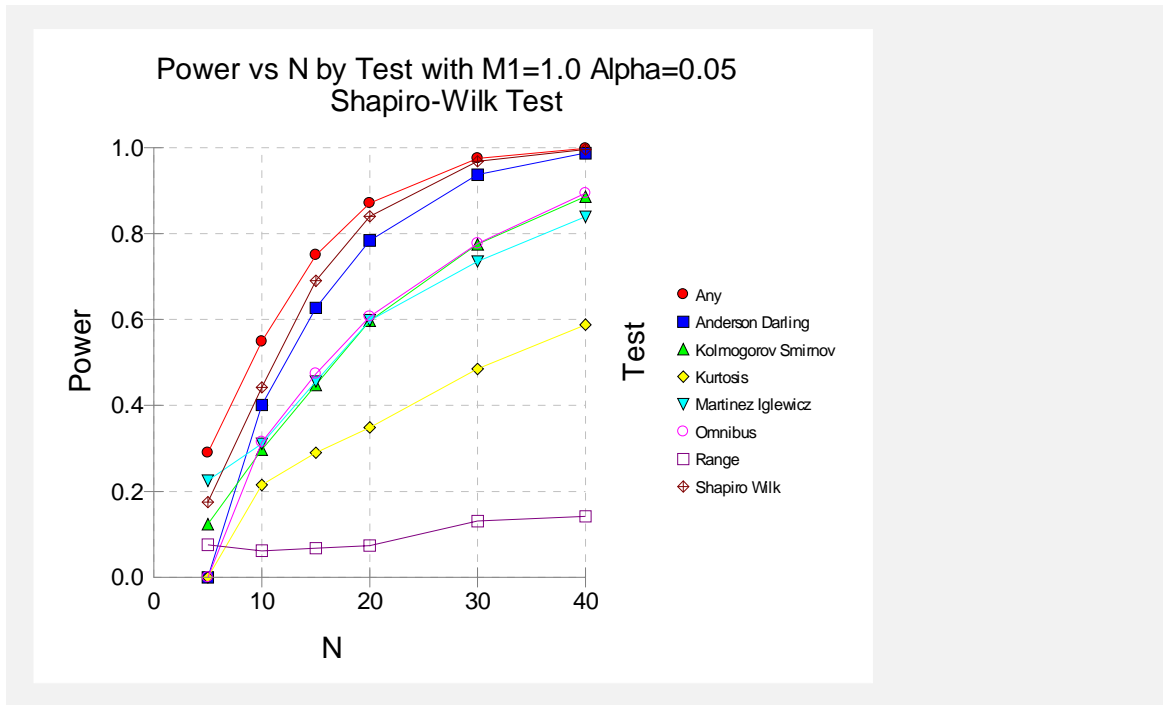
This report shows the estimated power with its confidence interval (using the binomial distribution) for each combination each of the normality tests. The last test, called *Any Test*, is calculated by combining the results of the eight normality tests into a single test.

Note the variation in the power values. In this case, the Martinez-Iglewicz is the champion for $N = 5$, but the Shapiro-Wilk test has greater power for $N = 10$.

Plots Section



This plot shows the relationship between sample size and power.



This plot shows performance of each of the normality tests. The champion is the Shapiro-Wilk test. The range test, which was designed for uniform alternatives, does very poorly with exponential alternatives.

Example 2 – Validation using the Normal Distribution

We will validate this procedure by setting the normal distribution as the alternative distribution. In this case, the power should be equal to the alpha value. Alpha is set to 0.05. The sample sizes will be set to 10, 30, and 50. The number of simulation iterations to 10,000.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Normality Tests (Simulation)** procedure window by clicking on **Means**, then **Normality Tests (Simulation)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power
Search/Report Test.....	Shapiro-Wilk
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	10 30 50
Simulations.....	10000
Data Distribution.....	N(M1 S)

670-12 Normality Tests (Simulation)

Data Tab (continued)

M1 1
S 1

Reports Tab

Show Numeric Report **Checked**
Show Comparative Report **Checked**
Show Confidence Interval for Power **Checked**
Show Definitions **Checked**
Show Plots **Checked**
Number of Summary Statements 1

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for the Shapiro-Wilk Normality Test Actual Distribution: N(M1 S)

	Lower 95% C.L. of Power	Upper 95% C.L. of Power	N	Alpha	Beta	M1	S	No.
Power								
0.052	0.047	0.056	10	0.050	0.948	1.0	1.0	1
0.047	0.043	0.051	30	0.050	0.953	1.0	1.0	2
0.053	0.048	0.057	50	0.050	0.948	1.0	1.0	3

Power of Various Normality Tests Actual Distribution: N(M1 S)

Sequence No.	N	Anderson Darling	Kolmo'v Smirnov	Kur- tosis	Martinez Iglewicz	Omni- bus	Range /SD	Shapiro Wilk	Skew- ness	Any Test
1	10	0.053	0.027	0.043	0.051	0.058	0.049	0.052	0.052	0.142
		0.048	0.024	0.039	0.046	0.053	0.045	0.047	0.048	0.135
		0.057	0.030	0.047	0.055	0.062	0.053	0.056	0.057	0.148
2	30	0.048	0.029	0.053	0.049	0.053	0.053	0.047	0.047	0.151
		0.044	0.026	0.049	0.045	0.048	0.049	0.043	0.043	0.145
		0.053	0.033	0.058	0.054	0.057	0.058	0.051	0.052	0.159
3	50	0.052	0.032	0.055	0.049	0.057	0.054	0.053	0.048	0.159
		0.048	0.028	0.050	0.045	0.052	0.050	0.048	0.044	0.152
		0.056	0.035	0.059	0.053	0.061	0.059	0.057	0.053	0.167

Simulations: 10000. Run Time: 77.42 seconds.

This report confirms that the simulation results are quite accurate. Most of the power values are very close to 0.05. The Kolmogorov-Smirnov test appears to be low and the skewness, kurtosis, and omnibus appear to be high. Perhaps this is the reason that the Shapiro-Wilk and the Anderson-Darling tests are usually recommended for general normality testing.

Chapter 700

Logrank Tests (Freedman)

Introduction

This module allows the sample size and power of the logrank test to be analyzed under the assumption of proportional hazards. Time periods are not stated. Rather, it is assumed that enough time elapses to allow for a reasonable proportion of responses to occur. If you want to study the impact of accrual and follow-up time, you should use the one of the other logrank modules also contained in *PASS*. The formulas used in this module come from Machin *et al.* (1997).

A clinical trial is often employed to test the equality of survival distributions for two treatment groups. For example, a researcher might wish to determine if Beta-Blocker A enhances the survival of newly diagnosed myocardial infarction patients over that of the standard Beta-Blocker B. The question being considered is whether the pattern of survival is different.

The two-sample t-test is not appropriate for two reasons. First, the data consist of the length of survival (time to failure), which is often highly skewed, so the usual normality assumption cannot be validated. Second, since the purpose of the treatment is to increase survival time, it is likely (and desirable) that some of the individuals in the study will survive longer than the planned duration of the study. The survival times of these individuals are then said to be *censored*. These times provide valuable information, but they are not the actual survival times. Hence, special methods have to be employed which use both regular and censored survival times.

The logrank test is one of the most popular tests for comparing two survival distributions. It is easy to apply and is usually more powerful than an analysis based simply on proportions. It compares survival across the whole spectrum of time, not at just one or two points.

The power calculations used here assume an underlying exponential distribution. However, we are rarely in a position to assume exponential survival times in an actual clinical trial. How do we justify the exponential survival time assumption? First, the logrank test and the test derived using the exponential distribution have nearly the same power when the data are in fact exponentially distributed. Second, under the proportional hazards model (which is assumed by the logrank test), the survival distribution can be transformed to be exponential and the logrank test remains the same under monotonic transformations.

Technical Details

In order to define the input parameters, we will present below some rather complicated looking formulas. You need not understand the formulas. However, you should understand the individual parameters used in these formulas.

We assume that a study is to be made comparing the survival of an existing (control) group with an experimental group. The control group consists of patients that will receive the existing treatment. In cases where no existing treatment exists, the control group consists of patients that will receive a placebo. This group is arbitrarily called group one. The experimental group will receive the new treatment. It is called group two.

We assume that the critical event of interest is death and that two treatments have survival distributions with instantaneous death (hazard) rates, λ_1 and λ_2 . These hazard rates are an subject's probability of death in a short period of time. We want to test hypotheses about these hazard rates.

Hazard Ratio

There are several ways to express the difference between two hazard rates. One way is to calculate the difference, $\lambda_1 - \lambda_2$. Another way is to form the hazard ratio (HR):

$$HR = \frac{\lambda_2}{\lambda_1}.$$

Note that since HR is formed by dividing the hazard rate of the experimental group by that of the control group, a treatment that does better than the control will have a hazard ratio that is less than one.

The hazard ratio may be formulated in other ways. If the proportions surviving during the study are called $S1$ and $S2$ for the control and experimental groups, the hazard ratio is given by

$$HR = \frac{\text{Log}(S2)}{\text{Log}(S1)}.$$

Furthermore, if the median survival times of the two groups are $M1$ and $M2$, the hazard ratio is given by

$$HR = \frac{M1}{M2}.$$

Each of these expressions for the difference between hazards rates is useful in some situations. There is no one best way, so you will have to be a little flexible.

Logrank Test

We assume that the logrank test will be used to analyze the data once they are collected. Basing the power calculations on the logrank test, we arrive at the following formula that gives the power based on several other parameters:

$$z_{1-\beta} = \frac{|HR - 1| \sqrt{N(1-w)\varphi[(1-S1) + \varphi(1-S2)] / (1+\varphi)}}{(1+\varphi HR)} - z_{1-\alpha/k}$$

where k is 1 for a one-sided hypothesis test or 2 for a two-sided test, α and β are the error rates defined as usual, the z 's are the usual points from the standard normal distribution, w is the proportion that are lost to follow up, and φ represents the sample size ratio between the two groups. That is, $p1$ is the proportion of the total sample size that is in the control group, φ is given by

$$\varphi = \frac{1 - p1}{p1}$$

Note that the hypothesis being tested is that the hazard rates are equal:

$$H_0: \lambda_1 = \lambda_2$$

This formulation assumes that the hazard rates are proportional. It does not assume that the failure times are exponentially distributed.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *Proportion Surviving 1*, *Proportion Surviving 2*, *N*, *Alpha*, or *Power and Beta*. Under most situations, you will select either *Power and Beta* or *N*.

Select *N* when you want to calculate the sample size needed to achieve a given power and alpha level.

Select *Power and Beta* when you want to calculate the power of an experiment that has already been run.

If you select *Survive 1*, the search is made for values that are less than *S2*. Likewise, if you select *Survive 2*, the search is made for values that are greater than *S1*.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal survival curves when in fact the curves are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis of equal survival curves when in fact the curves are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size – Total Sample Size

N (Total Sample Size)

The combined sample size of both groups.

Sample Size – Sample Proportions

Proportion in Group 1

The proportion of the sample size (N) that is assigned to the first group. The rest of the patients are assigned to the second group. When exact values are not available, the number in group one is found by rounding up to the next integer. The number in group two is found by subtraction.

Proportion Lost During Follow Up

This is the proportion of patients in the trial that are lost to follow-up. Since they are lost, you can never obtain their outcome. The sample size is inflated to account for this loss using the formula:

$$N(\text{adjusted}) = N/(1-p)$$

where p is the proportion lost.

Effect Size

S1 (Proportion Surviving in Group 1)

$S1$ is the proportion of patients belonging to group 1 (controls) that are expected to survive during the study. Since $S1$ is a proportion, it must be between zero and one. A value for $S1$ must be determined either from a pilot study or from previous studies.

S2 (Proportion Surviving in Group 2)

$S2$ is the proportion of patients belonging to group 2 (experimental) that are expected to survive during the study. Since $S2$ is a proportion, it must be between zero and one.

This value is not necessarily the expected survival proportion under the treatment. Rather, you may set it to that value that, if achieved, would be of special interest. Values below this amount would not be of interest.

For example, if the standard 1-year survival proportion is 0.2 and the new treatment raises this proportion to 0.3 (a 50% increase in the proportion surviving), others may be interested in using it.

Sometimes, researchers wish to state the alternative hypothesis in terms of the hazard ratio, HR , rather than the value of $S2$. Using the fact that

$$HR = \text{Log}(S2)/\text{Log}(S1)$$

we can solve for $S2$ to obtain

$$S2 = \text{Exp}(\text{Log}(S1)HR).$$

Using this equation, you can determine an appropriate value for $S2$ from a value for HR .

For example, suppose $S1$ is 0.4 and you have determined that if the hazard rate is reduced by half, the improvement is sufficient to justify a change in treatment. The appropriate value for $S2$ is $\text{Exp}(\text{Log}(0.4)(0.5)) = 0.632$.

When you do not have an anticipated value for HR , you can find an estimate base on the Median Survival Times ($M1$ and $M2$) since $HR = M1/M2$.

Test
One-Sided Test

Specify whether the test is one-sided (checked) or two-sided (unchecked). When a two-sided test is selected, the value of alpha is divided by two.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations
Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Finding the Power

An experiment has been conducted to test the effectiveness of a new treatment on a total of 100 patients. The current treatment for this disease achieves 50% survival after two years. The proportion in the treatment group that survived two years was 0.70. Testing was done at the 0.05 significance level. Even though there was an increase of 0.20 for 0.50 to 0.70, the logrank test did not reject the hypothesis of equal hazard rates for the two treatments. Study the power of this test.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Freedman)** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Freedman)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05 0.10
N (Total Sample Size)	50 to 300 by 50
Proportion in Group 1	0.5
Proportion Lost During Follow Up	0
S1 (Proportion Surviving in Group 1)	0.5
S2 (Proportion Surviving in Group 2)	0.7
One-Sided Test	Not Checked

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results								
Power	N	N1	N2	S1	S2	Hazard Ratio	Two-Sided Alpha	Beta
0.2992	50	25	25	0.5000	0.7000	0.5000	0.0500	0.7008
0.4162	50	25	25	0.5000	0.7000	0.5000	0.1000	0.5838
0.5267	100	50	50	0.5000	0.7000	0.5000	0.0500	0.4733
0.6488	100	50	50	0.5000	0.7000	0.5000	0.1000	0.3512
0.6994	150	75	75	0.5000	0.7000	0.5000	0.0500	0.3006
0.7989	150	75	75	0.5000	0.7000	0.5000	0.1000	0.2011
0.8177	200	100	100	0.5000	0.7000	0.5000	0.0500	0.1823
0.8891	200	100	100	0.5000	0.7000	0.5000	0.1000	0.1109
0.8934	250	125	125	0.5000	0.7000	0.5000	0.0500	0.1066
0.9406	250	125	125	0.5000	0.7000	0.5000	0.1000	0.0594
0.9395	300	150	150	0.5000	0.7000	0.5000	0.0500	0.0605
0.9690	300	150	150	0.5000	0.7000	0.5000	0.1000	0.0310

Report Definitions

Power is the probability of rejecting a false null hypothesis.

N is the combined sample size.

N1 sample size in group 1.

N2 sample size in group 2.

S1 is the proportion surviving in group 1

S2 is the proportion surviving in group 2

The Hazard Ratio is the ratio of hazard2 and hazard1. It is $\text{Log}(S2)/\text{Log}(S1)$.

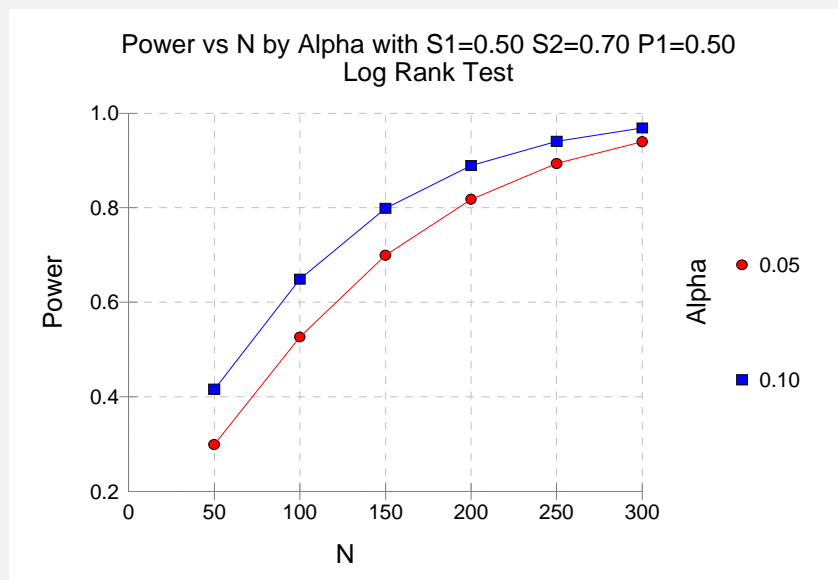
Alpha is the probability of rejecting a true null hypothesis.

Beta is the probability of accepting a false null hypothesis.

Summary Statements

A two-sided log rank test with an overall sample size of 50 subjects (of which 25 are in group 1 and 25 are in group 2) achieves 30% power at a 0.0500 significance level to detect a difference of 0.2000 between 0.5000 and 0.7000--the proportions surviving in groups 1 and 2, respectively. This corresponds to a hazard ratio of 0.5000. The proportion of patients lost during follow up was 0.0000. These results are based on the assumption that the hazard rates are proportional.

This report shows the values of each of the parameters, one scenario per row. The values from this table are in the chart below.

Plots Section

This plot shows the relationship between alpha and power in this example. We notice that for $\alpha = 0.05$, a power of 0.80 is reached when the sample size is about 200. A power of 90% is reached when the sample size is 250. Hence, we realize that this study should have had at least twice the number of patients that it had.

Example 2 – Finding the Sample Size

Continuing with our example, the researcher decides that he wants to do it right this time. He believes that if the survival proportion of the treatment group is 0.6 or better, people will begin to use his treatment. He wants to know how many subjects he needs.

Setup

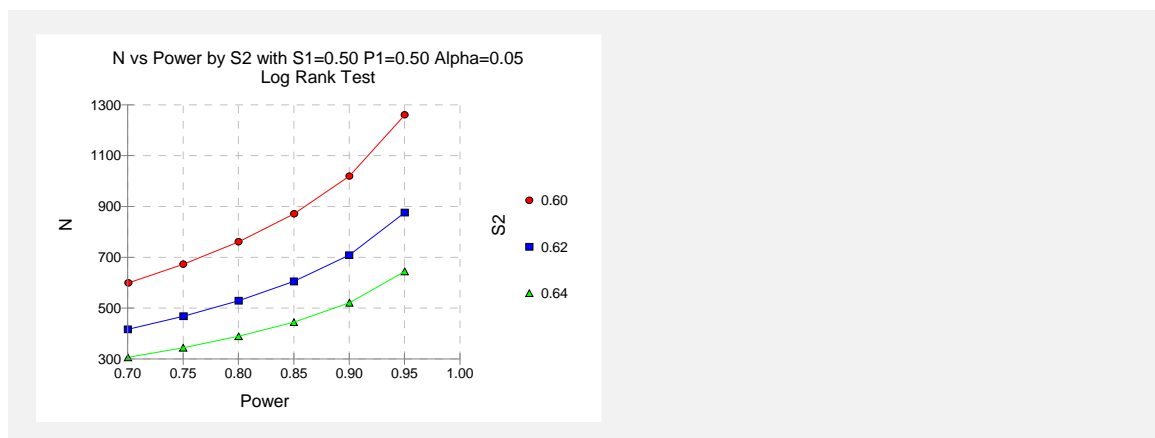
This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Freedman)** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Freedman)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.70 to 0.95 by 0.05
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Group 1	0.5
Proportion Lost During Follow Up	0
S1 (Proportion Surviving in Group 1)	0.5
S2 (Proportion Surviving in Group 2)	0.60 0.62 0.64
One-Sided Test	Not Checked

Output

Click the Run button to perform the calculations and generate the following output.

Plots Section



We will consider the chart since that allows us to understand the patterns more quickly. We note that changing S_2 from 0.60 (the top line) to 0.62 (the middle line) decreases the sample size requirements by almost half. Our researcher decides to take a sample of 500 patients. This will achieve almost 80% power in detecting a shift from 0.50 to 0.62 in survivability.

Example 3 – Validation using Machin

Machin *et al.* (1997) page 180 gives an example in which $S1$ is 0.25, $S2$ is 0.50, the one-sided significance level is 0.05, and the power is 90%. The total sample size is 124. We will now run this example through *PASS*.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the *PASS* Home window, load the **Logrank Tests (Freedman)** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Freedman)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.9
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Group 1	0.5
Proportion Lost During Follow Up	0
S1 (Proportion Surviving in Group 1)	0.25
S2 (Proportion Surviving in Group 2)	0.50
One-Sided Test	Checked

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results									
Power	N	N1	N2	S1	S2	Hazard Ratio	One-Sided Alpha	Beta	
0.9014	124	62	62	0.2500	0.5000	0.5000	0.0500	0.0986	

PASS also calculated the total sample size to be 124.

700-10 Logrank Tests (Freedman)

Chapter 705

Logrank Tests (Lachin and Foulkes)

Introduction

This module computes the sample size and power of the logrank test for equality (or non-inferiority) of survival distributions under the assumption of proportional hazards. Accrual time and follow-up time are included among the parameters to be set.

A clinical trial is often employed to test the equality of survival distributions for two treatment groups. For example, a researcher might wish to determine if Beta-Blocker A enhances the survival of newly diagnosed myocardial infarction patients over that of the standard Beta-Blocker B. The question being considered is whether the pattern of survival is different.

The two-sample t-test is not appropriate for two reasons. First, the data consist of the length of survival (time to failure), which is often highly skewed, so the usual normality assumption cannot be validated. Second, since the purpose of the treatment is to increase survival time, it is likely (and desirable) that some of the individuals in the study will survive longer than the planned duration of the study. The survival times of these individuals are then said to be *censored*. These times provide valuable information, but they are not the actual survival times. Hence, special methods have to be employed which use both regular and censored survival times.

The logrank test is one of the most popular tests for comparing two survival distributions. It is easy to apply and is usually more powerful than an analysis based simply on proportions. It compares survival across the whole spectrum of time, not just at one or two points. This module allows the sample size and power of the logrank test to be analyzed under very general conditions.

Power and sample size calculations for the logrank test have been studied by several authors.

PASS uses the methods of Lachin and Foulkes (1986) because of their generality. Although small differences among methods can be found depending upon which assumptions are adopted, there is little practical difference between the techniques.

The power calculations used here assume an underlying exponential distribution. However, this assumption is rarely accurate in an actual clinical trial. How is this assumption justified? First, the logrank test and the test derived using the exponential distribution have nearly the same power when the data are in fact exponentially distributed. Second, under the proportional hazards model

(which is assumed by the logrank test), the survival distribution can be transformed to exponential and the logrank test remains the same under monotonic transformations.

Technical Details

Some rather complicated formulas are used to define the input parameters. You need not understand the formulas. However, you should understand the individual parameters used in these formulas.

Basic Model

Suppose a clinical trial consists of two independent treatment groups labeled “1” and “2” (you could designate group one as the control group and group two as the treatment group). If the total sample size is N , the sizes of the two groups are n_1 and n_2 . Usually, you would plan to have $n_1 = n_2$. Define the proportion of the total sample in each group as

$$Q_i = \frac{n_i}{N}, \quad i = 1, 2$$

Individuals are recruited during an accrual period of R years (or months or days). They are followed for an additional period of time until a total of T years is reached. Hence, the follow-up period is $T-R$ years. At the end of the study, the logrank test is conducted at significance level α with power $1 - \beta$.

If we assume an exponential model with hazard rates λ_1 and λ_2 for the two groups, Lachin and Foulkes (1986, Eq. 2.1) establish the following equation relating N and power:

$$\sqrt{N}/|\lambda_1 - \lambda_2| = Z_\alpha \sqrt{\phi(\bar{\lambda}) \left(\frac{1}{Q_1} + \frac{1}{Q_2} \right)} + Z_\beta \sqrt{\frac{\phi(\lambda_1)}{Q_1} + \frac{\phi(\lambda_2)}{Q_2}}$$

where

$$\bar{\lambda} = Q_1 \lambda_1 + Q_2 \lambda_2$$

$$\phi(\lambda) = N \sigma^2(\hat{\lambda})$$

$$Z_\theta = \Phi(1 - \theta)$$

$\Phi(z)$ is the area to the left of z under the standard normal density, $\hat{\lambda}$ is the maximum likelihood estimate of λ , and that $\sigma^2(\hat{\lambda})$ represents the variance of $\hat{\lambda}$.

Exponential Distribution

The hazard rate from the exponential distribution, λ , is usually estimated using maximum likelihood techniques. In the planning stages, you have to obtain an estimate of this parameter. To see how to accomplish this, let's briefly review the exponential distribution. The density function of the exponential is defined as

$$f(t) = \lambda e^{-\lambda t}, \quad t \geq 0, \lambda > 0.$$

The cumulative survival distribution function is

$$S(t) = e^{-\lambda t}, \quad t \geq 0.$$

Solving this for λ yields

$$\lambda = -\frac{\log(S(t))}{t}$$

Note that $S(t)$ gives the probability of surviving t years. To obtain a planning estimate of λ , you need only know the proportion surviving during a particular time period. You can then use the above equation to calculate λ .

Patient Entry

Patients are enrolled during the accrual period. **PASS** lets you specify the pattern in which subjects are enrolled. Suppose patient entry times are distributed as $g(t)$ where t_i is the entry time of the i^{th} individual and $0 \leq t_i \leq R$. Let $g(t)$ follow the truncated exponential distribution with parameter A , which has the density

$$g(t) = \begin{cases} \frac{A e^{-At}}{1 - e^{-AR}}, & 0 \leq t \leq R, \quad A \neq 0 \\ 1/R, & 0 \leq t \leq R, \quad A = 0 \end{cases}$$

Note that R is accrual time. The corresponding cumulative distribution function is

$$G(t) = \begin{cases} \frac{1 - e^{-At}}{1 - e^{-AR}}, & 0 \leq t \leq R, \quad A \neq 0 \\ t/R, & 0 \leq t \leq R, \quad A = 0 \end{cases}$$

A is interpreted as follows.

$A > 0$ results in a convex (faster than expected) entry distribution.

$A < 0$ results in a concave (slower than expected) entry distribution.

$A = 0$ results in the uniform entry distribution in which $g(t) = 1/R$.

Rather than specify A directly, **PASS** has you enter the percentage of the accrual time that will be needed to enroll 50% of the subjects. Using an iterative search, the value of A corresponding to this percentage is calculated and used in the calculations.

Losses to Follow-Up

The staggered patient entry over the accrual period results in censoring times ranging from $T - R$ to T years during the follow-up period. This is often referred to as administrative censoring, since it is caused by the conclusion of the study rather than by some random factor working on an individual. To model the losses to follow-up which come from other causes, we use the exponential distribution with hazard rates η_1 and η_2 . Since these rates are difficult and confusing to define directly, **PASS** lets you input the proportion lost due to other causes over a specified period of time and uses the following equation to determine the hazard rates:

$$\eta = -\frac{\log(1 - P(t))}{t}$$

General Model

Combining all these parameters into the model results in

$$\phi(\lambda, \eta, \gamma) = \lambda^2 \left(\frac{\lambda}{\lambda + \eta} + \frac{\lambda \gamma e^{-(\lambda + \eta)T} [1 - e^{(\lambda + \eta - \gamma)R}]}{(1 - e^{-\gamma R})(\lambda + \eta)(\lambda + \eta - \gamma)} \right)^{-1}.$$

This expression may then be used in an equation that relates these parameters to sample size and power

$$\sqrt{N}/|\lambda_1 - \lambda_2| = Z_\alpha \sqrt{\phi(\bar{\lambda}) \left(\frac{1}{Q_1} + \frac{1}{Q_2} \right)} + Z_\beta \sqrt{\frac{\phi(\lambda_1)}{Q_1} + \frac{\phi(\lambda_2)}{Q_2}}$$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *Proportion Surviving 1*, *Proportion Surviving 2*, *Accrual*, *Follow-Up*, *Alpha*, *Beta*, and *N*. Under most situations, you will select either *Power and Beta* or *N*.

Select *N* when you want to calculate the sample size needed to achieve a given power and alpha level.

Select *Power and Beta* when you want to calculate the power.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal survival curves when in fact the curves are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis of equal survival curves when in fact the curves are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size – Total Sample Size

N (Total Sample Size)

This is the combined sample size of both groups. This amount is divided between the two groups using the value of the Proportion in Group 1.

Sample Size – Sample Proportion

Proportion in Group 1

This is the proportion of N in group one. If this value is labeled p_1 , the sample size of group one is Np_1 and the sample size of group two is $N - Np_1$. Note that the value of Np_1 is rounded to the nearest integer.

Sample Size – Proportions Lost to Follow-Up

Group 1 and Group 2

This is the proportion in this group that is lost to follow-up during the Fixed Time Period, $T0$. Using this value, the lost-to-follow-up hazard rate is calculated using the exponential distribution as it is with $S(t)$. The equation used to convert these proportions into η_1 and η_2 is:

$$\eta_i = -\frac{\log(1 - P_i(T0))}{T0}$$

705-6 Logrank Tests (Lachin and Foulkes)

Here $P(T0)$ is the proportion lost to follow up from the beginning of the study until $T0$. So if $T0$ is 3 years and 10% of the patients in group one are lost to follow-up

$$\begin{aligned}\eta_1 &= -\log(1.0 - 0.1) / 3 \\ &= 0.03512\end{aligned}$$

Values between zero and one are valid. Zero is used to indicate no loss to follow-up.

Effect Size

S1 and S2 (Proportion Surviving Past T0)

Specify the proportion of patients in each group that survive until after $T0$. These quantities are called $S1$ and $S2$. The value of $T0$ is given in the Fixed Time Point box. For example, if the Fixed Time parameter is 3 and this value is set to 0.7, then 70% of the patients survive at least 3 years. Note that since this is a proportion, values must be between zero and one.

These quantities are used as a convenient method of entering the hazard rates λ_1 and λ_2 . The first group is arbitrarily designated as the control group and the second group is arbitrarily designated as the treatment group.

The hazard rates are calculated from the proportions surviving using the formulas

$$\lambda_1 = -\frac{\log(S1)}{T0} \text{ and } \lambda_2 = -\frac{\log(S2)}{T0}$$

If your sample size problem is cast in terms of hazard rates, the proportions surviving are calculated using the formula

$$S1 = \exp(-\lambda_1 t) \text{ and } S2 = \exp(-\lambda_2 t)$$

T0 (Fixed Time Point)

This is the time period used to convert the proportions given in Proportion Surviving 1, Proportion Surviving 2, Follow-Up Loss1, and Follow-Up Loss2 into hazard rates. If you enter 0.4 in Proportion Survive1 and a 3 here, you are indicating that 40% survive longer than 3 years.

Duration

R (Accrual Time)

The accrual time is the length of time during which patients enter the study. It is the value of R .

% Time Until 50% Accrual

This specifies the percentage of the accrual time needed to enroll 50% of the patients. This value is converted into a value for A , the patient entry parameter. Use this option to indicate how subjects are enrolled during the accrual period. For example, in one study, it may be possible to enroll a number of patients early on, while in another study, most of the subjects will be enrolled near the end of the accrual period.

Values between 1 and 97 may be entered.

If you expect uniform patient entry, enter 50. Unless you know that patient enrollment will not be uniform during the accrual period, you should use this amount.

If you expect more patients to enter during the early part of the accrual period, enter an amount less than 50 such as 30. A 30 here means that 50% of the patients will have been enrolled when 30% of the accrual time has elapsed.

If you expect more patients to enter during the later part of the accrual period, enter an amount greater than 50 such as 70. A 70 here means that 50% of the patients will have been enrolled when 70% of the accrual time has elapsed.

PASS assumes that patient entry times follow the truncated exponential distribution. This parameter controls the shape and scale of that distribution.

Follow-Up Time, T-R

The *follow-up time* is the length of time between the entry of the last individual into the study and the end of the study. Since T is the total length of the study and R is the accrual time, the follow-up time is $T-R$.

Test

Alternative Hypothesis

This option specifies the alternative hypothesis in terms of the proportion surviving in each group. This implicitly specifies the direction of the hypothesis test. The null hypothesis is always $H_0: S_1 = S_2$.

Note that the alternative hypothesis enters into power calculations by specifying the rejection region of the hypothesis test. Its accuracy is critical.

Possible selections are:

- **Ha: $S_1 \neq S_2$**
This is the usual selection. It yields the *two-tailed* t-test. Use this option when you are testing whether the survival curves are different, but you do not want to specify beforehand which curve is better.
- **Ha: $S_1 < S_2$**
This option yields a *one-tailed* test. Use it when you are only interested in the case in which the survival in group one is less than that for group two.
- **Ha: $S_1 > S_2$**
This option yields a *one-tailed* test. Use it when you are only interested in the case in which the survival in group one is greater than that for group two.

Example 1 – Finding the Power

A researcher is planning a clinical trial using a parallel, two-group, equal sample allocation design to compare the survivability of a new treatment with that of the current treatment. The proportion surviving one-year after the current treatment is 0.50. The new treatment will be adopted if the proportion surviving after one year increases to 0.75.

The trial will include a recruitment period of one-year after which participants will be followed for an additional two-years. It is assumed that patients will enter the study uniformly over the accrual period. The researcher estimates a loss-to-follow rate of 15% during the first year in both the control and the experimental groups.

The researcher decides to investigate various sample sizes between 10 and 250 at both the 0.01 and 0.05 significance levels.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lachin and Foulkes)** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lachin and Foulkes)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.01 0.05
N (Total Sample Size)	10 25 50 100 150 200 250
Proportion in Group 1	0.5
Proportion Lost to Follow-Up 1	0.15
Proportion Lost to Follow-Up 2	0.15
S1 (Proportion Surviving in Group 1)	0.5
S2 (Proportion Surviving in Group 2)	0.75
T0 (Fixed Time Point).....	1
R (Accrual Time)	1
% Time Until 50% Accrual.....	50
Follow-Up Time, T-R	2
Alternative Hypothesis	Ha: S1 <> S2

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results with Ha: S1<>S2

Power	N	Surviving Group 1	Surviving Group 2	Accrual Time	Follow Up Time	Prop in Group 1	Alpha	Beta
0.06718	10	0.50000	0.75000	1.00	2.00	0.50000	0.01000	0.93282
0.18406	10	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.81594
0.17527	25	0.50000	0.75000	1.00	2.00	0.50000	0.01000	0.82473
0.36633	25	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.63367
0.38357	50	0.50000	0.75000	1.00	2.00	0.50000	0.01000	0.61643
0.61606	50	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.38394
0.72756	100	0.50000	0.75000	1.00	2.00	0.50000	0.01000	0.27244
0.88428	100	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.11572
0.90273	150	0.50000	0.75000	1.00	2.00	0.50000	0.01000	0.09727
0.97052	150	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.02948
0.96998	200	0.50000	0.75000	1.00	2.00	0.50000	0.01000	0.03002
0.99328	200	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.00672
0.99167	250	0.50000	0.75000	1.00	2.00	0.50000	0.01000	0.00833
0.99858	250	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.00142

Base Time 1.00
 Proportion loss to follow up in group 1 0.15000
 Proportion loss to follow up in group 2 0.15000
 Percent of accrual time until 50% enrollment is reached: 50.00%

Report Definitions

Power is the probability of rejecting a false null hypothesis. Power should be close to one.
 Alpha is the probability of rejecting a true null hypothesis. It should be small.
 Beta is the probability of accepting a false null hypothesis. It should be small.

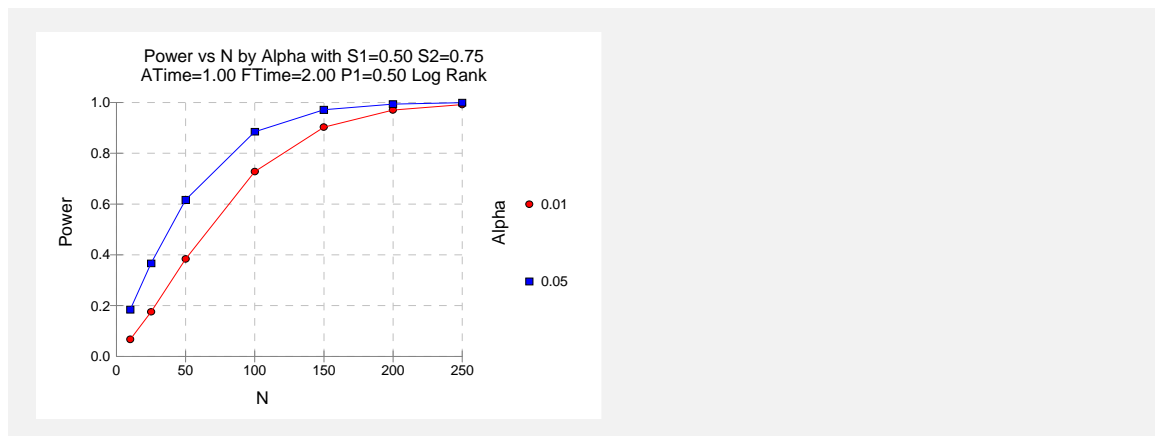
Summary Statements

A two-sided log rank test with an overall sample size of 10 subjects (of which 5 are in group 1 and 5 are in group 2) achieves 7% power at a 0.01000 significance level to detect a difference of 0.25000 between 0.50000 and 0.75000--the proportions surviving in groups 1 and 2 after 1.00 time periods. Patients entered the study during an accrual period of 2.00 time periods. 50% of the enrollment was complete when 50.00% of the accrual time had past. A follow-up period of 1.00 time periods had a 15.0% loss from group 1 and a 15.0% loss from group 2.

This report shows the values of each of the parameters, one scenario per row. Note that approximately 100 patients, 50 per group, will be needed to achieve about 90% power at the 0.05 significance level.

705-10 Logrank Tests (Lachin and Foulkes)

Plots Section



This plot shows the relationship between sample size and power for the two significance levels.

Example 2 – Finding the Sample Size

Continuing with the previous example, the researcher wants to investigate the sample size necessary to achieve 80% or 90% power for various values of the proportion surviving in the treatment group from 0.55 to 0.80 at the 0.05 significance level. All other parameters will remain the same.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lachin and Foulkes)** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lachin and Foulkes)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	0.8 0.9
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Group 1	0.5
Proportion Lost to Follow-Up 1	0.15
Proportion Lost to Follow-Up 2	0.15
S1 (Proportion Surviving in Group 1)	0.5
S2 (Proportion Surviving in Group 2)	0.55 to 0.80 by 0.05
T0 (Fixed Time Point).....	1
R (Accrual Time)	1
% Time Until 50% Accrual.....	50
Follow-Up Time, T-R	2
Alternative Hypothesis	Ha: S1 <> S2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots

Numeric Results with Ha: S1<>S2

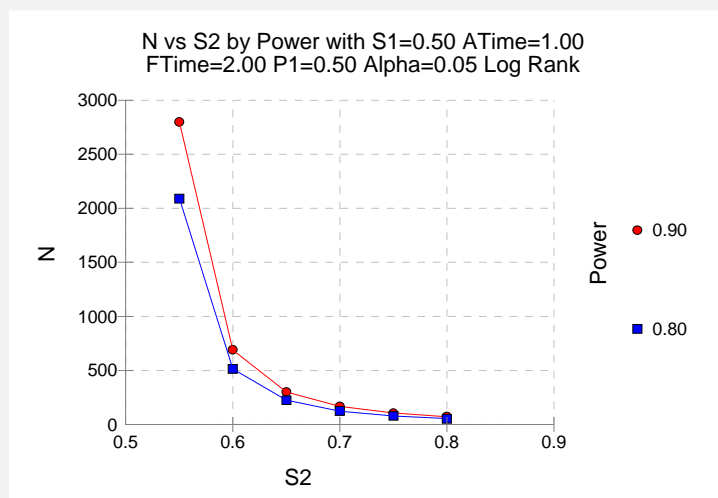
Power	N	Surviving Group 1	Surviving Group 2	Accrual Time	Follow Up Time	Prop in Group 1	Alpha	Beta
0.90004	2798	0.50000	0.55000	1.00	2.00	0.50000	0.05000	0.09996
0.80017	2090	0.50000	0.55000	1.00	2.00	0.50000	0.05000	0.19983
0.90024	690	0.50000	0.60000	1.00	2.00	0.50000	0.05000	0.09976
0.80050	515	0.50000	0.60000	1.00	2.00	0.50000	0.05000	0.19950
0.90001	302	0.50000	0.65000	1.00	2.00	0.50000	0.05000	0.09999
0.80010	225	0.50000	0.65000	1.00	2.00	0.50000	0.05000	0.19990
0.90098	168	0.50000	0.70000	1.00	2.00	0.50000	0.05000	0.09902
0.80177	125	0.50000	0.70000	1.00	2.00	0.50000	0.05000	0.19823
0.90107	106	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.09893
0.80357	79	0.50000	0.75000	1.00	2.00	0.50000	0.05000	0.19643
0.90274	73	0.50000	0.80000	1.00	2.00	0.50000	0.05000	0.09726
0.80432	54	0.50000	0.80000	1.00	2.00	0.50000	0.05000	0.19568

Base Time 1.00

Proportion loss to follow up in group 10.15000

Proportion loss to follow up in group 20.15000

Percent of accrual time until 50% enrollment is reached: 50.00%



This study shows the huge increase in sample size necessary to detect values of S_2 below 0.65. It also shows that roughly 35% more participants are required for 90% power than for 80% power in this situation.

Example 3 – Validation using Lachin and Foulkes

Lachin and Foulkes (1986), the developers of formulas used in this routine, give an example on page 509 in which $\lambda_1 = 0.30$, $\lambda_2 = 0.20$, $N = 378$, $\alpha = 0.05$ (one-sided), $R = 3$, $T = 3$, and $\beta = 0.10$. There is no loss to follow up and uniform patient entry is assumed.

The first step is to determine the proportion surviving at the end of one year for each group using the formula:

$$S(t) = e^{-\lambda t}$$

For group 1 we have

$$\begin{aligned} S_1(1) &= e^{-0.3(1)} \\ &= 0.74081822 \end{aligned}$$

For group 2 we have

$$\begin{aligned} S_2(1) &= e^{-0.2(1)} \\ &= 0.81873075 \end{aligned}$$

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lachin and Foulkes)** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests (Lachin and Foulkes)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find Setting</i>
Alpha	0.05
N (Total Sample Size)	378
Proportion in Group 1	0.5
Proportion Lost to Follow-Up 1	0.0
Proportion Lost to Follow-Up 2	0.0
S1 (Proportion Surviving in Group 1)	0.74081822
S2 (Proportion Surviving in Group 2)	0.81873075
T0 (Fixed Time Point).....	1
R (Accrual Time)	3
% Time Until 50% Accrual.....	50
Follow-Up Time, T-R	2
Alternative Hypothesis	Ha: S1 < S2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results with Ha: S1<S2

Power	N	Surviving Group 1	Surviving Group 2	Accrual Time	Follow Up Time	Prop in Group 1	Alpha	Beta
0.90123	378	0.74082	0.81873	3.00	2.00	0.50000	0.05000	0.09877

Base Time1.00

Proportion loss to follow up in group 10.00000

Proportion loss to follow up in group 20.00000

Percent of accrual time until 50% enrollment is reached: 50.00%

The power of 0.90 matches the value published in Lachin's article.

Example 4 – Non-Inferiority Test

This example will show how to use this module to calculate power and sample size for non-inferiority trials. Remember that non-inferiority is established by showing that the new treatment is no worse than the standard treatment, except for a small amount called the margin of equivalence.

Consider the following example. Suppose the median survival time of the standard drug is 15 months. Unfortunately, this drug has serious side effects. A promising new drug has been developed that has much milder side effects. A non-inferiority trial is to be designed to show that the new drug is not inferior to the standard drug. The margin of equivalence is set at 3 months, so, to establish non-inferiority, the study has to conclude that the median survival time of the new drug is at least 12 months.

The trial will accept subjects for 18 months. It will continue for an additional 6 months after the accrual period. Although the study planners anticipate some dropout, they want to begin their analysis without considering dropout. They set the significance level at 0.05. They want to determine the necessary sample size to achieve 80% and 90% power. The calculations proceed as follows.

Use the *Survival Parameter Conversion Tool* to convert the “median survival rates” to “proportions surviving” so that they can be entered into the panel. Press the **Survival Parameter Conversion Tool** button that is below the T0 box to display this window. Enter **24** for **T0**, **12** for **Median Survival Time (T1)**, and **15** for **Median Survival Time (T2)**. The screen will appear as follows.

705-14 Logrank Tests (Lachin and Foulkes)

Survival Parameter Conversion Tool

Survival Parameters | Proportions (Years to Months)

Use this window to convert among various survival distribution parameters (assumes an exponential survival distribution). The appropriate values may then be copied and pasted.

Survival Parameters

Time (T0): 24

Hazard Ratio (h1/h2): 1.25

Mortality Ratio (M1/M2): 1.49226322198727

Group 1

Hazard Rate (h1): 5.77622650466621E-02

Proportion Surviving Past T0 (S1): 0.25

Mortality until T0 (M1): 0.999999940395355

Median Survival Time (T1): 12

Changing a value will cause all appropriate values to be updated.

Hazard Ratio (h2/h1): 0.8

Mortality Ratio (M2/M1): 0.670123062249224

Group 2

Hazard Rate (h2): 4.62098120373297E-02

Proportion Surviving Past T0 (S2): 0.329876977693224

Mortality until T0 (M2): 0.670123022306776

Median Survival Time (T2): 15

HAZARD RATE:
This is the instantaneous hazard rate. Changing this value will cause the median survival time, the proportion surviving, and the hazard ratio to be updated.

Reset Close

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lachin and Foulkes)** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lachin and Foulkes)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.80 0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find Setting</i>
Proportion in Group 1	0.5
Proportion Lost to Follow-Up 1	0.0
Proportion Lost to Follow-Up 2	0.0
S1 (Proportion Surviving in Group 1)	0.25
S2 (Proportion Surviving in Group 2)	0.329876977693224
T0 (Fixed Time Point)	24
R (Accrual Time)	18
% Time Until 50% Accrual	50
Follow-Up Time, T-R	6
Alternative Hypothesis	Ha: S1 < S2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	N	Surviving Group 1	Surviving Group 2	Accrual Time	Follow Up Time	Prop in Group 1	Alpha	Beta
0.90018	1326	0.25000	0.32988	18.00	6.00	0.50000	0.05000	0.09982
0.80030	957	0.25000	0.32988	18.00	6.00	0.50000	0.05000	0.19970

Thus, a sample size of about 663 per group would be needed for 90% power, and a sample of 479 per group would be needed for 80% power.

Chapter 706

Logrank Tests for Non-Inferiority

Introduction

This module computes the sample size and power for *non-inferiority* tests under the assumption of proportional hazards. Accrual time and follow-up time are included among the parameters to be set. The non-inferiority logrank test is used for data analysis.

Sometimes, the objective of a study is to show that an experimental therapy is not inferior to (no worse than) the standard therapy. The experimental therapy may be cheaper, less toxic, or have fewer side effects. Such studies are often called non-inferiority trials and have a one-sided hypothesis.

Power and sample size calculations for the non-inferiority logrank test have been developed by Jung et al. (2005), and we use their results. These calculations assume an underlying exponential survival distribution with a uniform patient accrual pattern during the accrual period.

Technical Details

Test Statistic

Suppose a clinical trial consists of two independent groups. Designate group one as the standard group with hazard rate h_1 and sample size n_1 . Designate group two as the experimental group with hazard rate h_2 and sample size n_2 . The total sample size is $N = n_1 + n_2$. Usually, you would plan to have $n_1 = n_2$.

Define the proportion of the total sample in each group as

$$Q_i = \frac{n_i}{N}, \quad i = 1, 2$$

Individuals are recruited during an accrual period of R years (or months or days). They are followed for an additional period of time until a total of T years is reached. Hence, the follow-up period is $T-R$ years. At the end of the study, the non-inferiority logrank test is conducted at significance level α with power $1 - \beta$. Under the proportion hazards assumption, the hazard ratio $HR = h_2 / h_1$ is constant across time.

706-2 Logrank Tests for Non-Inferiority

For a given non-inferiority margin $HR_0 (>1)$ (the maximum ratio of clinical insignificance), the statistical hypotheses tested are

$$H_0 : HR \geq HR_0 \text{ vs. } H_1 : HR < HR_0$$

Define the partial score function as

$$W(HR) = HR \sum_{i=1}^{n_1} \frac{\delta_{1i} \sum_{j=1}^{n_2} I(X_{2j} \geq X_{1i})}{\sum_{j=1}^{n_1} I(X_{1j} \geq X_{1i}) + HR \sum_{j=1}^{n_2} I(X_{2j} \geq X_{1i})} - \sum_{i=1}^{n_2} \frac{\delta_{2i} \sum_{j=1}^{n_1} I(X_{1j} \geq X_{2i})}{\sum_{j=1}^{n_1} I(X_{1j} \geq X_{2i}) + HR \sum_{j=1}^{n_2} I(X_{2j} \geq X_{2i})}$$

and the information function as

$$\sigma_N^2(HR) = HR \sum_{k=1}^2 \sum_{i=1}^{n_k} \frac{\delta_{ki} \left\{ \sum_{j=1}^{n_1} I(X_{1j} \geq X_{ki}) \right\} \left\{ \sum_{j=1}^{n_2} I(X_{2j} \geq X_{ki}) \right\}}{\left\{ \sum_{j=1}^{n_1} I(X_{1j} \geq X_{ki}) + HR \sum_{j=1}^{n_2} I(X_{2j} \geq X_{ki}) \right\}^2}$$

where X_{ki} is the minimum of the survival time, and the censoring time, δ_{ki} , is an event indicator taking 1 if there was an event or 0 otherwise, and $I(\cdot)$ is an indicator function. Note that $W(1)$ is the standard logrank test statistic.

Under H_0 , $W(HR_0)/\sigma_N(HR_0)$ is asymptotically normal with mean 0 and variance 1. Reject H_0 in favor of H_1 if $W(HR_0)/\sigma_N(HR_0) > z_{1-\alpha}$ with one-sided type I error probability α .

The partial MLE, \hat{HR} , is obtained by solving $W(HR) = 0$. Let HR^* denote the true value of HR . It can be shown that \hat{HR} is asymptotically normal with mean HR^* and variance $\sigma_N^{-2}(HR^*)$.

An asymptotic $100(1-\alpha)\%$ confidence interval for HR is $\hat{HR} \pm z_{1-\alpha/2} \sigma_N^{-1}(\hat{HR})$.

Power Calculations

Jung (2005) shows that the power of the non-inferiority logrank test can be expressed as

$$1 - \beta = \Phi \left(\frac{(HR_0 - 1)DQ_1Q_2 - z_{1-\alpha}\sqrt{HR_0}}{Q_1 + Q_2HR_0} \right)$$

where D is the observed number of deaths (events). The total sample size N is obtained by inflating D according to the relationship $E(d)N = D$, where $E(d)$ is the expected death rate for the trial.

Following the proposal of Yateman and Skene (1992) and the results of Lakatos (1988), we compute $E(d)$ using the Markov Model given in chapter 715 as

$$E(d) = Q_1S_{1,2} + Q_2S_{2,2}$$

where $S_{1,2}$ and $S_{2,2}$ are the occupancy probabilities for the event state for the standard and experimental groups, respectively. This formulation allows the inclusion of loss to follow-up, noncompliance, and drop-in along with various accrual patterns.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *Power and Beta* or *N*. Select *N* when you want to calculate the sample size needed to achieve a given power and alpha level. Select *Power and Beta* when you want to calculate the power.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you do not reject the null hypothesis that the hazard ratio is greater than HR_0 when in fact it is.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis that the hazard ratio is greater than HR_0 when in fact it is not.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Total Sample Size)

This is the combined sample size of both groups. This amount is divided between the two groups using the value of the Proportion in Reference Group.

Proportion in Reference Group

This is the proportion of N in the reference (control) group. If this value is labeled Q_1 , the sample size of group one is NQ_1 and the sample size of group two is $N - NQ_1$. Note that the value of NQ_1 is rounded to the nearest integer.

Proportion Lost or Switching Groups

Reference (or Treatment) Lost

This is the proportion of subjects in the reference (treatment) group that disappear from the study during a single time period (month, year, etc.). Multiple entries, such as *0.01 0.03 0.05*, are allowed.

When you want to specify different proportions for different time periods, you would enter those rates into a column of the spreadsheet, one row per time period. You specify the column of the spreadsheet by beginning your entry with an equals sign. For example, if you have entered the proportions in column 5, you would enter `=C5` here.

Reference Switching to Treatment

This is the proportion of subjects in the reference group that change to a treatment regime similar in efficacy to the treatment group during a single time period (month, year, etc.). This is sometimes referred to as *drop in*. Multiple entries, such as *0.01 0.03 0.05*, are allowed.

When you want to specify different proportions for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column of the spreadsheet by beginning your entry with an equals sign. For example, if you have entered the proportions in column 1, you would enter `=C1` here.

Treatment Switching to Reference

This is the proportion of subjects in the treatment group that change to a treatment regime similar in efficacy to the reference group during a single time period (month, year, etc.). This is sometimes referred to as *noncompliance*. Multiple entries, such as *0.01 0.03 0.05*, are allowed.

When you want to specify different proportions for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column of the spreadsheet by beginning your entry with an equals sign. For example, if you have entered the proportions in column 2, you would enter `=C2` here.

Effect Size

HR0 (Equivalence)

This is the maximum value of the hazard ratio that still results in the conclusion of equivalence. It is referred to as *HR0* in this documentation. Assuming that events are bad (such as death), then this number should be $> one$.

For example, if you enter 1.20 here, you are saying that hazard ratios < 1.20 will result in the conclusion of non-inferiority when H_0 is rejected. In other words, hazard ratios up to 1.20 indicate that the treatment group is no worse than the reference group. Estimates of the hazard ratio may be obtained from median survival times, hazard rates, or from the proportion surviving past a certain time point. Pressing the Parameter Conversion button will load a tool for doing this.

h1 (Hazard Rate of Reference Group)

Specify one or more hazard rates (instantaneous failure rate) for the reference group. For an exponential distribution, the hazard rate is the inverse of the mean survival time. An estimate of the hazard rate may be obtained from the median survival time or from the proportion surviving to a certain time point. This calculation is automated by pressing the *Parameter Conversion* button.

Hazard rates must be greater than zero. Constant hazard rates are specified by entering them directly. Variable hazard rates are specified as columns of the spreadsheet. When you want to specify different hazard rates for different time periods, you would enter those rates into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the hazard rates in column 2, you would enter $=2$ here.

Duration

Accrual Time (Integers Only)

Enter one or more values for the number of time periods (months, years, etc.) during which subjects are entered into the study. The total duration of the study is equal to the Accrual Time plus the Follow-Up Time. These values must be integers.

Accrual times can range from 0 to the Total Time. That is, the accrual time must be less than or equal to the Total Time. Otherwise, the scenario is skipped.

Enter 0 when all subjects begin the study together.

Accrual Pattern

This contains the pattern of accrual (patient entry). Two types of entries are possible:

- **Uniform**

If you want to specify a uniform accrual rate for all time periods, enter *Equal* here.

- **Non-Uniform**

When you want to specify accrual patterns with different accrual proportions per time period, you would enter the pattern into a column of the spreadsheet, one row per time period and specify the column, or columns, here, beginning your entry with an equals sign. For example, if you have entered accrual patterns in columns 4 and 5, you would enter $=C4\ C5$.

Note that these values are standardized to sum to one. Thus, the accrual pattern $0.25\ 0.50\ 0.25$ would result in the same accrual pattern as $1\ 2\ 1$ or $25\ 50\ 25$.

Total Time (Integers Only)

Enter one or more values for the number of time periods (months, years, etc.) in the study. The follow-up time is equal to the Total Time minus the Accrual Time. These values must be integers.

Reports Tab

The Reports tab contains additional settings for this procedure.

Report Column Width

Report Column Width

This option sets the width of the each column of the numeric report.

The numeric report for this option necessarily contains many columns, so the maximum number of decimal places that can be displayed is four. If you try to increase that number, the numbers may run together. You can increase the width of each column using this option.

The recommended report column width for scenarios without large numbers of decimal places or extremely large sample sizes is 0.49.

Options Tab

The Options tab contains additional settings for this procedure.

Options

Number of Intervals within a Time Period

The algorithm requires that each time period be partitioned into a number of equal-width intervals. Each of these subintervals is assumed to follow an exponential distribution. This option controls the number of subintervals. All parameters such as hazard rates, loss to follow-up rates, and noncompliance rates are assumed to be constant within a subinterval.

Lakatos (1988) gives little input as to how the number of subintervals should be chosen. In a private communication, he indicated that 100 ought to be adequate. This seems to work when the hazard rate is less than 1.0.

As the hazard rate increases above 1.0, this number must increase. A value of 2000 should be sufficient as long as the hazard rates (h_1 and h_2) are less than 10. When the hazard rates are greater than 10, you may want to increase this value to 5000 or even 10000.

Example 1 – Finding the Power

A non-inferiority trial is planned in which the primary analysis will use the non-inferiority logrank test. After extensive discussion, the researchers have decided that the upper bound on non-inferiority is 1.3.

The trial will include a recruitment period of two-years after which participants will be followed for three more years. It is assumed that patients will enter the study uniformly over the accrual period. The researcher estimates a loss-to-follow rate of 5% per year in both the reference and experimental groups. Past experience leads to a base line hazard rate of 0.04. An equal sample allocation design will be used with a target power of 0.90 and significance level of 0.05.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests for Non-Inferiority** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests for Non-Inferiority**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Total Sample Size)	1000 to 5000 by 1000
Proportion in Reference Group	0.5
Proportion Lost - Reference	0.05
Proportion Lost - Treatment	0.0
HR0	1.3
HR	1.0
h1	0.04
Accrual Time	2
Accrual Pattern.....	Equal
Total Time	5
Reports Tab	
Show Detail Numeric Reports	Checked

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results in Terms of Sample Size

Power	N1	N2	N	Equiv Haz Ratio (HR0)	Actual Haz Ratio (HR)	Ref Haz Rate (h1)	Acc- rual Pat'n	Acc- rual Time/ Total Time	Ref Loss	Trt Loss	Ref to Trt	Trt to Ref	Alpha	Beta
.4510	500	500	1000	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.5490
.6920	1000	1000	2000	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.3080
.8366	1500	1500	3000	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.1634
.9169	2000	2000	4000	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.0831
.9591	2500	2500	5000	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.0409

References

Jung, Sin-Ho; Kang, Sun J.; McCall, Linda M.; Blumenstein, Brent. 2005. 'Sample Sizes Computation for Two-Sample Noninferiority Log-Rank Test', J. of Biopharmaceutical Statistics, Volume 15, pages 969-979.
Lakatos, Edward. 1988. 'Sample Sizes Based on the Log-Rank Statistic in Complex Clinical Trials', Biometrics, Volume 44, March, pages 229-241.

Report Definitions

Power is the probability of rejecting a false null hypothesis. Power should be close to one.

N1|N2|N are the sample sizes of the reference group, treatment group, and both groups, respectively.

E1|E2|E are the number of events in the reference group, the treatment group, and both groups, respectively.

Equivalence Haz Ratio (HR0) is the upper bound for the hazard ratio that still leads to the conclusion of non-inferiority.

Actual Haz Ratio (HR) is assumed to be the actual value of the hazard ratio.

Ref Haz Rate is the hazard (instantaneous failure) rate of the reference group. Its scale is events per time period.

Accrual Time is the number of time periods (years or months) during which accrual takes place.

Total Time is the total number of time periods in the study. Follow-up time = (Total Time) - (Accrual Time).

Ref Loss is the proportion of the reference group that is lost (drop out) during a single time period (year or month).

Trt Loss is the proportion of the treatment group that is lost (drop out) during a single time period (year or month).

Ref to Trt (drop in) is the proportion of the reference group that switch to a group with a hazard rate equal to the treatment group.

Trt to Ref (noncompliance) is the proportion of the treatment group that switch to a group with a hazard rate equal to the reference group.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

Beta is the probability of accepting a false null hypothesis. It should be small.

This report shows the values of each of the parameters, one scenario per row. We see that almost 4000 subjects will be required for this study.

Next, a report displaying the number of required events rather than the sample size is displayed.

Numeric Results in Terms of Events

Power	Ref Evts E1	Trt Evts E2	Total Evts E	Equiv Haz Ratio (HR0)	Actual Haz Ratio (HR)	Ref Haz Rate (h1)	Acc- rual Pat'n	Acc- rual Time/ Total Time	Ref Loss	Trt Loss	Ref to Trt	Trt to Ref	Alpha	Beta
.4510	66.8	66.8	133.6	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.5490
.6920	133.6	133.6	267.3	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.3080
.8366	200.4	200.4	400.9	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.1634
.9169	267.3	267.3	534.5	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.0831
.9591	334.1	334.1	668.1	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.0409

Most of this report is identical to the last report, except that the sample sizes are replaced by the number of required events.

Next, reports displaying the individual settings year-by-year for each scenario are displayed.

Detailed Input when Power=.4510 HR0=1.30 HR=1.00 N=1000 Alpha=.0500 Accrual/Total Time=2 / 5

Time Period	Reference Hazard Rate (H1)	Percent Accrual (Equal)	Percent Admin. Censored (Calc.)	Percent Reference Loss (.0500)	Percent Treatment Loss (.0500)	Switch Reference to Treatment (.0000)	Switch Treatment to Reference (.0000)
1	.0400	50.00	.00	.0500	.0500	.0000	.0000
2	.0400	50.00	.00	.0500	.0500	.0000	.0000
3	.0400	.00	.00	.0500	.0500	.0000	.0000
4	.0400	.00	50.00	.0500	.0500	.0000	.0000
5	.0400	.00	100.00	.0500	.0500	.0000	.0000

This report shows the individual settings for each time period (year). It becomes very useful when you want to document a study in which these parameters vary from year to year.

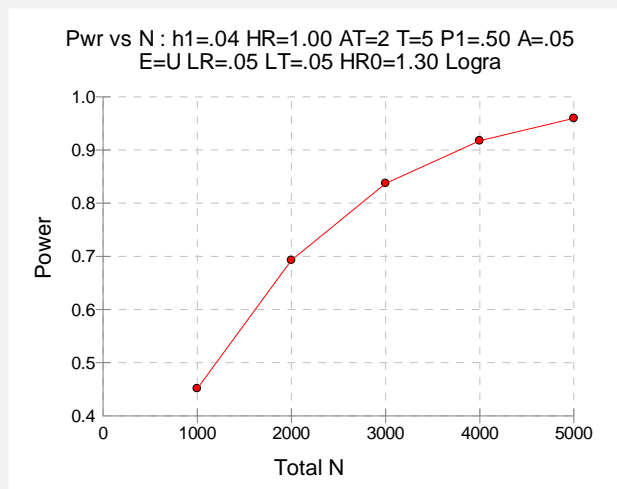
Next, summary statements are displayed.

Summary Statements

A non-inferiority logrank test with an overall sample size of 1000 subjects (500 in the reference group and 500 in the treatment group) achieves 45.1% power at a 0.050 significance level to detect an equivalence hazard ratio of 1.30 when the actual hazard ratio is an equivalence hazard ratio of 1.00 and the reference group hazard rate is 0.0400. The study lasts for 5 time periods of which subject accrual (entry) occurs in the first 2 time periods. The accrual pattern across time periods is uniform (all periods equal). The proportion dropping out of the reference group is 0.0500. The proportion dropping out of the treatment group is 0.0500. The proportion switching from the reference group to another group with a hazard rate equal to the treatment group is 0.0000. The proportion switching from the treatment group to another group with a hazard rate equal to the reference group is 0.0000.

Finally, a scatter plot of the results is displayed.

Plots Section



This plot shows the relationship between sample size and power. Note that for 90% power, a total sample size of about 4000 is required. The exact number will be found in Example 2.

Example 2 – Finding the Sample Size

Continuing with the previous example, the researcher wants to investigate the sample sizes necessary to achieve 80% and 90% power. All other parameters will remain the same as in Example 1.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests for Non-Inferiority** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests for Non-Inferiority**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.8 0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Reference Group	0.5
Proportion Lost - Reference	0.05
Proportion Lost - Treatment	0.05
HR0	1.3
HR	1.0
h1	0.04
Accrual Time	2
Accrual Pattern.....	Equal
Total Time	5

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results in Terms of Sample Size														
Power	N1	N2	N	Equiv	Actual	Ref	Acc- rual	Acc- rual	Ref	Trt	Ref to Trt	Trt to Ref	Alpha	Beta
				Haz Ratio (HR0)	Haz Ratio (HR)	Rate (h1)		Time/ Total Time						
.9000	1865	1866	3731	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.1000
.8000	1344	1345	2689	1.30	1.00	.0400	Equal	2 / 5	.0500	.0500	.0000	.0000	.0500	.2000

Example 3 – Validation using Jung

Jung et al. (2005) pages 974-975 present an example that will be used to validate this procedure. In this article, an 8.8-year trial is presented in which patient accrual occurs the first 3.8 years. The baseline hazard rate is 0.0446. The value of HR_0 is 1.3 and the value of HR is 1.0. Equal allocation between groups is used and uniform accrual is assumed. The significance level is 0.05 and the desired power is 0.90. Given these values, the number of events is found to be 499 and the sample size is 1891.

Since this procedure using integer values for the accrual and trial time, the accrual time and total time will be set to 4 and 9 years, respectively.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests for Non-Inferiority** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests for Non-Inferiority**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Reference Group	0.5
Proportion Lost - Reference	0
Proportion Lost - Treatment	0
HR_0	1.3
HR	1.0
h_1	0.0446
Accrual Time	4
Accrual Pattern.....	Equal
Total Time	9

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results in Terms of Sample Size														
Power	N1	N2	N	Equiv Haz Ratio (HR0)	Actual Haz Ratio (HR)	Ref Haz Rate (h1)	Acc- rual Pat'n	Acc- rual Time/ Total Time	Ref Loss	Trt Loss	Ref to Trt	Trt to Ref	Alpha	Beta
.9000	933	933	1866	1.30	1.00	.0446	Equal	4 / 9	.0000	.0000	.0000	.0000	.0500	.1000

Numeric Results in Terms of Events														
Power	Ref Evts E1	Trt Evts E2	Total Evts E	Equiv Haz Ratio (HR0)	Actual Haz Ratio (HR)	Ref Haz Rate (h1)	Acc- rual Pat'n	Acc- rual Time/ Total Time	Ref Loss	Trt Loss	Ref to Trt	Trt to Ref	Alpha	Beta
.9000	249.3	249.3	498.6	1.30	1.00	.0446	Equal	4 / 9	.0000	.0000	.0000	.0000	.0500	.1000

Note that the number of events (499) matches Jung's results exactly. The sample size of 1866 is slightly less than Jung's 1891. This difference occurs because these results were obtained for 4 years of accrual, not 3.8, and because we used Lakatos' method for transforming the number of events into the sample size.

Example 4 – Inputting Time-Dependent Hazard Rates from a Spreadsheet

Time-dependent parameters (hazard rates, losses to follow-up, etc) may be entered. Example 4 of Chapter 715 presents an extensive example of how this is done for the logrank test, so we will not duplicate that documentation here.

Chapter 710

Group-Sequential Logrank Tests

Introduction

Clinical trials are longitudinal. They accumulate data sequentially through time. The participants cannot be enrolled and randomized on the same day. Instead, they are enrolled as they enter the study. It may take several years to enroll enough patients to meet sample size requirements. Because clinical trials are long term studies, it is in the interest of both the participants and the researchers to monitor the accumulating information for early convincing evidence of either harm or benefit. This permits early termination of the trial.

Group sequential methods allow statistical tests to be performed on accumulating data while a phase III clinical trial is ongoing. Statistical theory and practical experience with these designs have shown that making four or five *interim analyses* is almost as effective in detecting large differences between treatment groups as performing a new analysis after each new data value. Besides saving time and resources, such a strategy can reduce the experimental subject's exposure to an inferior treatment and make superior treatments available sooner.

When repeated significance testing occurs on the same data, adjustments have to be made to the hypothesis testing procedure to maintain overall significance and power levels. The landmark paper of Lan & DeMets (1983) provided the theory behind the *alpha spending function* approach to group sequential testing. This paper built upon the earlier work of Armitage, McPherson, & Rowe (1969), Pocock (1977), and O'Brien & Fleming (1979). **PASS** implements the methods given in Reboussin, DeMets, Kim, & Lan (1992) to calculate the power and sample sizes of various group sequential designs.

This module calculates sample size and power for group sequential designs used to compare two survival curves. Other modules perform similar analyses for the comparison of means and proportions. The program allows you to vary number and times of interim tests, type of alpha spending function, and test boundaries. It also gives you complete flexibility in solving for power, significance level, sample size, or effect size. The results are displayed in both numeric reports and informative graphics.

Technical Details

In many clinical trials, patients are recruited and randomized to receive a particular treatment, either experimental or control, and then monitored until either a critical event occurs or the study is ended. The length of time the patient is monitored until the critical event occurs is called the *follow-up time*. After the study is ended, the follow-up times of the patients in the two groups are compared using the *logrank test* in what is often called a *survival analysis*.

When the critical event does not occur for a patient by the time the study is ended, the follow-up time is said to have been *censored*. Although the actual event time is not known for this patient, it is known that the event time will be greater than the follow-up time. Hence, some information is gleaned from these participants. Because of this censoring, the usual tests of means or proportions cannot be used. The logrank test was developed to provide a statistical test comparing the efficacy of the two treatments.

The Hazard Ratio (HR)

Suppose the critical event is death. The survival distribution of each treatment can be characterized by the *instantaneous death rates*, λ_1 and λ_2 . An instantaneous death rate, often called the *hazard*, is the probability of death in a short interval of time. The comparison of the efficacies of the two treatments is often formalized by considering the ratio of the hazard, or *hazard ratio* (HR).

$$HR = \frac{\lambda_2}{\lambda_1}$$

Although the logrank test concerns the ratio of the hazard rates of the two groups, when planning a study, it may be easier to obtain information about the expected proportion surviving during the trial. It turns out that the hazard ratio can be computed from the survival proportions, $S1$ and $S2$, using the equation

$$HR = \frac{\log(S2)}{\log(S1)}$$

when the hazard ratio is constant through time.

Assuming that group one is the control group, it may be easiest during the planning stages of a study to find $S1$ and state the minimum value of HR that would make the experimental treatment useful. The last equation can then be manipulated to calculate a value for $S2$ as follows

$$S2 = \exp\{HR(\log(S1))\}.$$

Sometimes it is more convenient to state hazard ratio in terms of the median survival times. In this case, the hazard ratio is estimated using

$$HR = \frac{M_1}{M_2}$$

when the hazard ratio is constant for different times.

The Logrank Statistic

The following results are excerpted from Reboussin (1992). The logrank statistic is given by the equation

$$L(d) = \sum_{i=1}^d \left(\frac{x_i r_{ic}}{r_{ic} + r_{it}} \right)$$

where d is the number of events, x_i is 1 if the event at time t_i is in the control group and 0 if it is in the treatment group, r_{ic} is the number of patients in the control group at risk just before t_i , and r_{it} is the corresponding number of patients at risk in the treatment group.

If $r_{ic} \approx r_{it}$ and HR is close to 1, then the sequential logrank statistic is (approximately)

$$z_k = 2 \frac{L(d_k)}{\sqrt{d_k}}.$$

The subscript k indicates that the computations use all data that are available at the time of the k^{th} interim analysis or k^{th} look (k goes from 1 to K).

Spending Functions

Lan and DeMets (1983) introduced alpha spending functions, $\alpha(\tau)$, that determine a set of boundaries b_1, b_2, \dots, b_K for the sequence of test statistics z_1, z_2, \dots, z_K . These boundaries are the critical values of the sequential hypothesis tests. That is, after each interim test, the trial is continued as long as $|z_k| < b_k$. When $|z_k| \geq b_k$, the hypothesis of equal means is rejected and the trial is stopped early.

The time argument τ either represents the proportion of elapsed time to the maximum duration of the trial or the proportion of the sample that has been collected. When elapsed time is used, it is referred to as *calendar time*. When time is measured in terms of the sample, it is referred to as *information time*. Since it is a proportion, τ can only vary between zero and one.

Alpha spending functions have the characteristics:

$$\alpha(0) = 0$$

$$\alpha(1) = \alpha$$

The last characteristic guarantees a fixed α level when the trial is complete. That is,


$$\Pr(|z_1| \geq b_1 \text{ or } |z_2| \geq b_2 \text{ or } \dots \text{ or } |z_K| \geq b_K) = \alpha(\tau)$$

This methodology is very flexible since neither the times nor the number of analyses must be specified in advance. Only the functional form of $\alpha(\tau)$ must be specified.


PASS provides five popular spending functions plus the ability to enter and analyze your own boundaries. These are calculated as follows:

710-4 Group-Sequential Logrank Tests


1. O'Brien-Fleming $2 - 2\Phi\left(\frac{Z_{\alpha/2}}{\sqrt{\tau}}\right)$




2. Pocock $\alpha \ln(1 + (e - 1)\tau)$




3. Alpha * time $\alpha\tau$



4. Alpha * time^{1.5} $\alpha\tau^{3/2}$



5. Alpha * time² $\alpha\tau^2$



6. User Supplied

A custom set of boundaries may be entered.

The O'Brien-Fleming boundaries are commonly used because they do not significantly increase the overall sample size and because they are conservative early in the trial. Conservative in the sense that the means must be extremely different before statistical significance is indicated. The Pocock boundaries are nearly equal for all times. The Alpha*t boundaries use equal amounts of alpha when the looks are equally spaced. You can enter your own set of boundaries using the User Supplied option.

Sequential Theory

A detailed account of the methodology is contained in Lan & DeMets (1983), DeMets & Lan (1984), Lan & Zucker (1993), and DeMets & Lan (1994). A brief summary of the theoretical basis of the method will be presented here.

Group sequential procedures for interim analysis are based on their equivalence to discrete boundary crossing of a Brownian motion process with drift parameter θ . The test statistics z_k follow the multivariate normal distribution with means $\theta\sqrt{\tau_k}$ and, for $j \leq k$, covariances $\sqrt{\tau_k / \tau_j}$. The drift is related to the parameters of the z-test through one of the equations

$$\theta = \frac{\log(HR)\sqrt{d_k}}{2} \quad (\text{exponential survival}) \quad \text{or} \quad \theta = \frac{|1 - HR|\sqrt{d_k}}{(1 + HR)} \quad (\text{proportional hazards}).$$

710-6 Group-Sequential Logrank Tests

These equations may be solved for d_k , the required number of events, giving

$$d_k = \frac{4\theta^2}{[\log(HR)]^2} \text{ (exponential survival) or } d_k = \left(\frac{(1 + HR)\theta}{1 - HR} \right)^2 \text{ (proportional hazards).}$$

In survival analysis, the size of a sample is measured in terms of number of events rather than number of patients because it is probable that many of the patients will be censored—their event times are not known. In order to ensure that the sample size produces the required number of events, it must be inflated by the event rates.

The expected number of events can be computed from the proportion surviving using the equation

$$d_k = \frac{N(1 - S_1) + N(1 - S_2)}{2}$$

where N is the total sample size (assumed to be split evenly between groups). This can be solved for N to give the sample size as

$$N = \frac{2d_k}{2 - S_1 - S_2}.$$

Hence, the algorithm is as follows:

1. Compute boundary values based on a specified spending function and alpha value.
2. Calculate the drift parameter based on those boundary values and a specified power value.
3. Use the drift parameter and the above equation to calculate the appropriate event size per group d_k .
4. Use the event size to compute the appropriate sample size, N .

Procedure Tabs

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains the parameters associated with the z-test such as the survival rates, sample sizes, alpha, and power.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are $S1$, $S2$, $Alpha$, $Power$ and $Beta$, or N . Under most situations, you will select either $Power$ and $Beta$ or N .

Select N when you want to calculate the sample size needed to achieve a given power and alpha level.

Select *Power and Beta* when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal survival curves when in fact the curves are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis of equal survival curves when in fact the curves are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Total Sample Size)

Enter a value (or range of values) for the total sample size. Each group is assumed to have a sample size of $N/2$. You may enter a range of values such as *100 to 1000 by 100*.

Note that the sample size is based implicitly on the length of the study since we used the equation

$$N = \frac{2d_k}{2 - S_1 - S_2}$$

to estimate the necessary sample size.

Test

Alternative Hypothesis

Specify whether the test is one-sided or two-sided. When a two-sided hypothesis is selected, the value of alpha is halved. Everything else remains the same.

Note that the accepted procedure is to use Two-Sided unless you can justify using a one-sided test.

Survival Time Assumption

Calculations can be based on the assumption of exponential survival times or proportional hazards. The proportional hazards assumption is the less restrictive and the one that is recommended.

- **Exponential Survival**

The survival curves follow exponential distributions. This is the most restrictive assumption.

- **Proportional Hazards**

The hazards of the two groups are proportional.

Effect Size

S1 (Proportion Surviving 1)

$S1$ is the proportion of patients belonging to group 1 (controls) that are expected to survive during the study. Since $S1$ is a proportion, it must be between zero and one. In many studies, this is the proportion surviving in the regular population with the standard treatment.

You may enter a range of values such as 0.1 , 0.2 , 0.3 or 0.1 to 0.9 by 0.2 .

S2 (Proportion Surviving 2)

$S2$ is the proportion of patients in the experimental group that survive during the study. This is not necessarily the expected proportion. Rather, you may set it to that proportion that, if achieved, would be of special interest. Values below (or above) this amount would not be of interest. For example, if the standard one-year survival proportion is 0.2 and the new treatment raises this proportion to 0.3 (a 50% increase in the proportion surviving), others may be interested in adopting this new treatment.

Sometimes, researchers wish to state the alternative hypothesis in terms of the hazard ratio, HR , rather than the value of $S2$. Using the fact that

$$HR = \frac{\log(S2)}{\log(S1)},$$

An appropriate value for $S2$ may be calculated from $S1$ and HR using the equation

$$S2 = \exp\{HR(\log(S1))\}.$$

Sometimes it is more convenient to state hazard ratio in terms of the median survival times. In this case, the hazard ratio is estimated using

$$HR = \frac{M_1}{M_2}$$

and the above equation is used to find $S2$.

Since $S2$ is a proportion, it must remain between zero and one. You may enter a range of values such as 0.1 , 0.2 , 0.3 or 0.1 to 0.9 by 0.2 .

Look Details

The Sequential tab contains the parameters associated with Group Sequential Design such as the type of spending function, the times, and so on.

Number of Looks

This is the number of interim analyses (including the final analysis). For example, a five here means that four interim analyses will be run in addition to the final analysis.

Boundary Truncation

You can truncate the boundary values at a specified value. For example, you might decide that no boundaries should be larger than 4.0. If you want to implement a boundary limit, enter the value here.

If you do not want a boundary limit, enter *None* here.

Spending Function

Specify which alpha spending function to use. The most popular is the O'Brien-Fleming boundary that makes early tests very conservative. Select *User Specified* if you want to enter your own set of boundaries.

Max Time

This is the total running time of the trial. It is used to convert the values in the Times box to fractions. The units (months or years) do not matter, as long as they are consistent with those entered in the Times box.

For example, suppose Max Time = 3 and Times = 1, 2, 3. Interim analyses would be assumed to have occurred at 0.33, 0.67, and 1.00.

Times

Enter a list of time values here at which the interim analyses will occur. These values are scaled according to the value of the Max Time option.

For example, suppose a 48-month trial calls for interim analyses at 12, 24, 36, and 48 months. You could set Max Time to 48 and enter *12,24,36,48* here or you could set Max Time to *1.0* and enter *0.25,0.50,0.75,1.00* here.

The number of times entered here must match the value of the Number of Looks.

- **Equally Spaced**

If you are planning to conduct the interim analyses at equally spaced points in time, you can enter *Equally Spaced* and the program will generate the appropriate time values for you.

Informations

You can weight the interim analyses on the amount of information obtained at each time point rather than on actual calendar time. If you would like to do this, enter the information amounts here. Usually, these values are the sample sizes obtained up to the time of the analysis.

For example, you might enter *50, 76, 103, 150* to indicate that 50 individuals were included in the first interim analysis, 76 in the second, and so on.

Upper and Lower Boundaries (Spending = User)

If the Spending Function is set to *User Supplied* you can enter a set of lower test boundaries, one for each interim analysis. The lower boundaries should be negative and the upper boundaries should be positive. Typical entries are 4,3,3,3,2 and 4,3,2,2,2.

- **Symmetric**

If you only want to enter the upper boundaries and have them copied with a change in sign to the lower boundaries, enter *Symmetric* for the lower boundaries.

Bnd Plot Axes Tab

The Bnd Plot Axes tab, short for Boundary Plot Axes tab, allows the axes of the spending function plots to be set separately from those of the power plots. The options are identical to those of the Axes tab.

Options Tab

The Options tab controls the convergence of the various iterative algorithms used in the calculations.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations to be run before the search for the criterion of interest (Alpha, Beta, etc.) is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank.

Recommended: 500 (or more).

Maximum Iterations (Lan-Demets algorithm)

This is the maximum number of iterations used in the Lan-DeMets algorithm during its search routine. We recommend a value of at least 200.

Tolerance

Probability Tolerance

During the calculation of the probabilities associated with a set of boundary values, probabilities less than this are assumed to be zero.

We suggest a value of 0.00000000001.

Power Tolerance

This is the convergence level for the search for the spending function values that achieve a certain power. Once the iteration changes are less than this amount, convergence is assumed. We suggest a value of 0.0000001.

If the search is too time consuming, you might try increasing this value.

Alpha Tolerance

This is the convergence level for the search for a given alpha value. Once the changes in the computed alpha value are less than this amount, convergence is assumed and iterations stop. We suggest a value of 0.0001.

This option is only used when you are searching for alpha.

If the search is too time consuming, you may try increasing this value.

Example 1 – Finding the Sample Size

A clinical trial is to be conducted over a two-year period to compare the hazard rate of a new treatment to that of the current treatment. The proportion surviving for two years using the current treatment is 0.3. The health community will be interested in the new treatment if the proportion surviving is increased to 0.45, a 50% increase. So that the sample size requirements for several survival proportions can be compared, it is also of interest to compute the sample size at response rates of 0.30, 0.35, 0.40, and 0.50. Assume the survival times are exponential.

Testing will be done at the 0.05 significance level and the power should be set to 0.90. A total of four tests are going to be performed on the data as they are obtained. The O'Brien-Fleming boundaries will be used.

Find the necessary sample sizes and test boundaries assuming equal sample sizes for each arm and two-sided hypothesis tests.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Group-Sequential Logrank Tests** procedure window by clicking on **Group-Sequential Tests**, then **Two Survival Curves (Logrank Test)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Alternative Hypothesis	Two-Sided
Survival Time Assumption.....	Exponential Survival
S1 (Proportion Surviving 1)	0.3
S2 (Proportion Surviving 2)	0.35, 0.40, 0.45, 0.50
Number of Looks	4
Spending Function	O'Brien-Fleming
Max Time	2
Times.....	Equally Spaced

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided LogRank Test (Assuming Exponential Survival)

Power	Total Sample Size (N)	Total Required Events	Alpha	Beta	Proportion Surv. (S1)	Proportion Surv. (S2)	Hazard Ratio
0.900003	3378	2280	0.050000	0.099997	0.3000	0.3500	0.8720
0.900267	884	575	0.050000	0.099733	0.3000	0.4000	0.7611
0.900626	407	254	0.050000	0.099374	0.3000	0.4500	0.6632
0.900018	234	140	0.050000	0.099982	0.3000	0.5000	0.5757

Report Definitions

Power is the probability of rejecting a false null hypothesis. Power should be close to one.

N is the number of items sampled from each group.

Events is the number of events that must occur in each group.

Alpha is the probability of rejecting a true null hypothesis in at least one of the sequential tests.

Beta is the probability of accepting a false null hypothesis at the conclusion of all tests.

S1 is the proportion surviving in group 1.

S2 is the proportion surviving in group 2.

HR is the hazard ratio. It is calculated using $\text{Log}(S2)/\text{Log}(S1)$.

Summary Statements

A total sample size of 3378 (split equally between the two groups), or 2280 events, achieves 90% power to detect a hazard rate of 0.8720 when the proportions surviving in each group are 0.3000 and 0.3500 at a significance level (alpha) of 0.050000 using a two-sided log rank test.

These results assume that 4 sequential tests are made using the O'Brien-Fleming spending function to determine the test boundaries and that the survival times are exponential.

This report shows the values of each of the parameters, one scenario per row. Note that 254 events are required when $S2 = 0.45$. Based on the expected survival proportions, this many events will occur if the overall sample size is 407.


Total Sample Size (N)

This is the estimated sample size that is needed to obtain the necessary number of events.

Total Required Events

This is the number of events (deaths, etc.) that are required to achieve the desired power levels.

Plots Section



This plot shows that an increase in sample size from under 1000 to well over 3000 is necessary when the detectable proportion surviving is reduced from 0.4 to 0.35.

Details Section

Details when Spending = O'Brien-Fleming, N = 407, d = 254, S1 = 0.3000, S2 = 0.4500

Look	Time	Lower Bndry	Upper Bndry	Nominal Alpha	Inc Alpha	Total Alpha	Inc Power	Total Power
1	0.5000	-4.33263	4.33263	0.000015	0.000015	0.000015	0.003516	0.003516
2	1.0000	-2.96311	2.96311	0.003045	0.003036	0.003051	0.255183	0.258699
3	1.5000	-2.35902	2.35902	0.018323	0.016248	0.019299	0.427665	0.686364
4	2.0000	-2.01406	2.01406	0.044003	0.030701	0.050000	0.214262	0.900626
Drift	3.27466							

This report shows information about the individual interim tests. One report is generated for each scenario.

Look

These are the sequence numbers of the interim tests.

Time

These are the time points at which the interim tests are conducted. Since the Max Time was set to 2 (for two years), these time values are in years. Hence, the first interim test is at half a year, the second at one year, and so on.

We could have set Max Time to 24 so that the time scale was in months.

Lower and Upper Boundary

These are the test boundaries. If the computed value of the test statistic z is between these values, the trial should continue. Otherwise, the trial can be stopped.

Nominal Alpha

This is the value of alpha for these boundaries if they were used for a single, standalone, test. Hence, this is the significance level that must be found for this look in a standard statistical package that does not adjust for multiple looks.

Inc Alpha

This is the amount of alpha that is *spent* by this interim test. It is close to, but not equal to, the value of alpha that would be achieved if only a single test was conducted. For example, if we lookup the third value, 2.35902, in normal probability tables, we find that this corresponds to a (two-sided) alpha of 0.018323. However, the entry is 0.016248. The difference is due to the correction that must be made for multiple tests.

Total Alpha

This is the total amount of alpha that is used up to and including the current test.

Inc Power

These are the amounts that are added to the total power at each interim test. They are often called the exit probabilities because they give the probability that significance is found and the trial is stopped, given the alternative hypothesis.

710-14 Group-Sequential Logrank Tests

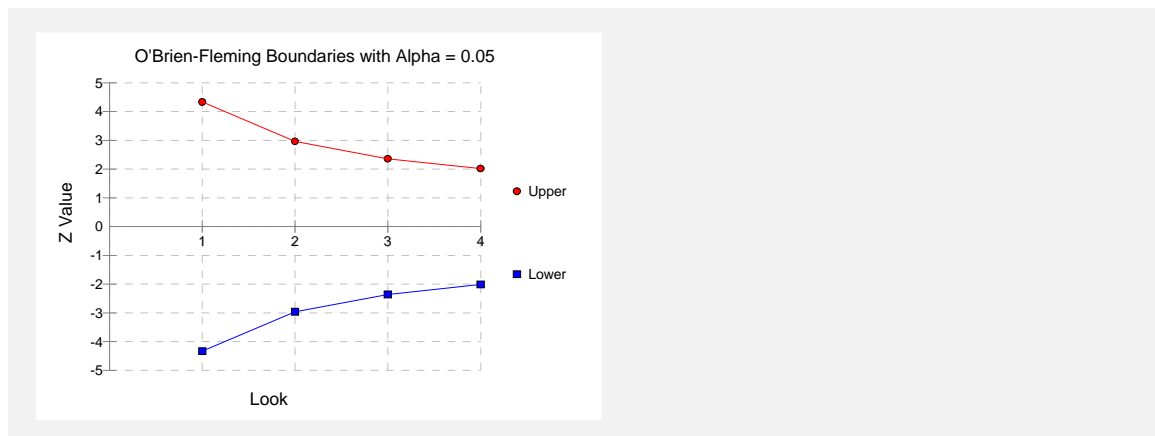
Total Power

These are the cumulative power values. They are also the cumulative exit probabilities. That is, they are the probability that the trial is stopped at or before the corresponding time.

Drift

This is the value of the Brownian motion drift parameter.

Boundary Plots



This plot shows the interim boundaries for each look. This plot shows very dramatically that the results must be extremely significant at early looks, but that they are near the single test boundary (1.96 and -1.96) at the last look.

Example 2 – Finding the Power

Continuing the scenario began in Example1, the researcher wishes to calculate the power of the design at sample sizes 50, 250, 450, 650, and 850. Testing will be done at the 0.01, 0.05, 0.10 significance levels and the overall power will be set to 0.10. Find the power of these sample sizes and test boundaries assuming equal sample sizes per arm and two-sided hypothesis tests.

Proceeding as in Example1, we decide to translate the mean and standard deviation into a percent of mean scale.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Group-Sequential Logrank Tests** procedure window by clicking on **Group-Sequential Tests**, then **Two Survival Curves (Logrank Test)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

Option

Value

Data Tab

Find (Solve For) **Power and Beta**

Power *Ignored since this is the Find setting*

Data Tab (continued)


Alpha **0.01, 0.05, 0.10**
 N (Total Sample Size) **50 to 850 by 200**
 Alternative Hypothesis **Two-Sided**
 Survival Time Assumption..... **Exponential Survival**
 S1 (Proportion Surviving 1) **0.3**
 S2 (Proportion Surviving 2) **0.45**
 Number of Looks **4**
 Spending Function **O'Brien-Fleming**
 Max Time **2**
 Times..... **Equally Spaced**

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results**Numeric Results for Two-Sided LogRank Test (Assuming Exponential Survival)**

Power	Total Sample Size (N)	Total Required Events	Alpha	Beta	Proportion Surv. (S1)	Proportion Surv. (S2)	Hazard Ratio
0.075632	50	31	0.010000	0.924368	0.3000	0.4500	0.6632
0.491011	250	156	0.010000	0.508989	0.3000	0.4500	0.6632
0.802921	450	281	0.010000	0.197079	0.3000	0.4500	0.6632
0.938923	650	406	0.010000	0.061077	0.3000	0.4500	0.6632
0.983780	850	531	0.010000	0.016220	0.3000	0.4500	0.6632
0.205032	50	31	0.050000	0.794968	0.3000	0.4500	0.6632
0.719225	250	156	0.050000	0.280775	0.3000	0.4500	0.6632
0.926894	450	281	0.050000	0.073106	0.3000	0.4500	0.6632
0.984044	650	406	0.050000	0.015956	0.3000	0.4500	0.6632
0.996907	850	531	0.050000	0.003093	0.3000	0.4500	0.6632
0.305105	50	31	0.100000	0.694895	0.3000	0.4500	0.6632
0.812103	250	156	0.100000	0.187897	0.3000	0.4500	0.6632
0.960490	450	281	0.100000	0.039510	0.3000	0.4500	0.6632
0.992807	650	406	0.100000	0.007193	0.3000	0.4500	0.6632
0.998800	850	531	0.100000	0.001200	0.3000	0.4500	0.6632



These data show the power for various sample sizes and alphas. It is interesting to note that once the sample size is greater than about 450, the value of alpha has comparatively little difference on the value of power.

Example 3 – Effect of Number of Looks

Continuing with examples one and two, it is interesting to determine the impact of the number of looks on power. **PASS** allows only one value for the Number of Looks parameter per run, so it will be necessary to run several analyses. To conduct this study, set alpha to 0.05, N to 407, and leave the other parameters as before. Run the analysis with Number of Looks equal to 1, 2, 3, 4, 6, 8, 10, and 20. Record the power for each run.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Group-Sequential Logrank Tests** procedure window by clicking on **Group-Sequential Tests**, then **Two Survival Curves (Logrank Test)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Total Sample Size)	407
Alternative Hypothesis	Two-Sided
Survival Time Assumption.....	Exponential Survival
S1 (Proportion Surviving 1)	0.3
S2 (Proportion Surviving 2)	0.45
Number of Looks	1 (Also run with 2, 3, 4, 6, 8, 10, and 20)
Spending Function	O'Brien-Fleming
Max Time.....	2
Times.....	Equally Spaced

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Logrank Test (Assuming Exponential Survival)

Power	N	Events	Alpha	Beta	S1	S2	Looks
0.905693	407	254	0.050000	0.094307	0.3000	0.4500	1
0.904719	407	254	0.050000	0.095281	0.3000	0.4500	2
0.902412	407	254	0.050000	0.097588	0.3000	0.4500	3
0.900626	407	254	0.050000	0.099374	0.3000	0.4500	4
0.898243	407	254	0.050000	0.101757	0.3000	0.4500	6
0.896763	407	254	0.050000	0.103237	0.3000	0.4500	8
0.895758	407	254	0.050001	0.104242	0.3000	0.4500	10
0.893398	407	254	0.050001	0.106602	0.3000	0.4500	20

This analysis shows how little the number of looks impacts the power of the design. The power of a study with no interim looks is 0.905693. When twenty interim looks are made, the power falls to 0.893398—a very small change.

Example 4 – Studying a Boundary Set

Continuing with the previous examples, suppose that you are presented with a set of boundaries and want to find the quality of the design (as measured by alpha and power). This is easy to do with *PASS*. Suppose that the analysis is to be run with five interim looks at equally spaced time points. The upper boundaries to be studied are 3.5, 3.5, 3.0, 2.5, 2.0. The lower boundaries are symmetric. The analysis would be run as follows.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the *PASS* Home window, load the **Group-Sequential Logrank Tests** procedure window by clicking on **Group-Sequential Tests**, then **Two Survival Curves (Logrank Test)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05 (will be calculated from boundaries)
N (Total Sample Size)	407
Alternative Hypothesis	Two-Sided
Survival Time Assumption.....	Exponential Survival
S1 (Proportion Surviving 1)	0.30
S2 (Proportion Surviving 2)	0.45
Number of Looks.....	5
Spending Function	User Supplied
Max Time	2
Times.....	Equally Spaced
Lower Boundaries	Symmetric
Upper Boundaries	3.5, 3.5, 3.0, 2.5, 2.0

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Logrank Test (Assuming Exponential Survival)								
Power	Total Sample Size (N)	Total Required Events	Alpha	Beta	Proportion Surv. (S1)	Proportion Surv. (S2)	Hazard Ratio	
0.900746	407	254	0.048157	0.099254	0.3000	0.4500	0.6632	
Details when Spending = User Supplied, N = 407, d=254, S1 = 0.3000, S2 = 0.4500								
Look	Time	Lower Bndry	Upper Bndry	Nominal Alpha	Inc Alpha	Total Alpha	Inc Power	Total Power
1	0.4000	-3.50000	3.50000	0.000465	0.000465	0.000465	0.020899	0.020899
2	0.8000	-3.50000	3.50000	0.000465	0.000408	0.000874	0.063132	0.084031
3	1.2000	-3.00000	3.00000	0.002700	0.002410	0.003284	0.243522	0.327553
4	1.6000	-2.50000	2.50000	0.012419	0.010331	0.013615	0.343072	0.670625
5	2.0000	-2.00000	2.00000	0.045500	0.034542	0.048157	0.230122	0.900746
Drift	3.27466							

The power for this design is about 0.90. This value depends on both the boundaries and the sample size. The alpha level is about 0.048. This value only depends on the boundaries.

Example 5 – Validation using O’Brien-Fleming Boundaries

Reboussin (1992) presents an example for binomial distributed data for a design with two-sided O’Brien-Fleming boundaries, looks = 3, alpha = 0.05, beta = 0.10, S1 = 0.30, S2 = 0.786 (which gives a hazard ratio of 0.20). They compute a drift of 3.261 and the number of events at 16.42. The upper boundaries are: 4.8769, 3.3569, 2.6803, 2.2898, 2.0310.

To test that *PASS* provides the same result, enter the following.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Group-Sequential Logrank Tests** procedure window by clicking on **Group-Sequential Tests**, then **Two Survival Curves (Logrank Test)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example5** from the Template tab on the procedure window.

Option	Value
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Alternative Hypothesis	Two-Sided
Survival Time Assumption.....	Exponential Survival

Data Tab (continued)

S1 (Proportion Surviving 1) **0.30**
 S2 (Proportion Surviving 2) **0.786**
 Number of Looks **5**
 Spending Function **O'Brien-Fleming**
 Max Time **1**
 Times **Equally Spaced**

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results**Numeric Results for Two-Sided Logrank Test (Assuming Exponential Survival)**

Power	Total Sample Size (N)	Total Required Events	Alpha	Beta	Proportion Surv. (S1)	Proportion Surv. (S2)	Hazard Ratio
0.905173	37	17	0.050000	0.094827	0.3000	0.7860	0.2000

Details when Spending = O'Brien-Fleming, N = 37, d = 17, S1 = 0.3000, S2 = 0.7860

Look	Time	Lower Bndry	Upper Bndry	Nominal Alpha	Inc Alpha	Total Alpha	Inc Power	Total Power
1	0.2000	-4.87688	4.87688	0.000001	0.000001	0.000001	0.000341	0.000341
2	0.4000	-3.35695	3.35695	0.000788	0.000787	0.000788	0.102760	0.103101
3	0.6000	-2.68026	2.68026	0.007357	0.006828	0.007616	0.352450	0.455551
4	0.8000	-2.28979	2.28979	0.022034	0.016807	0.024424	0.298953	0.754504
5	1.0000	-2.03100	2.03100	0.042255	0.025576	0.050000	0.150669	0.905173
Drift	3.30902							

The number of events, rounded to 17, matches the 16.42 reported in Reboussin (1992).

Chapter 715

Logrank Tests (Lakatos)

Introduction

This module computes the sample size and power of the logrank test for equality of survival distributions under very general assumptions. Accrual time, follow-up time, loss during follow up, noncompliance, and time-dependent hazard rates are parameters that can be set.

A clinical trial is often employed to test the equality of survival distributions for two treatment groups. For example, a researcher might wish to determine if Beta-Blocker A enhances the survival of newly diagnosed myocardial infarction patients over that of the standard Beta-Blocker B. The question being considered is whether the pattern of survival is different.

The two-sample t-test is not appropriate for two reasons. First, the data consist of the length of survival (time to failure), which is often highly skewed, so the usual normality assumption cannot be validated. Second, since the purpose of the treatment is to increase survival time, it is likely (and desirable) that some of the individuals in the study will survive longer than the planned duration of the study. The survival times of these individuals are then said to be *censored*. These times provide valuable information, but they are not the actual survival times. Hence, special methods have to be employed which use both regular and censored survival times.

The logrank test is one of the most popular tests for comparing two survival distributions. It is easy to apply and is usually more powerful than an analysis based simply on proportions. It compares survival across the whole spectrum of time, not just at one or two points. This module allows the sample size and power of the logrank test to be analyzed under very general conditions.

Power and sample size calculations for the logrank test have been studied by several authors. This *PASS* module uses the method of Lakatos (1988) because of its generality. This method is based on a Markov model that yields the asymptotic mean and variance of the logrank statistic under very general conditions.

Four Procedures Documented Here

There are four closely-related procedures that are documented in this chapter. These procedures are identical except for the parameterization of the effect size. The parameterization can be in terms of hazard rates, median survival time, proportion surviving, and mortality (proportion dying). Each of these options is listed separately.

The Markov process methodology divides the total study time into K equal-length intervals. The value of K is large enough so that the distribution within an interval can be assumed to follow the

715-2 Logrank Tests (Lakatos)

exponential distribution. The next section presents pertinent results for the exponential distribution.

Exponential Distribution

The density function of the exponential is defined as

$$f(t) = he^{-ht}$$

The probability of surviving the first t years is

$$S(t) = e^{-ht}$$

The mortality (probability of dying during the first t years) is

$$M(t) = 1 - e^{-ht}$$

For an exponential distribution, the mean survival is $1/h$ and the median is $\ln(2)/h$.

Notice that it is easy to translate between the hazard rate, the proportion surviving, the mortality, and the median survival time. The choice of which parameterization is used is arbitrary and is selected according to the convenience of the user.

Hazard Rate Parameterization

In this case, the hazard rates for the control and treatment groups are specified directly.

Median Survival Time Parameterization

Here, the median survival time is specified. These are transformed to hazard rates using the relationship $h = \ln(2) / MST$.

Proportion Surviving Parameterization

In this case, the proportion surviving until a given time $T0$ is specified. These are transformed to hazard rates using the relationship $h = -\ln(S(T0)) / T0$. Note that when separate proportions surviving are given for each time period, $T0$ is taken to be the time period number.

Mortality Parameterization

Here, the mortality until a given time $T0$ is specified. These are transformed to hazard rates using the relationship $h = -\ln(1 - M(T0)) / T0$. Note that when separate mortalities are given for each time period, $T0$ is taken to be the time period number.

Comparison of Lakatos Procedures to the other *PASS* Logrank Procedures

The follow chart lists the capabilities and assumptions of each of the logrank procedures available in *PASS*.

Feature/Capability	Algorithm		
	Simple (Freedman)	Advanced (Lachin)	Markov Process (Lakatos)
Test Statistic	Logrank statistic	Mean hazard difference*	Logrank statistic
Hazard Ratio	Constant	Constant	Any pattern including time-dependent
Basic Time Distribution	Constant hazard ratio**	Constant hazard ratio (exponential)	Any distribution
Loss to Follow Up Parameters	Yes	Yes	Yes
Accrual Parameters	No	Yes	Yes
Drop In Parameters	No	No	Yes
Noncompliance Parameters	No	No	Yes
Duration Parameters	No	Yes	Yes
Input Hazard Ratios	No	No	Yes
Input Median Survival Times	No	No	Yes
Input Proportion Surviving	Yes	Yes	Yes
Input Mortality Rates	No	No	Yes

*Simulation shows power similar to logrank statistic

**Not necessarily exponential

Comparison of Results

It is informative to calculate sample sizes for various scenarios using several of the methods. The scenario used to compare the various methods was $S1 = 0.5$, $S2 = 0.7$, $T0 = 4$, Loss to Follow Up = 0.05, Accrual Time = 2, Total Time = 4, and $N = 200$. Note that the Freedman method in *PASS* does not allow the input of $T0$, Accrual Time, or Total Time, so it is much less comparable. The Lachin/Foulkes and Lakatos values are very similar.

Computation Method	S1	S2	T0	Loss to Follow Up	Accrual Time	Total Time	N	Power
PASS (Freedman)	0.5	0.7	?	0.05	0	?	200	0.7979
PASS (Lachin/Foulkes)	0.5	0.7	4	0.05	2	4	200	0.7219
PASS (Lakatos)	0.5	0.7	4	0.05	2	4	200	0.7144

Technical Details

The logrank statistic L is defined as

$$L = \frac{\sum_{i=1}^d \left(X_i - \frac{n_{1i}}{n_{1i} + n_{2i}} \right)}{\left[\sum_{i=1}^d \frac{n_{1i}n_{2i}}{(n_{1i} + n_{2i})^2} \right]^{\frac{1}{2}}}$$

where X_i is an indicator for the control group, n_{1i} is the number at risk in the experimental group just before the i^{th} event (death), and n_{2i} is the number at risk in the control group just before the i^{th} event (death).

Following Freedman (1982) and Lakatos (1988), the trial is partitioned into K equal intervals. The distribution of L is asymptotically normal with mean E and variance V given by

$$E = \frac{\sum_{k=1}^K \sum_{i=1}^{d_k} \left(\frac{\phi_{ki} \theta_{ki}}{1 + \phi_{ki} \theta_{ki}} - \frac{\phi_{ki}}{1 + \phi_{ki}} \right)}{\left[\sum_{k=1}^K \sum_{i=1}^{d_k} \frac{\phi_{ki}}{(1 + \phi_{ki})^2} \right]^{\frac{1}{2}}}$$

$$V = \frac{\sum_{k=1}^K \sum_{i=1}^{d_k} \frac{\phi_{ki} \theta_{ki}}{(1 + \phi_{ki} \theta_{ki})^2}}{\sum_{k=1}^K \sum_{i=1}^{d_k} \frac{\phi_{ki}}{(1 + \phi_{ki})^2}}$$

where

$$\phi_{ki} = \frac{n_{1ki}}{n_{2ki}}, \quad \theta_{ki} = \frac{h_{1ki}}{h_{2ki}}$$

and h_{1ki} and h_{2ki} are the hazards of dying in the treatment and control groups respectively, just before the i^{th} death in the k^{th} interval. d_k is the number of deaths in the k^{th} interval.

Next, assume that the intervals are short enough so that the parameters are constant within an interval. That is, so that

$$\phi_{ki} = \phi_k, \quad \theta_{ki} = \theta_k, \quad h_{1ki} = h_{1k}, \quad h_{2ki} = h_{2k}$$

The values of E and V then reduce to

$$\begin{aligned}
 E &= \sqrt{d} \frac{\sum_{k=1}^K \left[\left(\frac{d_k}{d} \right) \left(\frac{\phi_k \theta_k}{1 + \phi_k \theta_k} - \frac{\phi_k}{1 + \phi_k} \right) \right]}{\sqrt{\sum_{k=1}^K \left[\left(\frac{d_k}{d} \right) \left(\frac{\phi_k}{(1 + \phi_k)^2} \right) \right]}} \\
 &= \sqrt{d} \frac{\sum_{k=1}^K \left[\rho_k \left(\frac{\phi_k \theta_k}{1 + \phi_k \theta_k} - \frac{\phi_k}{1 + \phi_k} \right) \right]}{\sqrt{\sum_{k=1}^K \left[\rho_k \left(\frac{\phi_k}{(1 + \phi_k)^2} \right) \right]}} \\
 V &= \frac{\sum_{k=1}^K \left[\left(\frac{d_k}{d} \right) \frac{\phi_k \theta_k}{(1 + \phi_k \theta_k)^2} \right]}{\sum_{k=1}^K \left[\left(\frac{d_k}{d} \right) \frac{d_k \phi_k}{(1 + \phi_k)^2} \right]} \\
 &= \frac{\sum_{k=1}^K \left[\rho_k \frac{\phi_k \theta_k}{(1 + \phi_k \theta_k)^2} \right]}{\sum_{k=1}^K \left[\rho_k \frac{d_k \phi_k}{(1 + \phi_k)^2} \right]}
 \end{aligned}$$

where

$$d = \sum_{k=1}^K d_k$$

and ρ_k is the proportion of the events (deaths) that occur in interval k .

The intervals mentioned above are constructed to correspond to a non-stationary Markov process, one for each group. This Markov process is defined as follows

$$S_{1,k} = T_{1,k,k-1} S_{1,k-1}$$

where $S_{1,k}$ is a vector giving the occupancy probabilities for each of the four possible states of the process: lost, dead, active complier, or active non-complier and $T_{1,k,k-1}$ is the transition matrix constructed so that each element gives the probability of transferring from state $j1$ to state $j2$ in the treatment group. A similar formulation is defined for the control group.

At each iteration

$$S_{1,k} = \begin{bmatrix} S_{1,k,1} \\ S_{1,k,2} \\ S_{1,k,3} \\ S_{1,k,4} \end{bmatrix}, \quad S_{2,k} = \begin{bmatrix} S_{2,k,1} \\ S_{2,k,2} \\ S_{2,k,3} \\ S_{2,k,4} \end{bmatrix}$$

715-6 Logrank Tests (Lakatos)

At the beginning of the trial

$$S_{1,0} = \begin{bmatrix} 0 \\ 0 \\ q_1 \\ 0 \end{bmatrix}, S_{2,0} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 - q_1 \end{bmatrix}$$

where q_1 is the control proportion of the total sample.

The transition matrices may be different for each group, but this does not need to be so. Its elements are as follows (the first row and column contains labels which are not part of the actual matrix).

$$T_{1,k,k-1} = \begin{bmatrix} \text{States} & \text{Lost} & \text{Event} & \text{Complier} & \text{Non-complier} \\ \text{Lost} & 1 & 0 & p_{loss,k} & p_{loss,k} \\ \text{Event} & 0 & 1 & p_{event1,k} & p_{event2,k} \\ \text{Complier} & 0 & 0 & 1 - sum_c & p_{drop-in,k} \\ \text{Non-complier} & 0 & 0 & p_{noncomp,k} & 1 - sum_n \end{bmatrix}$$

where sum_c and sum_n represent the sum of the other elements of their columns.

These values represent parameters of the population such as event rates, loss to follow-up rates, and recruitment rates.

The parameters ϕ_k , θ_k , and d_k are estimated from the occupancy probabilities as follows

Events (deaths)

$$d_{1,k} = s_{1,k,2} - s_{1,k-1,2}$$

$$d_{2,k} = s_{2,k,2} - s_{2,k-1,2}$$

Censored

$$c_{1,i} = s_{1,k,1} - s_{1,k-1,1}$$

$$c_{2,i} = s_{2,k,1} - s_{2,k-1,1}$$

At Risk

$$a_{1,k} = (s_{1,k-1,3} + s_{1,k-1,4})$$

$$a_{2,k} = (s_{2,k-1,3} + s_{2,k-1,4})$$

Hazard

$$h_{1,k} = d_{1,k} / a_{1,k}$$

$$h_{2,k} = d_{2,k} / a_{2,k}$$

Finally, the interval parameters are given by

$$\phi_k = \frac{s_{1,k-1,3} + s_{1,k-1,4}}{s_{2,k-1,3} + s_{2,k-1,4}}$$

$$\theta_k = \frac{h_{1,k,3}}{h_{2,k,3}}$$

$$d_k = d_{1,k} + d_{2,k}$$

Power Calculation

1. Find z_α such that $1 - \Phi(z_\alpha) = \alpha$, where $\Phi(x)$ is the area under the standardized normal curve to the left of x .
2. Calculate E_0 and V_0 assuming the two transition matrices are the same (H0). Also, calculate E_1 and V_1 assuming the two transition matrices are different (H1)
3. Calculate: $X_\alpha = E_0 + z_\alpha V_0$
4. Calculate: $z_\beta = \frac{X_\alpha - E_1}{V_1}$
5. Calculate beta and power: $\beta = \Phi(z_\beta)$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with. This chapter covers four procedures, each of which has different effect size options. However, many of the options are common to all four procedures. These common options will be displayed first, followed by the various effect size options.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *Power and Beta*, *N*, and *{effect size}*. Note that the *effect size* corresponds to the parameterization that is chosen.

Select *N* when you want to calculate the sample size needed to achieve a given power and alpha level.

715-8 Logrank Tests (Lakatos)

Select *Power and Beta* when you want to calculate the power.

Error Rates

Power (1 – Beta)

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal survival curves when in fact the curves are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis of equal survival curves when in fact the curves are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for a two-sided test and 0.025 has been used for a one-sided test. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Total Sample Size)

This is the combined sample size of both groups. This amount is divided between the two groups using the value of the Proportion in Control Group. You can enter a single value or a list of Sample Sizes such as *50 100 150* or *50 to 450 by 100*.

Proportion in Control Group

Enter one or more values for the proportion of N in the control group. If this value is labeled p_1 , the sample size of the control group is Np_1 and the sample size of treatment group is $N - Np_1$. Note that the value of Np_1 is rounded to the nearest integer.

The value of 0.5 results in equal sample sizes per group.

Proportion Lost or Switching Groups

Controls (or Treatment) Lost

This is the proportion of subjects in the control (treatment) group that disappear from the study during a single time period (month, year, etc.). Multiple entries, such as *0.01 0.03 0.05*, are allowed.

When you want to specify different proportions for different time periods, you would enter those rates into a column of the spreadsheet, one row per time period. You specify the column of the

spreadsheet by beginning your entry with an equals sign. For example, if you have entered the proportions in column 5, you would enter =C5 here.

Controls Switching to Treatments

This is the proportion of subjects in the control group that change to a treatment regime similar in efficacy to the treatment group during a single time period (month, year, etc.). This is sometimes referred to as *drop in*. Multiple entries, such as 0.01 0.03 0.05, are allowed.

When you want to specify different proportions for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column of the spreadsheet by beginning your entry with an equals sign. For example, if you have entered the proportions in column 1, you would enter =C1 here.

Treatments Switching to Controls

This is the proportion of subjects in the treatment group that change to a treatment regime similar in efficacy to the control group during a single time period (month, year, etc.). This is sometimes referred to as *noncompliance*. Multiple entries, such as 0.01 0.03 0.05, are allowed.

When you want to specify different proportions for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column of the spreadsheet by beginning your entry with an equals sign. For example, if you have entered the proportions in column 2, you would enter =C2 here.

Effect Size (Hazard Rate)

h1 (Hazard Rate – Control Group)

Specify one or more hazard rates (instantaneous failure rate) for the control group. For an exponential distribution, the hazard rate is the inverse of the mean survival time. An estimate of the hazard rate may be obtained from the median survival time or from the proportion surviving to a certain time point. This calculation is automated by pressing the *Parameter Conversion* button.

Hazard rates must be greater than zero. Constant hazard rates are specified by entering them directly. Variable hazard rates are specified as columns of the spreadsheet. When you want to specify different hazard rates for different time periods, you would enter those rates into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the hazard rates in column 2, you would enter =2 here.

The following examples assume an exponential survival distribution.

Median Survival Time	Hazard Rate
0.5	1.386
1.0	0.693
2.0	0.347
3.0	0.231
4.0	0.173
5.0	0.139

Treatment Group Parameter

Specify which of the parameters below will be used to specify the treatment group hazard rate.

h2 (Hazard Rate – Treatment Group)

Specify one or more hazard rates (instantaneous failure rate) for the treatment group. An estimate of the hazard rate may be obtained from the median survival time or from the proportion surviving to a certain time point. This calculation is automated by pressing the *Parameter Conversion* button.

Hazard rates must be greater than zero. Constant hazard rates are specified by entering them directly. Variable hazard rates are specified as columns of the spreadsheet. When you want to specify different hazard rates for different time periods, you would enter those rates into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the hazard rates in column 3, you would enter =3 here.

HR (Hazard Ratio = h_2/h_1)

Specify one or more values for the hazard ratio, $HR = h_2/h_1$. Hazard ratios must be greater than zero. The null hypothesis is that the hazard ratio is 1.0. Typical values of the hazard ratio are from 0.25 to 4.0.

Constant hazard ratios are specified by entering them directly. Variable hazard ratios are specified as columns of the spreadsheet.

An estimate of the hazard ratio may be obtained from the median survival times, from the hazard rates, or from the proportion surviving past a certain time point by pressing the *Parameter Conversion* button.

When you want to specify different hazard ratios for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning your entry with an equals sign.

For example, if you have entered the hazard ratios in column 3, you would enter =3 here.

Effect Size (Median Survival Time)

T1 (Median Survival Time – Control)

Specify one or more median survival times for the control group. These values must be greater than zero.

Constant median survival times are specified by entering them directly. Variable median survival times are specified as columns of the spreadsheet. When you want to specify different median survival times for different time periods, you would enter those times into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the median survival times in column 2, you would enter =2 here.

The following examples assume an exponential survival distribution.

Median Survival Time	Hazard Rate
0.5	1.386
1.0	0.693
2.0	0.347
3.0	0.231
4.0	0.173
5.0	0.139

Treatment Group Parameter

Specify which of the parameters below will be used to specify the treatment group median survival time.

T2 (Median Survival Time – Treatment)

Specify one or more median survival times for the treatment group. These values must be greater than zero.

Constant median survival times are specified by entering them directly. Variable median survival times are specified as columns of the spreadsheet. When you want to specify different median survival times for different time periods, you would enter those times into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the median survival times in column 1, you would enter =1 here.

HR (Hazard Ratio = T1/T2)

Specify one or more values for the hazard ratio, $HR = T1/T2 = h2/h1$. Hazard ratios must be greater than zero. The null hypothesis is that the hazard ratio is 1.0. Typical values of the hazard ratio are from 0.25 to 4.0.

Constant hazard ratios are specified by entering them directly. Variable hazard ratios are specified as columns of the spreadsheet.

An estimate of the hazard ratio may be obtained from the median survival times, from the hazard rates, or from the proportion surviving past a certain time point by pressing the *Parameter Conversion* button.

When you want to specify different hazard ratios for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning your entry with an equals sign.

For example, if you have entered the hazard ratios in column 3, you would enter =3 here.

Effect Size (Proportion Surviving)

S1 (Proportion Surviving – Control)

Specify one or more proportions surviving for the control group. These values must be between zero and one. Constant proportions surviving are specified by entering them directly. The values represent the proportions surviving until time T0.

715-12 Logrank Tests (Lakatos)

Variable proportions surviving are specified as columns of the spreadsheet. When you want to specify different median survival times for different time periods, you would enter those times into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the median survival times in column 2, you would enter =2 here.

Treatment Group Parameter

Specify which of the parameters below will be used to specify the proportion surviving in the treatment group.

S2 (Proportion Surviving – Treatment)

Specify one or more proportions surviving for the treatment group. These values must be between zero and one. Constant proportions surviving are specified by entering them directly. The values represent the proportions surviving until time T0.

Variable proportions surviving are specified as columns of the spreadsheet. When you want to specify different proportions surviving for different time periods, you would enter those times into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the proportions surviving in column 3, you would enter =3 here.

HR (Hazard Ratio)

Specify one or more values for the hazard ratio, $HR = h_2/h_1$. Hazard ratios must be greater than zero. The null hypothesis is that the hazard ratio is 1.0. Typical values of the hazard ratio are from 0.25 to 4.0.

Constant hazard ratios are specified by entering them directly. Variable hazard ratios are specified as columns of the spreadsheet.

An estimate of the hazard ratio may be obtained from the median survival times, from the hazard rates, or from the proportion surviving past a certain time point by pressing the *Parameter Conversion* button.

When you want to specify different hazard ratios for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning your entry with an equals sign.

For example, if you have entered the hazard ratios in column 3, you would enter =3 here.

T0 (Survival Time)

This is the time corresponding to the proportions surviving. It must be a value greater than zero.

When you say 0.40 survive, you must indicate the number of time periods (years) to which they survive. That is, you must say 40% survive over five years. For example, a value of 3 here and proportion surviving of 0.4 means that 40% survive over three years.

This value is only used when S1 and S2 are entered as numbers. It is not used when a proportion surviving is entered as a column because, in that case, the time period is different for each row.

Effect Size (Mortality)

M1 (Mortality – Control)

Specify one or more mortality values for the control group. These values must be between zero and one. Constant mortalities are specified by entering them directly. The values represent the proportions dying until time T0.

Variable mortalities are specified as columns of the spreadsheet. When you want to specify different mortalities for different time periods, you would enter those times into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the mortalities in column 2, you would enter =2 here.

Treatment Group Parameter

Specify which of the parameters below will be used to specify the proportion dying in the treatment group.

M2 (Mortality – Treatment)

Specify one or more mortalities (proportions dying) for the treatment group. These values must be between zero and one. Constant mortalities are specified by entering them directly. The values represent the mortalities until time T0.

Variable mortalities are specified as columns of the spreadsheet. When you want to specify different mortalities for different time periods, you would enter those times into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning the entry with an equals sign. For example, if you have entered the mortalities in column 3, you would enter =3 here.

MR (Mortality Ratio = M2/M1)

Specify one or more values for the mortality ratio, $MR = M2/M1$. Mortality ratios must be greater than zero. The null hypothesis is that the mortality ratio is 1.0. Typical values of the mortality ratio are from 0.25 to 4.0.

Constant mortality ratios are specified by entering them directly. Variable mortality ratios are specified as columns of the spreadsheet.

An estimate of the mortality ratio may be obtained from median survival times, from hazard rates, or from the proportions surviving past a certain time point by pressing the *Parameter Conversion* button.

When you want to specify different mortality ratios for different time periods, you would enter those values into a column of the spreadsheet, one row per time period. You specify the column (or columns) by beginning your entry with an equals sign.

For example, if you have entered the hazard ratios in column 3, you would enter =3 here.

T0 (Survival Time)

This is the time corresponding to the mortality. It must be a value greater than zero.

When you say 0.40 die, you must indicate the number of time periods (years) on which this is based. For example, a value of 3 here and mortality of 0.4 means that 40% die over the first three years.

This value is only used when M1 and M2 are entered as numbers. It is not used when mortality is entered as a column because, in that case, the time period is different for each row.

Duration

Accrual Time (Integers Only)

Enter one or more values for the number of time periods (months, years, etc.) during which subjects are entered into the study. The total duration of the study is equal to the Accrual Time plus the Follow-Up Time. These values must be integers.

Accrual times can range from 0 to the Total Time. That is, the accrual time must be less than or equal to the Total Time. Otherwise, the scenario is skipped.

Enter 0 when all subjects begin the study together.

Accrual Pattern

This contains the pattern of accrual (patient entry). Two types of entries are possible:

- **Uniform**

If you want to specify a uniform accrual rate for all time periods, enter *Equal* here.

- **Non-Uniform**

When you want to specify accrual patterns with different accrual proportions per time period, you would enter the pattern into a column of the spreadsheet, one row per time period and specify the column, or columns, here, beginning your entry with an equals sign. For example, if you have entered accrual patterns in columns 4 and 5, you would enter =C4 C5.

Note that these values are standardized to sum to one. Thus, the accrual pattern 0.25 0.50 0.25 would result in the same accrual pattern as 1 2 1 or 25 50 25.

Total Time (Integers Only)

Enter one or more values for the number of time periods (months, years, etc.) in the study. The follow-up time is equal to the Total Time minus the Accrual Time. These values must be integers.

Test

Alternative

Specify whether the statistical test is two-sided or one-sided.

- **Two-Sided**

This option tests whether the two hazards rates, median survival times, survival proportions, or mortalities are different ($H_1: h_1 \neq h_2$). This is the option that is usually selected.

- **One-Sided**

When this option is used and the value of h_1 is less than h_2 , rejecting the null hypothesis results in the conclusion that the control hazard rate (h_1) is less than the treatment hazard rate (h_2). When h_1 is greater than h_2 , rejecting the null hypothesis results in the conclusion that the control hazard rate (h_1) is greater than the treatment hazard rate (h_2).

When you use a one-sided test, you should divide your alpha level by two.

Reports Tab

The Reports tab contains additional settings for this procedure.

Report Column Width

Report Column Width

This option sets the width of the each column of the numeric report.

The numeric report for this option necessarily contains many columns, so the maximum number of decimal places that can be displayed is four. If you try to increase that number, the numbers may run together. You can increase the width of each column using this option.

The recommended report column width for scenarios without large numbers of decimal places or extremely large sample sizes is 0.49.

Options Tab

The Options tab contains additional settings for this procedure.

Options

Number of Intervals within a Time Period

The algorithm requires that each time period be partitioned into a number of equal-width intervals. Each of these subintervals is assumed to follow an exponential distribution. This option controls the number of subintervals. All parameters such as hazard rates, loss to follow-up rates, and noncompliance rates are assumed to be constant within a subinterval.

Lakatos (1988) gives little input as to how the number of subintervals should be chosen. In a private communication, he indicated that 100 ought to be adequate. This seems to work when the hazard is less than 1.0.

As the hazard rate increases above 1.0, this number must increase. A value of 2000 should be sufficient as long as the hazard rates (h_1 and h_2) are less than 10. When the hazard rates are greater than 10, you may want to increase this value to 5000 or even 10000.

Example 1 – Finding the Power using Proportion Surviving

A researcher is planning a clinical trial using a parallel, two-group, equal sample allocation design to compare the survivability of a new treatment with that of the current treatment. The proportion surviving one-year after the current treatment is 0.50. The new treatment will be adopted if the proportion surviving after one year can be shown to be higher than the current treatment. The researcher wishes to determine the power of the logrank test to detect a difference in survival when the true proportion surviving in the new treatment group at one year is 0.70. To obtain a better understanding of the relationship between power and survivability, the researcher also wants to see the results when the proportion surviving is 0.65 and 0.75.

The trial will include a recruitment period of one-year after which participants will be followed for an additional two-years. It is assumed that patients will enter the study uniformly over the accrual period. The researcher estimates a loss-to-follow-up rate of 5% per year in both the control and the experimental groups. Past experience has lead to estimates of noncompliance and drop in of 4% and 3%, respectively.

The researcher decides to investigate various sample sizes between 50 and 250 at a significance level of 0.05.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lakatos) [Proportion Surviving]** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lakatos) using Proportion Surviving**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Total Sample Size)	50 to 250 by 50
Proportion in Control Group	0.5
Proportion of Controls Lost	0.05
Proportion of Treatment Lost	0.05
Proportion of Controls Switch	0.03
Proportion of Treatment Switch	0.04
S1	0.50
Treatment Group Parameter	S2
S2	0.65 0.70 0.75
T0	1
Accrual Time	1
Accrual Pattern.....	Equal
Total Time	3
Test	Two-Sided

Reports TabShow Detail Numeric Reports **Checked****Annotated Output**

Click the Run button to perform the calculations and generate the following output.

Numeric Results**Numeric Results in Terms of Sample Size when the Test is Two-Sided and T0 is 1**

Power	N1	N2	N	Haz Ratio (HR)	Ctrl Prop Surv (S1)	Trt Prop Surv (S2)	Acc-rual Pat'n	Acc-rual Time/Total Time	Ctrl Loss	Trt Loss	Ctrl to Trt	Trt to Ctrl	Alpha	Beta
.2608	25	25	50	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.7392
.4615	50	50	100	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.5385
.6262	75	75	150	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.3738
.7500	100	100	200	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.2500
.8378	125	125	250	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.1622
.4320	25	25	50	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.5680
.7162	50	50	100	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.2838
.8732	75	75	150	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.1268
.9477	100	100	200	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0523
.9796	125	125	250	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0204
.6293	25	25	50	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.3707
.9010	50	50	100	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0990
.9784	75	75	150	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0216
.9959	100	100	200	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0041
.9993	125	125	250	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0007

References

Lakatos, Edward. 1988. 'Sample Sizes Based on the Log-Rank Statistic in Complex Clinical Trials', Biometrics, Volume 44, March, pages 229-241.

Lakatos, Edward. 2002. 'Designing Complex Group Sequential Survival Trials', Statistics in Medicine, Volume 21, pages 1969-1989.

Report Definitions

Power is the probability of rejecting a false null hypothesis. Power should be close to one.

N1|N2|N are the sample sizes of the control group, treatment group, and both groups, respectively.

E1|E2|E are the number of events in the control group, treatment group, and both groups, respectively.

Hazard Ratio (HR) is the treatment group's hazard rate divided by the control group's hazard rate.

Proportion Surviving is the proportion surviving past time T0.

Accrual Time is the number of time periods (years or months) during which accrual takes place.

Total Time is the total number of time periods in the study. Follow-up time = (Total Time) - (Accrual Time).

Ctrl Loss is the proportion of the control group that is lost (drop out) during a single time period (year or month).

Trt Loss is the proportion of the treatment group that is lost (drop out) during a single time period (year or month).

Ctrl to Trt (drop in) is the proportion of the control group that switch to a group with a hazard rate equal to the treatment group.

Trt to Ctrl (noncompliance) is the proportion of the treatment group that switch to a group with a hazard rate equal to the control group.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

Beta is the probability of accepting a false null hypothesis. It should be small.

This report shows the values of each of the parameters, one scenario per row. In addition to the parameters that were set on the template, the hazard ratio is displayed.

715-18 Logrank Tests (Lakatos)

Next, a report displaying the number of required events rather than the sample size is displayed.

Numeric Results in Terms of Events when the Test is Two-Sided and T0 is 1															
	Ctrl Evts	Trt Evts	Total Evts	Haz Ratio	Ctrl Prop Surv	Trt Prop Surv	Acc- rual	Acc- rual Time/Total	Ctrl Loss	Trt Loss	Ctrl to Trt	Trt to Ctrl	Alpha	Beta	
Power	E1	E2	E	(HR)	(S1)	(S2)	Pat'n	Time							
.2608	19.5	15.8	35.2	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.7392	
.4615	38.9	31.6	70.5	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.5385	
.6262	58.4	47.4	105.7	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.3738	
.7500	77.8	63.1	140.9	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.2500	
.8378	97.3	78.9	176.2	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.1622	
.4320	19.4	14.2	33.6	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.5680	
.7162	38.8	28.4	67.2	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.2838	
.8732	58.2	42.7	100.9	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.1268	
.9477	77.6	56.9	134.5	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0523	
.9796	97.0	71.1	168.1	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0204	
.6293	19.4	12.5	31.8	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.3707	
.9010	38.7	25.0	63.7	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0990	
.9784	58.1	37.5	95.5	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0216	
.9959	77.4	49.9	127.4	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0041	
.9993	96.8	62.4	159.2	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0007	

Most of this report is identical to the last report, except that the sample sizes are replaced by the number of required events.

Next, reports displaying the individual settings year-by-year for each scenario are displayed.

Detailed Input when Power=.2608 N1=25 N2=25 N=50 Alpha=.0500 Accrual/Total Time=1 / 3										
Time Period	Control Prop Surviving (.5000)	Treatment Prop Surviving (.6500)	Hazard Ratio (HR) (.6215)	Percent Accrual (Equal)	Percent Admin. Censored (Calc.)	Control Loss (.0500)	Treatment Loss (.0500)	Control to Treatment (.0300)	Treatment to Control (.0400)	
1	.5000	.6500	.7692	100.00	.00	.0500	.0500	.0300	.0400	
2	.5000	.6500	.7692	.00	.00	.0500	.0500	.0300	.0400	
3	.5000	.6500	.7692	.00	100.00	.0500	.0500	.0300	.0400	

This report shows the individual settings for each time period (year). It becomes very useful when you want to document a study in which these parameters vary from year to year.

One subtle point should be mentioned here. Note that the proportion surviving is the same for each of the three years. Obviously, if deaths occur, the proportion surviving must decrease as time passes. The reason that these stay the same is that this proportion surviving was entered as a single value. Hence, it is converted to a hazard rate using T0 (which in this example is 1.0), not the row number. Since the power is based on the hazard rates, the proportions surviving can be identical.


Next, summary statements are displayed.

Summary Statements

A two-sided logrank test with an overall sample size of 50 subjects (25 in the control group and 25 in the treatment group) achieves 26.1% power at a 0.050 significance level to detect a hazard ratio of 0.6215 when the proportion surviving in the control group is 0.5000. The study lasts for 3 time periods of which subject accrual (entry) occurs in the first time period. The proportion dropping out of the control group is 0.0500. The proportion dropping out of the treatment group is 0.0500. The proportion switching from the control group to another group with a hazard ratio equal to that of the treatment group is 0.0300. The proportion switching from the treatment group to another group with a hazard equal to that of the control group is 0.0400.

Finally, a scatter plot of the results is displayed.

Plots Section



This plot shows the relationship between sample size and power for the three values of S2. Note that for 90% power, a total sample size of about 160 is required. The exact number will be found in Example 2.

Example 2 – Finding the Sample Size

Continuing with the previous example, the researcher wants to investigate the sample sizes necessary to achieve 80% and 90% power. All other parameters will remain the same as in Example 1.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lakatos) [Proportion Surviving]** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lakatos) using Proportion Surviving**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.80 0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Control Group	0.5
Proportion of Controls Lost	0.05
Proportion of Treatment Lost	0.05

715-20 Logrank Tests (Lakatos)

Data Tab (continued)

Proportion of Controls Switch0.03
 Proportion of Treatment Switch0.04
 S10.50
 Treatment Group ParameterS2
 S20.65 0.70 0.75
 T01
 Accrual Time1
 Accrual Pattern.....Equal
 Total Time3
 TestTwo-Sided

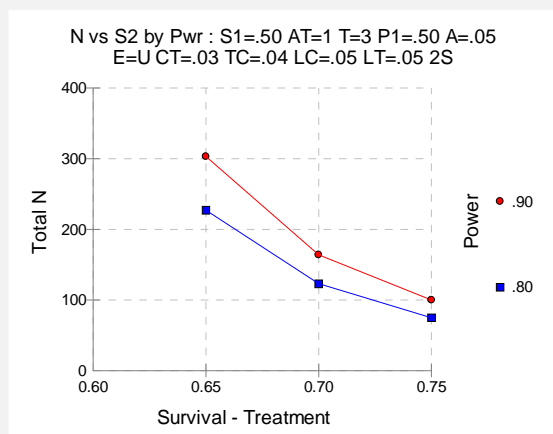
Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results in Terms of Sample Size when the Test is Two-Sided and T0 is 1

Power	N1	N2	N	Haz Ratio (HR)	Ctrl Prop Surv (S1)	Trt Prop Surv (S2)	Acc-rual Pat'n	Acc-rual Time/Total Time	Ctrl Loss	Trt Loss	Ctrl to Trt	Trt to Ctrl	Alpha	Beta
.9002	151	152	303	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0998
.8014	113	114	227	.6215	.5000	.6500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.1986
.9004	82	82	164	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0996
.8019	61	62	123	.5146	.5000	.7000	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.1981
.9010	50	50	100	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.0990
.8022	37	38	75	.4150	.5000	.7500	Equal	1 / 3	.0500	.0500	.0300	.0400	.0500	.1978



The total sample size need to achieve 90% power when the proportion surviving with the new treatment is 0.70, is 164. It is apparent that as the proportion surviving (effect size) increases, the sample size decreases.

Example 3 – Validation using Lakatos

Lakatos (1988) pages 231-234 presents an example that will be used to validate this procedure. In this example, a two-year trial is investigated. All subjects begin the trial together, so there is no accrual period. The hazard rates are 1.0 and 0.5 for the control and treatment groups, respectively. The yearly loss to follow-up is 3% per year in both groups. Noncompliance and drop-in rates are assumed to be 4% and 5%, respectively. The power is set to 90%. A two-sided logrank test with alpha set to 0.05 is assumed. Equal allocation of the sample to both control and experiment groups is used. Lakatos obtains a sample size of 139.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lakatos) [Hazard Ratio]** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lakatos) using Hazard Rate**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find	N
Power	0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Control Group	0.5
Proportion of Controls Lost	0.03
Proportion of Treatment Lost	0.03
Proportion of Controls Switch	0.04
Proportion of Treatment Switch	0.05
h1	1.0
Treatment Group Parameter	h2
h2	0.5
Accrual Time	0
Accrual Pattern.....	Equal
Total Time	2
Test	Two-Sided

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results in Terms of Sample Size when the Test is Two-Sided and T0 is 1

Power	N1	N2	N	Haz Ratio (HR)	Ctrl Haz Rate (h1)	Trt Haz Rate (h2)	Acc- rual Pat'n	Acc- rual Time/ Total Time	Ctrl Loss	Trt Loss	Ctrl to Trt	Trt to Ctrl	Alpha	Beta
.9014	69	70	139	.5000	1.00	.5000	Equal	0 / 2	.0300	.0300	.0400	.0500	.0500	.0986

The total sample size of 139 matches the value published in Lakatos' article.

Example 4 – Inputting Time-Dependent Hazard Rates from a Spreadsheet

This example shows how time-dependent hazard rates and other parameters can be input directly from a spreadsheet.

A pre-trial study indicates that a newly developed treatment will cut the hazard rate in half, when compared to the current treatment. A 5-year trial is being designed to confirm the finding of the pre-trial study. The goal for this portion of the study design is to determine the sample size needed to detect a decrease in hazard rate with 90% power.

The pre-trial study showed that the hazard rate immediately following either treatment (during the first year) is high, drops considerably during the second year, and then gradually increases. Fifty percent of the study participants will be enrolled during the first year, followed by 25% each of the second and third years. The following table shows the time-dependent parameters for the 5-year trial, based on the pre-trial study.

PRETRIAL dataset

Year	H1	Ls1	Ls2	NCom	Acc
1	0.08	0.04	0.06	0.04	50
2	0.04	0.04	0.06	0.04	25
3	0.05	0.05	0.07	0.05	25
4	0.06	0.06	0.07	0.06	
5	0.07	0.07	0.08	0.07	

The column H1 refers to the anticipated hazard rates for each of the five years. Ls1 and Ls2 refer to the proportions lost to follow-up in the control group and the treatment group, respectively. The proportion that are noncompliant are also expected to increase after the second year according to the proportions shown. The final column specifies the accrual rate as outlined in the previous paragraph.

Following the 5-year trial, a two-sided logrank test with alpha equal to 0.05, will be used to determine the evidence of difference among the current and new treatments.

Setup

In order to run this example the **PRETRIAL.S0** data must be loaded into the spreadsheet. To open the spreadsheet window from the PASS Home Window, click on the **Tools** menu and select **Spreadsheet**. Once the spreadsheet is open, the **PRETRIAL.S0** data is loaded by clicking the **File** menu and selecting **Open**. The **PRETRIAL.S0** file is then selected from the **DATA** folder (the default location for this folder is *C:\...\[My] Documents\NCSS\PASS2008*). Then click **Open**.

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lakatos) [Hazard Ratio]** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lakatos) using Hazard Rate**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Total Sample Size)	<i>Ignored since this is the Find setting</i>
Proportion in Control Group	0.5
Proportion of Controls Lost	=Ls1
Proportion of Treatment Lost	=Ls2
Proportion of Controls Switch	0.02
Proportion of Treatment Switch	=NCom
h1	=H1
Treatment Group Parameter	HR (Hazard Ratio = h2/h1)
HR	0.5
Accrual Time	3
Accrual Pattern.....	=Acc
Total Time	5
Alternative Hypothesis	Two-Sided

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results in Terms of Sample Size when the Test is Two-Sided
Using Spreadsheet: C:\PASS2008\DATA\PRETRIAL.S0

Power	N1	N2	N	Haz Ratio (HR)	Ctrl Haz Rate (h1)	Trt Haz Rate (h2)	Acc-rual Pat'n	Acc-rual Time/Total Time	Ctrl Loss Ls1	Trt Loss Ls2	Ctrl to Trt (.0200)	Trt to Ctrl (NCom)	Alpha	Beta
.9001	418	419	837	.5000	H1	Calc.	Acc	3 / 5	Ls1	Ls2	.0200	NCom	.0500	.0999

Numeric Results in Terms of Events when the Test is Two-Sided
Using Spreadsheet: C:\PASS2008\DATA\PRETRIAL.S0

Power	Ctrl Evts (E1)	Trt Evts (E2)	Total Evts (E)	Haz Ratio (HR)	Ctrl Haz Rate (h1)	Trt Haz Rate (h2)	Acc-rual Pat'n	Acc-rual Time/Total Time	Ctrl Loss Ls1	Trt Loss Ls2	Ctrl to Trt (.0200)	Trt to Ctrl (NCom)	Alpha	Beta
.9001	74.4	41.2	115.6	.5000	H1	Calc.	Acc	3 / 5	Ls1	Ls2	.0200	NCom	.0500	.0999

Detailed Input when Power=.9001 N1=418 N2=419 N=837 Alpha=.0500 Accrual/Total Time=3 / 5
Using Spreadsheet: C:\PASS2008\DATA\PRETRIAL.S0

Time Period	Control Hazard Rate (H1)	Treatment Hazard Rate (Calc.)	Hazard Ratio (HR) (.5000)	Percent Accrual (Acc)	Percent Admin. Censored (Calc.)	Control Loss (Ls1)	Treatment Loss (Ls2)	Control to Treatment (.0200)	Treatment to Control (NCom)
1	.0800	.0400	.5000	50.00	.00	.0400	.0600	.0200	.0400
2	.0400	.0200	.5000	25.00	.00	.0400	.0600	.0200	.0400
3	.0500	.0250	.5000	25.00	25.00	.0500	.0700	.0200	.0500
4	.0600	.0300	.5000	.00	33.33	.0600	.0700	.0200	.0600
5	.0700	.0350	.5000	.00	100.00	.0700	.0800	.0200	.0700

Summary Statements

A two-sided logrank test with an overall sample size of 837 subjects (418 in the control group and 419 in the treatment group) achieves 90.0% power at a 0.050 significance level to detect a hazard ratio of 0.5000 when the control group hazard rate is given in column H1. The study lasts for 5 time periods of which subject accrual (entry) occurs in the first 3 time periods. The accrual pattern across time periods is given in column Acc. The proportion dropping out of the control group is given in column Ls1. The proportion dropping out of the treatment group is given in column Ls2. The proportion switching from the control group to another group with a hazard rate equal to the treatment group is 0.0200. The proportion switching from the treatment group to another group with a hazard rate equal to the control group is given in column NCom.

For the 5-year study, the total sample size needed to detect a change in hazard rate, if the true hazard ratio is 0.5, is 837 subjects.

Example 5 – Finding the Power using Median Survival Time

A researcher is planning a clinical trial using a parallel, two-group, equal sample allocation design to compare the survivability of a new treatment with that of the current treatment. The median survival time for the current treatment is 1.6 years. The new treatment will be adopted if the median survival time can be shown to be higher than the current treatment. Because the true median survival time is unknown, the researcher wishes to determine the power of the logrank test to detect a difference in survival when the true median survival time for the new treatment is 2.0, 2.5, or 3.0 years.

The trial will include a recruitment period of one year, after which participants will be followed for an additional two years. It is assumed that patients will enter the study uniformly over the accrual period. The researcher estimates a loss-to-follow rate of 4% per year in both the control and the experimental groups. Past experience has led to estimates of noncompliance and drop in of 6% and 5%, respectively.

The researcher decides to investigate various sample sizes between 50 and 200 at a significance level of 0.05.

Setup

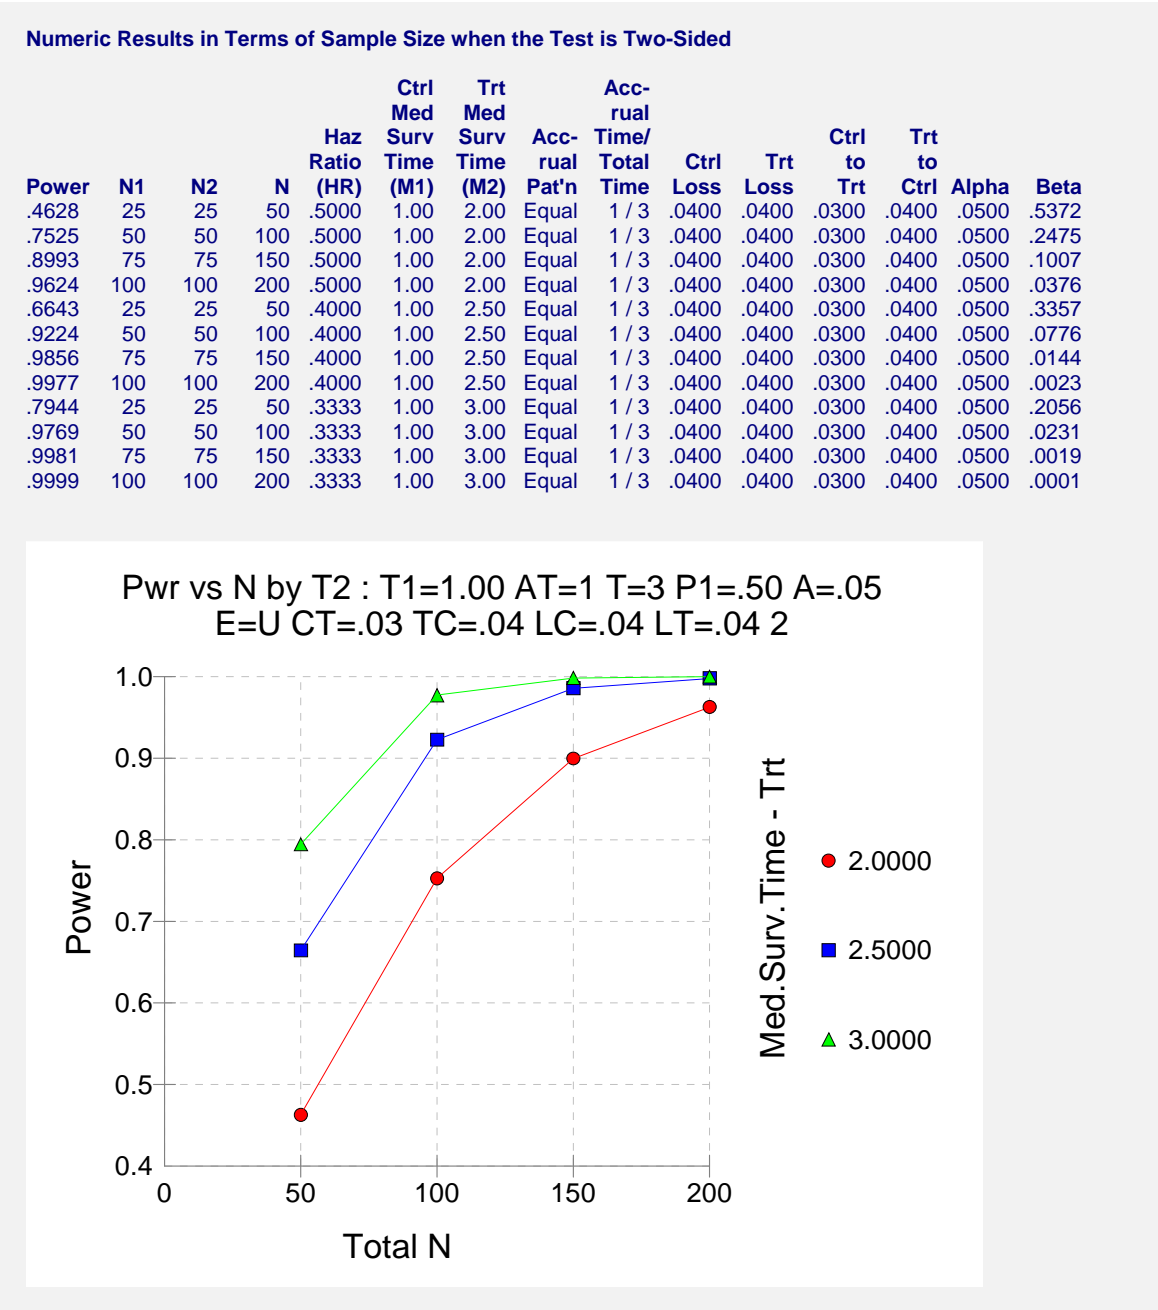
This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lakatos) [Median Survival Time]** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lakatos) using Median Survival Time**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example5** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Total Sample Size)	50 to 200 by 50
Proportion in Control Group	0.5
Proportion of Controls Lost	0.04
Proportion of Treatment Lost	0.04
Proportion of Controls Switch	0.05
Proportion of Treatment Switch	0.06
T1	1.6
Treatment Group Parameter	T2
T2	2.0 2.5 3.0
Accrual Time	1
Accrual Pattern.....	Equal
Total Time	3
Alternative Hypothesis	Two-Sided

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



This plot shows the relationship between sample size and power for the three median survival times.

Example 6 – Finding the Power using Mortality

A researcher is planning a clinical trial using a parallel, two-group, equal sample allocation design to compare the mortality rate of a new treatment with that of the current treatment. The mortality rate at one-year after the current treatment is 0.40. The new treatment will be adopted if the mortality rate after one year can be shown to be lower than the current treatment. The researcher wishes to determine the power of the logrank test to detect a difference in mortality when the true mortality rate in the new treatment group at one year is 0.20, 0.25, or 0.30.

The trial will include a recruitment period of one year, after which participants will be followed for an additional two years. It is assumed that patients will enter the study uniformly over the accrual period. The researcher estimates a loss-to-follow rate of 5% per year in both the control and the experimental groups. Past experience has lead to estimates of noncompliance and drop in of 3% and 4%, respectively.

The researcher decides to investigate various sample sizes between 50 and 250 at a significance level of 0.05.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lakatos) [Mortality]** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lakatos) using Mortality**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example6** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Total Sample Size)	50 to 250 by 50
Proportion in Control Group	0.5
Proportion of Controls Lost	0.05
Proportion of Treatment Lost	0.05
Proportion of Controls Switch	0.04
Proportion of Treatment Switch	0.03
M1	0.4
Treatment Group Parameter	M2
M2	0.20 0.25 0.30
T0 (Survival Time).....	1
Accrual Time	1
Accrual Pattern.....	Equal
Total Time	3
Alternative Hypothesis	Two-Sided

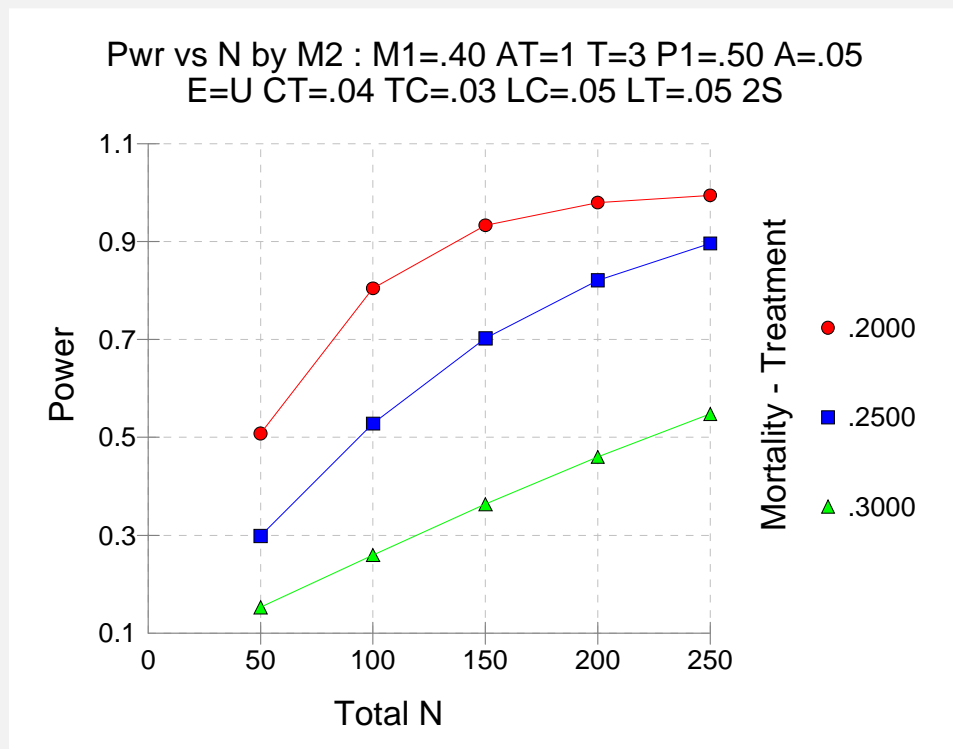
Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots

Numeric Results in Terms of Sample Size when the Test is Two-Sided and T0 is 1

Power	N1	N2	N	Mort Ratio (MR)	Ctrl Mort (M1)	Trt Mort (M2)	Acc- rual Pat'n	Acc- rual Time/ Total Time	Ctrl Loss	Trt Loss	Ctrl to Trt	Trt to Ctrl	Alpha	Beta
.5079	25	25	50	.5000	.4000	.2000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.4921
.8044	50	50	100	.5000	.4000	.2000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.1956
.9332	75	75	150	.5000	.4000	.2000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.0668
.9794	100	100	200	.5000	.4000	.2000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.0206
.9941	125	125	250	.5000	.4000	.2000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.0059
.2988	25	25	50	.6250	.4000	.2500	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.7012
.5281	50	50	100	.6250	.4000	.2500	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.4719
.7021	75	75	150	.6250	.4000	.2500	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.2979
.8207	100	100	200	.6250	.4000	.2500	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.1793
.8961	125	125	250	.6250	.4000	.2500	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.1039
.1529	25	25	50	.7500	.4000	.3000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.8471
.2597	50	50	100	.7500	.4000	.3000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.7403
.3635	75	75	150	.7500	.4000	.3000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.6365
.4603	100	100	200	.7500	.4000	.3000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.5397
.5479	125	125	250	.7500	.4000	.3000	Equal	1 / 3	.0500	.0500	.0400	.0300	.0500	.4521



This plot shows the relationship between sample size and power for the three mortality rates.

Example 7 – Converting Years to Months

A researcher is planning a clinical trial using a parallel, two-group, equal sample allocation design to compare the hazard rate of a new treatment with that of the current treatment. The hazard rate for the current treatment is 0.14. The new treatment will be adopted if the hazard rate after can be shown to be lower than the current treatment. The researcher wishes to determine the power of the logrank test to detect true hazard ratios for the new treatment of 0.4, 0.5, and 0.6.

The trial will include a recruitment period of four months, after which participants will be followed for an additional year and 8 months. It is assumed that patients will enter the study uniformly over the accrual period. The researcher estimates a loss-to-follow proportion of 4% per year in both the control and the experimental groups. Past experience has lead to estimates of noncompliance and drop in of 3% each.

The researcher decides to investigate various sample sizes between 50 and 350 at a significance level of 0.05.

Before entering the values into the Logrank Test (Hazard Ratio) window, the values stated above in terms of years must be converted to the corresponding monthly values. This can be done using the Proportions (Years to Months) tab of the Survival Parameter Conversion Tool.

Survival Parameter Conversion Tool

Survival Parameters | **Proportions (Years to Months)**

Use this window to convert proportions and rates from an annual basis to a monthly (other weekly) basis and vice versa.

Proportions

Changing a value will cause the corresponding value to be updated.

Main Proportion (e.g. Annual):	Sub Proportion (e.g. Monthly):
<input type="text" value="0.04"/>	<input type="text" value="0.00339605319892"/>
Main Rate (e.g. Annual):	Sub Rate (e.g. Monthly):
<input type="text" value="0.14"/>	<input type="text" value="0.011666666666667"/>

Number of Sub Time Units in One Main Time Unit:

MAIN RATE:
This rate will be transformed from an annual basis to a monthly basis.
It is transformed using the formula:
 $R(\text{monthly}) = R(\text{annual})/12$

715-30 Logrank Tests (Lakatos)

The number of sub time units in one main time unit is 12, since there are 12 months in a year. The yearly proportion 0.04 corresponding to the loss-to-follow 4% is converted to the monthly value of 0.00339605319892 using the relationship $P(\text{annual}) = 1 - (1 - P(\text{monthly}))^{12}$. Similarly, the yearly noncompliance and drop in values of 3% are converted to the monthly value of 0.00253504861384. The annual hazard rate of 0.14 is converted to the monthly hazard rate of 0.01166666666667 using the relationship $R(\text{monthly}) = R(\text{annual})/12$.

Setup

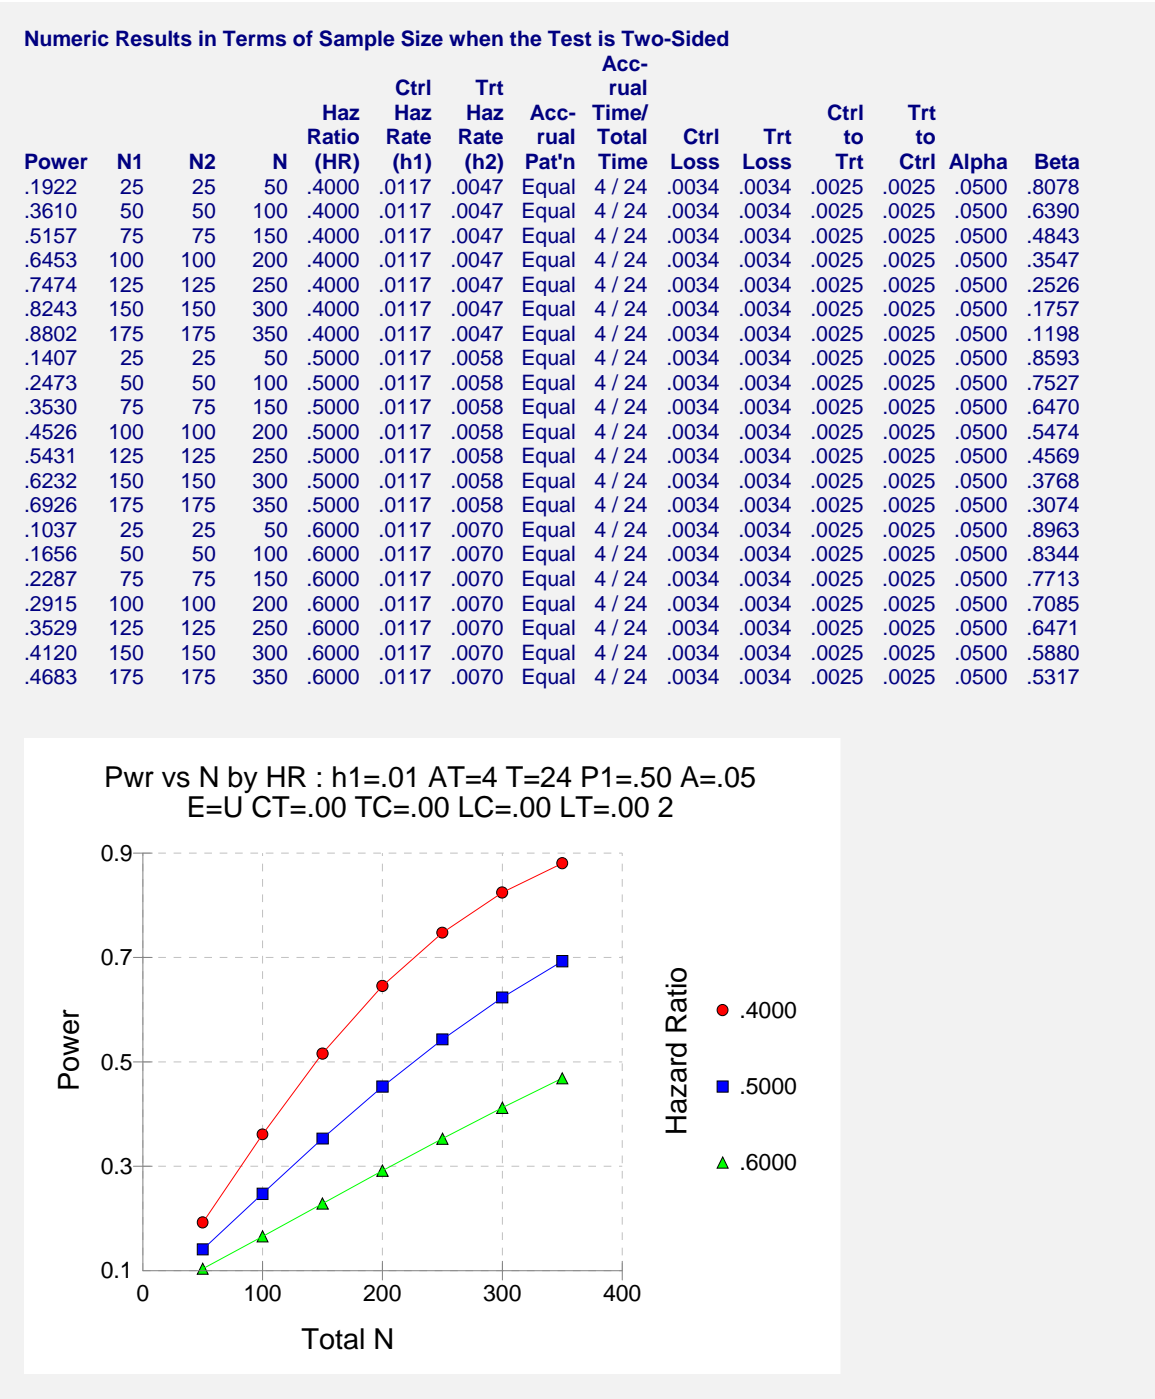
This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logrank Tests (Lakatos) [Hazard Ratio]** procedure window by clicking on **Survival Analysis and Reliability**, then **Logrank Tests**, then **Logrank Tests (Lakatos) using Hazard Ratio**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example7** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Total Sample Size)	50 to 350 by 50
Proportion in Control Group	0.5
Proportion of Controls Lost	0.00339605319892
Proportion of Treatment Lost	0.00339605319892
Proportion of Controls Switch	0.00253504861384
Proportion of Treatment Switch	0.00253504861384
h1	0.01166666666667
Treatment Group Parameter	HR (Hazard Ratio = h2/h1)
HR	0.4 0.5 0.6
Accrual Time	4
Accrual Pattern.....	Equal
Total Time	24
Alternative Hypothesis	Two-Sided

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



This plot shows the relationship between sample size and power for the three hazard ratios.

Chapter 800

Inequality Tests for One Correlation

Introduction

The correlation coefficient, ρ (rho), is a popular statistic for describing the strength of the relationship between two variables. The correlation coefficient is the slope of the regression line between two variables when both variables have been standardized by subtracting their means and dividing by their standard deviations. The correlation ranges between plus and minus one.

When ρ is used as a descriptive statistic, no special distributional assumptions need to be made about the variables (Y and X) from which it is calculated. When hypothesis tests are made, you assume that the observations are independent and that the variables are distributed according to the bivariate-normal density function. However, as with the t-test, tests based on the correlation coefficient are robust to moderate departures from this normality assumption.

The population correlation ρ is estimated by the sample correlation coefficient r . Note we use the symbol R on the screens and printouts to represent the population correlation.

Difference between Linear Regression and Correlation

The correlation coefficient is used when both X and Y are from the normal distribution (in fact, the assumption actually is that X and Y follow a bivariate normal distribution). The point is, X is assumed to be a random variable whose distribution is normal. In the linear regression context, no statement is made about the distribution of X. In fact, X is not even a random variable. Instead, it is a set of fixed values such as 10, 20, 30 or -1, 0, 1. Because of this difference in definition, we have included both Linear Regression and Correlation algorithms. This module deals with the Correlation (random X) case.

Test Procedure

The testing procedure is as follows. H_0 is the null hypothesis that the true correlation is a specific value, ρ_0 (usually, $\rho_0 = 0$). H_A represents the alternative hypothesis that the actual correlation of the population is ρ_1 , which is not equal to ρ_0 . Choose a value R_α , based on the

800-2 Inequality Tests for One Correlation

distribution of the sample correlation coefficient, so that the probability of rejecting H_0 when H_0 is true is equal to a specified value, α . Select a sample of n items from the population and compute the sample correlation coefficient, r_s . If $r_s > R_\alpha$ reject the null hypothesis that $\rho = \rho_0$ in favor of an alternative hypothesis that $\rho = \rho_1$, where $\rho_1 > \rho_0$. The power is the probability of rejecting H_0 when the true correlation is ρ_1 .

All calculations are based on the algorithm described by Guenther (1977) for calculating the cumulative correlation coefficient distribution.

Calculating the Power

Let $R(r|N, \rho)$ represent the area under a correlation density curve to the left of r . N is the sample size and ρ is the population correlation. The power of the significance test of $\rho_1 > \rho_0$ is calculated as follows:

1. Find r_α such that $1 - R(r_\alpha|N, \rho_0) = \alpha$.
2. Compute the power = $1 - R(r_\alpha|N, \rho_1)$.

Notice that the calculations follow the same pattern as for the t-test. First find the rejection region by finding the critical value (r_α) under the null hypothesis. Next, calculate the probability that a sample of size N drawn from the population defined by setting the correlation to ρ_1 is in this rejection region. This is the power.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be calculated from the values of the other parameters. Under most conditions, you would either select *Power and Beta* or *N*.

Select *N* when you want to determine the sample size needed to achieve a given power and alpha error level.

Select *Power and Beta* when you want to calculate the power.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal correlations when in fact they are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis of equal correlations when in fact they are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

The number of observations in the sample. Each observation is made up of two values: one for X and one for Y .

Effect Size

R0 (Baseline Correlation)

Specify the value of ρ_0 . Note that the range of the correlation is between plus and minus one. This value is usually set to zero.

R1 (Alternative Correlation)

Specify the value of ρ_1 , the population correlation under the alternative hypothesis. Note that the range of the correlation is between plus and minus one. The difference between R0 and R1 is being tested by this significance test.

You can enter a range of values separated by blanks or commas.

Test

Alternative Hypothesis

This option specifies the alternative hypothesis. This implicitly specifies the direction of the hypothesis test. The null hypothesis is $H_0: \rho_0 = \rho_1$.

800-4 Inequality Tests for One Correlation

Note that the alternative hypothesis enters into power calculations by specifying the rejection region of the hypothesis test. Its accuracy is critical.

Possible selections are:

- **Ha: $R_0 < R_1$**

This is the most common selection. It yields the *two-tailed* test. Use this option when you are testing whether the correlation values are different, but you do not want to specify beforehand which correlation is larger.

- **Ha: $R_0 < R_1$**

This option yields a *one-tailed* test.

- **Ha: $R_0 > R_1$**

This option also yields a *one-tailed* test.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Finding the Power

Suppose a study will be run to test whether the correlation between forced vital capacity (X) and forced expiratory value (Y) in a particular population is 0.30. Find the power when alpha is 0.01, 0.05, and 0.10 and the $N = 20, 60, 100$.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Correlation** procedure window by clicking on **Correlation**, then **One Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.01 0.05 0.10
N (Sample Size)	20 60 100
R0 (Baseline Correlation).....	0.0

Data Tab (continued)

R1 (Alternative Correlation) **0.3**
 Alternative Hypothesis **Ha: R0 <> R1**

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results**Numeric Results for Ha: R0<>R1**

Power	N	Alpha	Beta	R0	R1
0.09401	20	0.01000	0.90599	0.00000	0.30000
0.40755	60	0.01000	0.59245	0.00000	0.30000
0.68475	100	0.01000	0.31525	0.00000	0.30000
0.25394	20	0.05000	0.74606	0.00000	0.30000
0.65396	60	0.05000	0.34604	0.00000	0.30000
0.86524	100	0.05000	0.13476	0.00000	0.30000
0.37052	20	0.10000	0.62948	0.00000	0.30000
0.76282	60	0.10000	0.23718	0.00000	0.30000
0.92230	100	0.10000	0.07770	0.00000	0.30000

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N is the size of the sample drawn from the population. To conserve resources, it should be small.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

Beta is the probability of accepting a false null hypothesis. It should be small.


R0 is the value of the population correlation under the null hypothesis.

R1 is the value of the population correlation under the alternative hypothesis.

Summary Statements

A sample size of 20 achieves 9% power to detect a difference of -0.30000 between the null hypothesis correlation of 0.00000 and the alternative hypothesis correlation of 0.30000 using a two-sided hypothesis test with a significance level of 0.01000.

This report shows the values of each of the parameters, one scenario per row. The values from this table are plotted in the chart below.

Plots Section

This plot shows the relationship between alpha and power in this example.

Example 2 – Finding the Sample Size

Continuing with the last example, find the sample size necessary to achieve a power of 90% with a 0.05 significance level.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Correlation** procedure window by clicking on **Correlation**, then **One Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
R0 (Baseline Correlation).....	0.0
R1 (Alternative Correlation).....	0.3
Alternative Hypothesis	Ha: R0 <> R1

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Ha: R0<>R1					
Power	N	Alpha	Beta	R0	Ra
0.90081	112	0.05000	0.09919	0.00000	0.30000

The required sample size is 112. You would now experiment with the parameters to find out how much varying each will influence the sample size.

Example 3 – Validation using Zar

Zar (1984) page 312 presents an example in which the power of a correlation coefficient is calculated. If $N = 12$, $\alpha = 0.05$, $R_0 = 0$, and $R_1 = 0.866$, Zar calculates a power of 98% for a two-sided test.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Correlation** procedure window by clicking on **Correlation**, then **One Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	12
R0 (Baseline Correlation)	0.0
R1 (Alternative Correlation)	0.866
Alternative Hypothesis	Ha: R0 <> R1

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Ha: R0<>R1

Power	N	Alpha	Beta	R0	R1
0.98398	12	0.05000	0.01602	0.00000	0.86600

The power of 0.98 matches Zar's results.

Example 4 – Validation using Graybill

Graybill (1961) pages 211-212 presents an example in which the power of a correlation coefficient is calculated when the baseline correlation is different from zero. Let $N = 24$, $\alpha = 0.05$, and $R0 = 0.5$. Graybill calculates the power of a two-sided test when $R1 = 0.2$ and 0.3 to be 0.363 and 0.193.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Correlation** procedure window by clicking on **Correlation**, then **One Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	24
R0 (Baseline Correlation).....	0.5
R1 (Alternative Correlation).....	0.2 0.3
Alternative Hypothesis	Ha: R0 <> R1

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Ha: R0<>R1

Power	N	Alpha	Beta	R0	R1
0.36583	24	0.05000	0.63417	0.50000	0.20000
0.19950	24	0.05000	0.80050	0.50000	0.30000

The power values match Graybill's results to two decimal places.

Chapter 801

Confidence Intervals for One Correlation

Introduction

This routine calculates the sample size necessary to achieve a specified interval width or distance from the sample correlation to the confidence limit at a stated confidence level for a confidence interval for one correlation.

Caution: This procedure assumes that the correlation of the future sample will be the same as the correlation that is specified. If the sample correlation is different from the one specified when running this procedure, the interval width may be narrower or wider than specified.

Technical Details

Assuming a bivariate normal population with population correlation ρ , the transformation of the sample correlation, r ,

$$z_r = \frac{1}{2} \ln \left(\frac{1+r}{1-r} \right)$$

is approximately normally distributed with variance $1/(n - 3)$ (Fisher, 1921). The lower and upper confidence limits for ρ can be obtained by first computing

$$z_r \pm z_{1-\alpha/2} \sqrt{\frac{1}{n-3}}$$

to obtain z_L and z_U . The values of z_L and z_U can then be transformed to the correlation scale using the inverse transformations

$$r_L = \frac{e^{2z_L} - 1}{e^{2z_L} + 1}$$

and

$$r_U = \frac{e^{2z_U} - 1}{e^{2z_U} + 1}.$$

801-2 Confidence Intervals for One Correlation

One-sided limits may be obtained by replacing $\alpha/2$ by α .

For two-sided intervals, the distance from the sample correlation to each of the limits may be different. Thus, instead of specifying the distance to the limits we specify the width of the interval, W .

The basic equation for determining sample size for a two-sided interval when W has been specified is

$$W = r_U - r_L$$

For one-sided intervals, the distance from the sample correlation to limit, D , is specified.

The basic equation for determining sample size for a one-sided upper limit when D has been specified is

$$D = r_U - r$$

The basic equation for determining sample size for a one-sided lower limit when D has been specified is

$$D = r - r_L$$

Each of these equations can be solved for any of the unknown quantities in terms of the others.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population correlation is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population correlation is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as 0.90, 0.95 or 0.90 to 0.99 by 0.01.

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

Precision

Confidence Interval Width (Two-Sided)

This is the distance from the lower confidence limit to the upper confidence limit.

You can enter a single value or a list of values. The value(s) must be between 0 and 2.

Distance from R to Limit (One-Sided)

This is the distance from the sample correlation to the lower or upper limit of the confidence interval, depending on whether the Interval Type is set to Lower Limit or Upper Limit.

You can enter a single value or a list of values. The value(s) must be between 0 and 2.

Correlation

R (Sample Correlation)

Enter an estimate of the sample correlation. The sample size and width calculations assume that the value entered here is the correlation estimate that is obtained from the sample. If the sample correlation is different from the one specified here, the width may be narrower or wider than specified.

The range of the values of the sample correlation that can be entered is -1 to 1.

You can enter a range of values such as .1 .2 .3 or .1 to .5 by .1.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to construct a two-sided 95% confidence interval for the population correlation such that the width of the interval is no wider than 0.08. The researcher would like to examine a large range of correlation values to determine the effect of the correlation estimate on necessary sample size. Instead of examining only the interval width of 0.08, widths of 0.06 and 0.10 will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Correlation** procedure window by clicking on **Confidence Intervals**, then **Correlation and Regression**, then **One Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	0.06 0.08 0.10
R (Sample Correlation)	-0.9 to 0.9 by 0.1

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals for One Correlation

Confidence Level	Sample Size (N)	Target Width	Actual Width	Correlation (R)	Lower Limit	Upper Limit	Width if R = 0.0
0.950	161	0.060	0.060	-0.900	-0.926	-0.866	0.309
0.950	559	0.060	0.060	-0.800	-0.828	-0.768	0.166
0.950	1115	0.060	0.060	-0.700	-0.729	-0.669	0.117
0.950	1752	0.060	0.060	-0.600	-0.629	-0.569	0.094
0.950	2404	0.060	0.060	-0.500	-0.529	-0.469	0.080
0.950	3014	0.060	0.060	-0.400	-0.430	-0.370	0.071
0.950	3536	0.060	0.060	-0.300	-0.330	-0.270	0.066
0.950	3935	0.060	0.060	-0.200	-0.230	-0.170	0.062
0.950	4184	0.060	0.060	-0.100	-0.130	-0.070	0.061
0.950	4269	0.060	0.060	0.000	-0.030	0.030	0.060
0.950	4184	0.060	0.060	0.100	0.070	0.130	0.061
0.950	3935	0.060	0.060	0.200	0.170	0.230	0.062
0.950	3536	0.060	0.060	0.300	0.270	0.330	0.066
0.950	3014	0.060	0.060	0.400	0.370	0.430	0.071
0.950	2404	0.060	0.060	0.500	0.469	0.529	0.080
0.950	1752	0.060	0.060	0.600	0.569	0.629	0.094
0.950	1115	0.060	0.060	0.700	0.669	0.729	0.117
0.950	559	0.060	0.060	0.800	0.768	0.828	0.166
0.950	161	0.060	0.060	0.900	0.866	0.926	0.309
0.950	94	0.080	0.080	-0.900	-0.933	-0.853	0.405
0.950	317	0.080	0.080	-0.800	-0.836	-0.757	0.220
0.950	629	0.080	0.080	-0.700	-0.738	-0.658	0.156
.
.
.

References

- Cook, R. D. and Weisburg, S. 1999. Applied Regression Including Computing and Graphics. John Wiley and Sons, Inc.
- Ostle, B. and Malone, L.C. 1988. Statistics in Research. Iowa State University Press. Ames, Iowa.
- Zar, J. H. 1984. Biostatistical Analysis. Second Edition. Prentice-Hall. Englewood Cliffs, New Jersey.
- Fisher, R. A. 1921. 'On the probable error of a coefficient of correlation deduced from a small sample.' Metron, i (4), 1-32.

Report Definitions


- Confidence level is the proportion of confidence intervals (constructed with this same confidence level, sample size, etc.) that would contain the true correlation.
- Sample Size (N) is the size of the sample drawn from the population.
- Width is the distance from the lower limit to the upper limit.
- Target Width is the value of the width that is entered into the procedure.
- Actual Width is the value of the width that is obtained from the procedure.
- Correlation (R) is the assumed sample correlation.
- Lower and Upper Limit are the lower and upper limits of the confidence interval.
- Width if R = 0.0 is the maximum width for a confidence interval with sample size N.

Summary Statements

A sample size of 161 produces a two-sided 95% confidence interval with a width equal to 0.060 when the sample correlation is -0.900.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the sample correlation for the three confidence interval widths.

Example 2 – Validation using Cook and Weisberg

Cook and Weisberg (1999), page 78, give an example of a calculation for an two-sided confidence interval for a single correlation when the confidence level is 95%, the sample correlation is -0.889, and the confidence interval width is 0.428. The sample size is 9.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for One Correlation** procedure window by clicking on **Confidence Intervals**, then **Correlation and Regression**, then **One Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Confidence Interval Width (Two-Sided) ..	0.428
R (Sample Correlation)	-0.889

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Confidence Level	Sample Size (N)	Target Width	Actual Width	Correlation (R)	Lower Limit	Upper Limit	Width if R = 0.0
0.950	9	0.428	0.428	-0.889	-0.977	-0.549	1.328

PASS also calculated the sample size to be 9.

801-8 Confidence Intervals for One Correlation

Chapter 805

Inequality Tests for Two Correlations

Introduction

The correlation coefficient (or correlation), ρ , is a popular parameter for describing the strength of the association between two variables. The correlation coefficient is the slope of the regression line between two variables when both variables have been standardized. It ranges between plus and minus one. This chapter covers the case in which you want to test the difference between two correlations, each coming from a separate sample.

Since the correlation is the standardized slope between two variables, you could also apply this procedure to the case in which you want to test whether the slopes in two groups are equal.

Test Procedure

In the following discussion, ρ is the population correlation coefficient and r is the value calculated from a sample. The testing procedure is as follows. H_0 is the null hypothesis that $\rho_1 = \rho_2$. H_A represents the alternative hypothesis that $\rho_1 \neq \rho_2$ (one-tailed hypotheses are also available). To construct the hypothesis test, transform the correlations using the Fisher- z transformation.

$$z_i = \frac{1}{2} \log \left(\frac{1 + r_i}{1 - r_i} \right)$$

$$Z_i = \frac{1}{2} \log \left(\frac{1 + \rho_i}{1 - \rho_i} \right)$$

This transformation is used because the combined distribution of r_1 and r_2 is too difficult to work with, but the distributions of z_1 and z_2 are approximately normal.

Note that the reverse transformation is

$$r_i = \frac{e^{z_i} - e^{-z_i}}{e^{z_i} + e^{-z_i}}$$

805-2 Inequality Tests for Two Correlations

Once the correlations have been converted into z values, the normal distribution may be used to conduct the test of $Z_1 - Z_2$. The standard deviation of the difference is given by

$$\sigma_{z_1 - z_2} = \sqrt{\frac{1}{N_1 - 3} + \frac{1}{N_2 - 3}}$$

The test statistic is given by

$$z = \frac{(z_1 - z_2) - (Z_1 - Z_2)}{\sigma_{z_1 - z_2}}$$

Note that the lower case z 's represent the values calculated from the two samples and the upper case Z 's represent the hypothesized population values.

Calculating the Power

1. Find z_α such that $1 - \Phi(z_\alpha) = \alpha$, where $\Phi(x)$ is the area under the standardized normal curve to the left of x .
2. Calculate: $Z_1 = \frac{1}{2} \log\left(\frac{1 + \rho_1}{1 - \rho_1}\right)$
3. Calculate: $Z_2 = \frac{1}{2} \log\left(\frac{1 + \rho_2}{1 - \rho_2}\right)$
4. Calculate: $\sigma_{z_1 - z_2} = \sqrt{\frac{1}{N_1 - 3} + \frac{1}{N_2 - 3}}$
5. Calculate: $x_a = Z_1 - Z_2 + z_\alpha \sigma_{z_1 - z_2}$
6. Calculate: $z_a = \frac{x_a}{\sigma_{z_1 - z_2}}$
7. Calculate: Power = $1 - \Phi(z_a)$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are $R1$, $R2$, $N1$, $N2$, $Alpha$, and $Power$ and $Beta$. Under most situations, you will select either $Power$ and $Beta$ or $N1$.

Select $N1$ when you want to calculate the sample size needed to achieve a given power and alpha level.

Select $Power$ and $Beta$ when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal correlations when in fact they are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis of equal correlations when in fact they are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N1 (Sample Size Group 1)

Enter a value (or range of values) for the sample size of this group. Note that these values are ignored when you are solving for $N1$. You may enter a range of values such as *10 to 100 by 10*.

N2 (Sample Size Group 2)

Enter a value (or range of values) for the sample size of group 2 or enter *Use R* to base $N2$ on the value of $N1$. You may enter a range of values such as *10 to 100 by 10*.

- **Use R**

When *Use R* is entered here, $N2$ is calculated using the formula

$$N2 = [R(N1)]$$

where R is the Sample Allocation Ratio and $[Y]$ is the first integer greater than or equal to Y . For example, if you want $N1 = N2$, select *Use R* and set $R = 1$.

R (Sample Allocation Ratio)

Enter a value (or range of values) for R , the allocation ratio between samples. This value is only used when $N2$ is set to *Use R*.

When used, $N2$ is calculated from $N1$ using the formula: $N2 = [R(N1)]$ where $[Y]$ is the next integer greater than or equal to Y . Note that setting $R = 1.0$ forces $N2 = N1$.

Effect Size

R1 (Correlation Group 1)

Specify the value of the population correlation coefficient of group one. Possible values range between plus and minus one.

You can enter a single value or a range of values separated by commas or blanks.

Note that the power depends on the specific values of $R1$ and $R2$, not just their difference. Hence, $R1 = 0$ and $R2 = 0.3$ will have a different power from $R1 = 0.3$ and $R2 = 0.6$.

R2 (Correlation Group 2)

Specify the value of the population correlation coefficient from group two under the alternative hypothesis. Possible values range between plus and minus one.

You can enter a single value or a range of values separated by commas or blanks.

Test

Alternative Hypothesis

This option specifies the alternative hypothesis. This implicitly specifies the direction of the hypothesis test. The null hypothesis is always $H_0: \rho_1 = \rho_2$.

Possible selections are:

- **Ha: $R1 \neq R2$**

This is the most common selection. It yields the *two-tailed* test. Use this option when you are testing whether the correlation values are different, but you do not want to specify beforehand which value is larger.

- **Ha: $R1 < R2$**

This option yields a *one-tailed* test. When you use this option, you should be careful to enter values for $R1$ and $R2$ that follow this relationship.

- **Ha: $R1 > R2$**

This option yields a *one-tailed* test. When you use this option, you should be careful to enter values for $R1$ and $R2$ that follow this relationship.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Finding the Power

A researcher wants to compare the relationship between weight and heart rate in males and females. If the correlation between weight and heart rate is 0.3 in a sample of 100 males and 0.5 in a sample of 100 females, what is the power of a two sided test for the difference between correlations at the 0.01 and 0.05 significance levels? Also compute the power for samples of 20, 200, 300, 400, and 600.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Correlations** procedure window by clicking on **Correlation**, then **Two Correlations**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting.</i>
Alpha	0.01 0.05
N1 (Sample Size Group 1).....	20 100 200 300 400 600
N2 (Sample Size Group 2).....	Use R
R (Sample Allocation Ratio).....	1.0
R1 (Correlation Group 1).....	0.3
R2 (Correlation Group 2).....	0.5
Alternative Hypothesis	Ha: $R1 <> R2$

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when Ha: $R_1 \neq R_2$

Power	N1	N2	Allocation Ratio	R1	R2	Difference (R1-R2)	Alpha	Beta
0.03081	20	20	1.000	0.30000	0.50000	-0.20000	0.01000	0.96919
0.18250	100	100	1.000	0.30000	0.50000	-0.20000	0.01000	0.81750
0.42230	200	200	1.000	0.30000	0.50000	-0.20000	0.01000	0.57770
0.63541	300	300	1.000	0.30000	0.50000	-0.20000	0.01000	0.36459
0.78888	400	400	1.000	0.30000	0.50000	-0.20000	0.01000	0.21112
0.94144	600	600	1.000	0.30000	0.50000	-0.20000	0.01000	0.05856
0.10760	20	20	1.000	0.30000	0.50000	-0.20000	0.05000	0.89240
0.38603	100	100	1.000	0.30000	0.50000	-0.20000	0.05000	0.61397
0.66271	200	200	1.000	0.30000	0.50000	-0.20000	0.05000	0.33729
0.83200	300	300	1.000	0.30000	0.50000	-0.20000	0.05000	0.16800
0.92196	400	400	1.000	0.30000	0.50000	-0.20000	0.05000	0.07804
0.98548	600	600	1.000	0.30000	0.50000	-0.20000	0.05000	0.01452

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N1 and N2 are the sizes of the samples drawn from the two populations. To conserve resources, it should be small.

Allocation Ratio is N_1/N_2 so that $N_2 = N_1 \times R$.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

Beta is the probability of accepting a false null hypothesis. It should be small.

R1 is the value of both correlations under the null hypothesis.

R2 is the correlation in group two under the alternative hypothesis.


Summary Statements

Group sample sizes of 20 and 20 achieve 3% power to detect a difference of 0.20000 between the null hypothesis that both group correlations are 0.30000 and the alternative hypothesis that the correlation in group 2 is 0.50000 using a two-sided z test (which uses Fisher's z-transformation) with a significance level of 0.01000.

This report shows the values of each of the parameters, one scenario per row. The definitions of each column are given in the Report Definitions section of the report, so they will not be repeated here.

The values from this table are plotted in the chart below.

Plots Section



This plot shows the relationship between alpha, power, and sample size in this example.

Example 2 – Finding the Sample Size

Continuing with the previous example, suppose the researchers want to determine the exact sample size necessary to achieve 90% power at a 0.05 significance level.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Correlations** procedure window by clicking on **Correlation**, then **Two Correlations**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N1
Power	0.90
Alpha	0.05
N1 (Sample Size Group 1)	<i>Ignored since this is the Find setting</i>
N2 (Sample Size Group 2)	Use R
R (Sample Allocation Ratio)	1.0
R1 (Correlation Group 1)	0.3
R2 (Correlation Group 2)	0.5
Alternative Hypothesis	Ha: R1 <> R2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when Ha: R1<>R2								
	Allocation					Difference		
Power	N1	N2	Ratio	R1	R2	(R1-R2)	Alpha	Beta
0.90040	369	369	1.000	0.30000	0.50000	-0.20000	0.05000	0.09960

PASS has calculated the sample size as 369 per group.

Example 3 – Validation using Zar

Zar (1984) page 314 presents an example of calculating the power for a test of two correlations. In his example, when $N1 = 95$, $N2 = 98$, $R1 = 0.84$, $R2 = 0.78$, and $\alpha = 0.05$, the power is 22% for a two-sided test.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Correlations** procedure window by clicking on **Correlation**, then **Two Correlations**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
R1 (Correlation Group 1).....	0.84
R2 (Correlation Group 2).....	0.78
Alpha	0.05
Power	<i>Ignored since this is the Find setting</i>
Alternative Hypothesis	Ha: R1 <> R2
N1 (Sample Size Group 1).....	95
N2 (Sample Size Group 2).....	98
R (Sample Allocation Ratio).....	1.0

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when Ha: R1<>R2								
	N1	N2	Allocation Ratio	R1	R2	Difference (R1-R2)	Alpha	Beta
Power	95	98	1.032	0.84000	0.78000	0.06000	0.05000	0.77502

PASS has also calculated the power to be 22%.

Chapter 810

Inequality Tests for Intraclass Correlation

Introduction

The intraclass correlation coefficient is often used as an index of reliability in a measurement study. In these studies, there are N observations made on each of K individuals. These individuals represent a factor observed at random. This design arises when K subjects are each rated by N raters.

The intraclass correlation coefficient may be thought of as the correlation between any two observations made on the same subject. When this correlation is high, the observations on a subject tend to match, and the measurement reliability is ‘high.’

Technical Details

Our formulation comes from Walter, Eliasziw, and Donner (1998) and Winer (1991). Denote response j of subject i by Y_{ij} , where $i = 1, 2, \dots, K$ and $j = 1, 2, \dots, N$. The model for this situation is

$$Y_{ij} = \mu + a_i + e_{ij}$$

where the random subject effects a_i are normally distributed with mean 0 and variance σ_a^2 and the measurement errors, e_{ij} are normally distributed with mean 0 and variance σ_e^2 . We assume that the subject effects and the measurement errors are independent. The intraclass correlation is then defined as

$$\rho = \frac{\sigma_a^2}{\sigma_a^2 + \sigma_e^2}$$

The hypothesis test is stated formally as

$$\begin{aligned} H_0: \rho &= \rho_0 \\ H_1: \rho &= \rho_1 > \rho_0 \end{aligned}$$

810-2 Inequality Tests for Intraclass Correlation

This hypothesis is tested from the data of a one-way analysis of variance table using the value:

$\frac{MS_a}{MS_e}$. The critical value for the test statistic is

$$C(F_{1-\alpha/2, df1, df2})$$

where

$$C = 1 + \left[\frac{N\rho_0}{1 - \rho_0} \right]$$

$$df1 = K - 1$$

$$df2 = K(N - 1)$$

The power of this test procedure is given by

$$Power = 1 - P(F \geq C_0 F_{1-\alpha/2, df1, df2})$$

where

$$C_0 = \frac{1 + N\rho_0 / (1 - \rho_0)}{1 + N\rho_1 / (1 - \rho_1)}$$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *R0*, *R1*, *K*, *N*, *Alpha*, and *Power* (or *Beta*).

Under most situations, you will select either *Power and Beta* to calculate power or *N* to calculate sample size.

Note that the value selected here always appears as the vertical axis on the charts.

The program is set up to evaluate power directly. For the other parameters, a search is made using an iterative procedure until an appropriate value is found.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis when it is true.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

K (Number of Subjects)

Enter a value (or range of values) for the number of subjects, K , that are measured.

You may enter a range of values such as *50,150,250* or *50 to 300 by 50*.

N (Observations Per Subject)

Enter a value (or range of values) for the number of observations, N , per subjects. In a reliability study, this is the number of raters (assuming each subject is rated by all raters).

You may enter a range of values such as *2,3,4* or *2 to 12 by 2*.

Effect Size

R0 (Intraclass Correlation 0)

This is the value(s) of the intraclass correlation coefficient when the null hypothesis is true. You may enter a single value or a list of values. The range of $R0$ is between zero and $R1$.

The intraclass correlation is calculated as $V(A)/[V(A)+V(E)]$ where $V(E)$ is the variation within a subject and $V(A)$ is the variation between subjects. It is a measure of the extent to which the observations within a subject are similar (or dependent) relative to observations from other subjects.

R1 (Intraclass Correlation 1) > R0

This is the value(s) of the intraclass correlation coefficient when the alternative hypothesis is true. You may enter a single value or a list of values. The range of $R1$ is between $R0$ and one.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Power

Suppose that a study is to be conducted in which $R0 = 0.2$; $R1 = 0.3$; $K = 50$ to 250 by 100 ; Alpha = 0.05 ; and $N = 2$ to 5 by 1 and beta is to be calculated.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Intraclass Correlation** procedure window by clicking on **Correlation**, then **Intraclass Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
K (Number of Subjects)	50 to 250 by 100
N (Observations Per Subject)	2 to 5 by 1
R0 (Intraclass Correlation 0)	0.2
R1 (Intraclass Correlation 1) > R0	0.3

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results

Power	Number of Subjects	Observations Per Subject	Intraclass Correlation 0	Intraclass Correlation 1	Alpha	Beta
0.18333	50	2	0.200	0.300	0.05000	0.81667
0.29534	50	3	0.200	0.300	0.05000	0.70466
0.38528	50	4	0.200	0.300	0.05000	0.61472
0.45522	50	5	0.200	0.300	0.05000	0.54478
0.36558	150	2	0.200	0.300	0.05000	0.63442
0.60094	150	3	0.200	0.300	0.05000	0.39906
0.74538	150	4	0.200	0.300	0.05000	0.25462
0.83005	150	5	0.200	0.300	0.05000	0.16995
0.51549	250	2	0.200	0.300	0.05000	0.48451
0.78790	250	3	0.200	0.300	0.05000	0.21210
0.90522	250	4	0.200	0.300	0.05000	0.09478
0.95403	250	5	0.200	0.300	0.05000	0.04597

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

K is the number of subjects.

N is the number of observations per subject in the sample.

R0 is intraclass correlation assuming the null hypothesis.

R1 is intraclass correlation assuming the alternative hypothesis.

Alpha is the probability of rejecting a true null hypothesis. It should be small.


Beta is the probability of accepting a false null hypothesis. It should be small.

Summary Statements

A sample size of 50 subjects with 2 observations per subject achieves 18% power to detect an intraclass correlation of 0.300 under the alternative hypothesis when the intraclass correlation under the null hypothesis is 0.200 using an F-test with a significance level of 0.05000.

This report shows the power for each of the scenarios.

Plots Section



This plot shows the relation between power, number of subjects (K), and observations per subject (N).

Example 2 – Validation using Walter

Walter *et al.* (1998) page 106 give a table of sample sizes. When R_0 is 0.2, R_1 is 0.3, K is 544, N is 2, and Alpha is 0.05, Beta is 0.80.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Intraclass Correlation** procedure window by clicking on **Correlation**, then **Intraclass Correlation**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
K (Number of Subjects)	544
N (Observations Per Subject)	2
R_0 (Intraclass Correlation 0)	0.2
R_1 (Intraclass Correlation 1) > R_0	0.3

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results						
Power	Number of Subjects	Observations Per Subject	Intraclass Correlation 0	Intraclass Correlation 1	Alpha	Beta
0.80033	544	2	0.200	0.300	0.05000	0.19967

PASS has also calculated the power as 0.80.

Chapter 811

Kappa Test for Agreement Between Two Raters

Introduction

This module computes power and sample size for the test of agreement between two raters using the kappa statistic. The power calculations are based on the results in Flack, Afifi, Lachenbruch, and Schouten (1988). Calculations are based on ratings for k categories from two raters or judges. You are able to vary category frequencies on a single run of the procedure to analyze a wide range of scenarios all at once. For further information about kappa analysis, see chapter 18 of Fleiss, Levin, and Paik (2003).

Technical Details

Suppose that N subjects are each assigned independently to one of k categories by two separate judges or raters. The results are placed in a $k \times k$ contingency table. Each p_{ij} represents the proportion of subjects that Rater A classified in category i , but Rater B classified in category j , with $i, j = 1, 2, \dots, k$. The proportions $p_{i.}$ and $p_{.j}$ are the frequencies or marginal probabilities of assignment into categories i and j for Rater A and Rater B, respectively. For each rater, the category frequencies sum to one.

Rater A	Rater B				Total
	1	2	...	k	
1	p_{11}	p_{12}	...	p_{1k}	$p_{1.}$
2	p_{21}	p_{22}	...	p_{2k}	$p_{2.}$
\vdots	\vdots	\vdots	\ddots	\vdots	\vdots
k	p_{k1}	p_{k2}	...	p_{kk}	$p_{k.}$
Total	$p_{.1}$	$p_{.2}$...	$p_{.k}$	1

811-2 Kappa Test for Agreement Between Two Raters

The proportions on the diagonal, p_{ii} , represent the proportion of subjects in each category for which the two raters agreed on the assignment. The overall proportion of observed agreement is

$$p_o = \sum_{i=1}^k p_{ii} ,$$

and the overall proportion of agreement expected by chance is

$$p_e = \sum_{i=1}^k p_{i.} p_{.i} .$$

The overall value of kappa, which measures the degree of rater agreement, is then

$$\kappa = \frac{p_o - p_e}{1 - p_e} .$$

A kappa value of 1 represents perfect agreement between the two raters. A kappa value of 0 indicates no more rater agreement than that expected by chance. A kappa value of -1 would indicate perfect disagreement between the raters.

The true value of kappa can be estimated by replacing the observed and expected proportions by their sample estimates

$$\hat{\kappa} = \frac{\hat{p}_o - \hat{p}_e}{1 - \hat{p}_e} ,$$

where

$$\hat{p}_o = \sum_{i=1}^k \hat{p}_{ii}$$
$$\hat{p}_e = \sum_{i=1}^k \hat{p}_{i.} \hat{p}_{.i} .$$

The minimum possible value of $\hat{\kappa}$ depends on the marginal proportions. If the marginal proportions are such that $\hat{p}_e = 0.5$, then the minimum value is -1. Otherwise, the minimum value is between -1 and 0.

The standard error of $\hat{\kappa}$ is

$$s.e.(\hat{\kappa}) = \frac{\tau(\hat{\kappa})}{\sqrt{N}} ,$$

where

$$\tau(\hat{\kappa}) = \frac{1}{(1 - \hat{p}_e)^2} \left\{ p_o (1 - \hat{p}_e)^2 + (1 - p_o)^2 \sum_{i=1}^k \sum_{j=1}^k p_{ij} (p_{i.} + p_{.j})^2 \right. \\ \left. - 2(1 - p_o)(1 - \hat{p}_e) \sum_{i=1}^k p_{ii} (p_{i.} + p_{.i})^2 - (p_o \hat{p}_e - 2\hat{p}_e + p_o)^2 \right\}^{1/2} .$$

Again, an estimate of the standard error can be obtained by replacing the unknown values p_{ij} by their sample estimates \hat{p}_{ij} .

Hypothesis Tests

One- and two-sided hypothesis tests can be conducted using the test statistic

$$z = \frac{\hat{\kappa} - \kappa_0}{s.e.(\hat{\kappa})},$$

where κ_0 is the null hypothesized value of kappa, and the denominator is the estimated standard error. For a one-sided alternative, the test rejects H_0 if $|z| \geq z_\alpha$, where z_α is the value that leaves α in the upper tail of the standard normal distribution. For a two-sided alternative, the test rejects H_0 if $|z| \geq z_{\alpha/2}$.

Power Calculation

The standard error for the kappa statistic is based on values p_{ij} , which are unknown prior to conducting a study. Therefore, the power is computed at the maximum standard error based on given category frequencies or marginal probabilities. The following steps are taken to compute the power of the test.

1. Determine the category assignment frequencies for both raters. In practice, the category frequencies may not be equivalent, but the standard error maximization method of Flack, Afifi, Lachenbruch, and Schouten (1988) assumes that the category frequencies are equal for both raters. Therefore, only one set of frequencies is needed.
2. Determine the maximum standard error under the null and alternative hypotheses for the given marginal frequencies. This is equivalent to finding the maximum $\tau(\hat{\kappa})$ under the null and alternative hypotheses.
3. Find the critical value using the standard normal distribution. The critical value, $z_{critical}$, is that value of z that leaves exactly the target value of alpha (or alpha/2) in the upper tail of the standard normal distribution. For example, for an upper-tailed test with a target alpha of 0.05, the critical value is 1.645.
4. Without loss of generality, for a one-sided test of the alternative hypothesis that $\kappa > \kappa_0$, compute the power at an alternative value of kappa, κ_1 , as

$$\begin{aligned} 1 - \beta &= \Pr(z \geq z_{critical} | H_1) \\ &= 1 - \Phi(u) \end{aligned},$$

where $\Phi()$ is the cumulative standard normal distribution and

$$u = \frac{\sqrt{N}(\kappa_0 - \kappa_1) + z_{critical} \max \tau(\hat{\kappa} | \kappa = \kappa_0)}{\max \tau(\hat{\kappa} | \kappa = \kappa_1)}.$$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *Power and Beta*, *N (Sample Size)*, or *K1*.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal means when in fact the means are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when a true null hypothesis is rejected. In this procedure, a type-I error occurs when you reject the null hypothesis of equal means when in fact the means are equal.

Values must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size / Groups – Sample Size Multiplier

N (Sample Size)

Enter a value for the sample size (*N*). This is the number of subjects rated by the two judges in the study. You may enter a range such as *10 to 100 by 10* or a list of values separated by commas or blanks such as *50 60 70*.

Effect Size – Kappa

K1 (Kappa|H1)

This is the value of Kappa under the alternative hypothesis, H1. Values must be between -1 and 1. If a one-sided hypothesis is used, the alternative value(s) for Kappa in relation to the null value(s), K0, should match the direction of the alternative hypothesis. You can enter a list of values separated by blanks or commas such as *0.6 0.7 0.8* or *0.6 to 0.8 by 0.05*.

K0 (Kappa|H0)

This is the value of Kappa under the null hypothesis, H0. Values must be between -1 and 1. If a one-sided hypothesis is used, the alternative value(s) for Kappa in relation to the null value(s), K0, should match the direction of the alternative hypothesis. You can enter a list of values separated by blanks or commas such as *0.2 0.3 0.4* or *0.2 to 0.4 by 0.05*.

Effect Size – Categories

P (Category Frequencies)

Specify two or more category frequencies, proportions, or marginal proportions. These are the proportions of subjects assigned to each category by the two raters or judges. In practice the proportions may not be exactly equal for the two raters, but the power calculations assume that the category frequencies are equal for the two raters. Each proportion set must sum to 1.

Several sets of proportions can be entered by using the *PASS* spreadsheet. To launch the spreadsheet, click on the “Spreadsheet” button above the box. To select columns from the spreadsheet, click on the button with the arrow pointing down on the right of the box. Specify the column (or columns) to be used by beginning your entry with an equals sign, e.g. enter *=C1-C3*.

List Input

Specify a single set of proportions as a list. For example, with three groups you might enter *0.2 0.3 0.5*.

Spreadsheet Column Input

Specify more than one set of proportions using the column input syntax

= [column 1] [column 2] etc.

For example, if you have three proportion sets stored in the spreadsheet in columns C1, C2, and C3, you would enter *=C1 C2 C3* in the P (Category Frequencies) box.

Each column in the spreadsheet corresponds to a single set of proportions. The columns may contain different numbers of frequencies, but in all cases, the values in each column must sum to 1.

Test

Alternative Hypothesis (H1)

Specify the alternative hypothesis of the test. Since the null hypothesis is the opposite, specifying the alternative is all that is needed. The alternative hypothesis determines how the alternative value(s) of Kappa (K1) should be entered. Usually, the two-sided option is selected.

For a one-sided alternative hypothesis test of $K1 > K0$, all values for K1 should be greater than K0.

811-6 Kappa Test for Agreement Between Two Raters

For a one-sided alternative hypothesis test of $K_1 < K_0$, all values for K_1 should be less than K_0 .

For a two-sided test, the values for K_1 can be either greater than or less than K_0 .

Example 1 – Finding the Power

Suppose a study is being planned to measure the degree of inter-rater agreement for two psychiatrists. The two psychiatrists will independently classify each of a series of patients into one of three diagnostic categories: personality disorder, neurosis, or psychosis. The study will then determine how well the psychiatrists “agree” with a hypothesis test using the kappa statistic.

Before the data are collected, the organizers would like to study the relationship between sample size and power. From previous experience, they have determined to use frequencies of 0.4, 0.5, and 0.1 for the personality disorder, neurosis, and psychosis diagnoses, respectively. They would like to determine the power for detecting alternative kappa values of 0.5, 0.6, and 0.7 when the null value is 0.4. A two-sided hypothesis test will be conducted at $\alpha = 0.05$. What will be the power for a wide range of sample sizes?

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Kappa Test for Agreement Between Two Raters** procedure window by clicking on **Correlation**, then **Kappa Test for Rater Agreement**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	30 to 200 by 10
K1	0.5 0.6 0.7
K0	0.4
P (Category Frequencies)	0.4 0.5 0.1
Alternative Hypothesis (H1)	Two-Sided

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results

Test Type = Two-sided Z test.

H0: Kappa = K0 vs. H1: Kappa <> K0.

Power	Sample Size N	Kappa H0 K0	Kappa H1 K1	Alpha	Beta	- Rating Categories - k	Frequencies
0.07748	30	0.40	0.50	0.05000	0.92252	3	0.40, 0.50, 0.10
0.19421	30	0.40	0.60	0.05000	0.80579	3	0.40, 0.50, 0.10
0.43345	30	0.40	0.70	0.05000	0.56655	3	0.40, 0.50, 0.10
0.09199	40	0.40	0.50	0.05000	0.90801	3	0.40, 0.50, 0.10
0.26055	40	0.40	0.60	0.05000	0.73945	3	0.40, 0.50, 0.10
0.58208	40	0.40	0.70	0.05000	0.41792	3	0.40, 0.50, 0.10
0.10677	50	0.40	0.50	0.05000	0.89323	3	0.40, 0.50, 0.10
0.32746	50	0.40	0.60	0.05000	0.67254	3	0.40, 0.50, 0.10
0.70452	50	0.40	0.70	0.05000	0.29548	3	0.40, 0.50, 0.10
0.12180	60	0.40	0.50	0.05000	0.87820	3	0.40, 0.50, 0.10
0.39325	60	0.40	0.60	0.05000	0.60675	3	0.40, 0.50, 0.10
0.79842	60	0.40	0.70	0.05000	0.20158	3	0.40, 0.50, 0.10
0.13704	70	0.40	0.50	0.05000	0.86296	3	0.40, 0.50, 0.10
0.45663	70	0.40	0.60	0.05000	0.54337	3	0.40, 0.50, 0.10
0.86661	70	0.40	0.70	0.05000	0.13339	3	0.40, 0.50, 0.10
0.15246	80	0.40	0.50	0.05000	0.84754	3	0.40, 0.50, 0.10
0.51666	80	0.40	0.60	0.05000	0.48334	3	0.40, 0.50, 0.10
0.91404	80	0.40	0.70	0.05000	0.08596	3	0.40, 0.50, 0.10
.
.
.

(Report Continues)

References

Flack, V.F., Afifi, A.A., Lachenbruch, P.A., and Schouten, H.J.A. 1988. 'Sample Size Determinations for the Two Rater Kappa Statistic'. Psychometrika 53, No. 3, 321-325.

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N is the total sample size.

K0 is the value of Kappa under the null hypothesis, H0.

K1 is the value of Kappa under the alternative hypothesis, H1.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

Beta is the probability of accepting a false null hypothesis. It should be small.

k is the number of rating categories.

Frequencies lists the rating category frequencies. The number of frequencies is equal to k.

Summary Statements

In a test for agreement between two raters using the Kappa statistic, a sample size of 30 subjects achieves 8% power to detect a true Kappa value of 0.50 in a test of H0: Kappa = 0.40 vs. H1: Kappa <> 0.40 when there are 3 categories with frequencies equal to 0.40, 0.50, and 0.10. This power calculation is based on a significance level of 0.05000.

This report shows the numeric results of this power study. Following are the definitions of the columns of the report.

Power

The probability of rejecting a false null hypothesis.

N

The total sample size for the study.

811-8 Kappa Test for Agreement Between Two Raters

K0

The value of kappa under the null hypothesis.

K1

The value of kappa under the alternative hypothesis.

Alpha

The probability of rejecting a true null hypothesis. This is often called the significance level.

Beta

The probability of accepting a false null hypothesis.


k

The number of rating categories.

Frequencies

The rating category frequencies used.

Plots Section



This plot gives a visual presentation to the results in the Numeric Report. We can quickly see the impact on the power of increasing the sample size for the different values of K1.

When you create one of these plots, it is important to use trial and error to find an appropriate range for the horizontal variable so that you have results with both low and high power.

Example 2 – Finding the Sample Size

Continuing with the last example, we will determine how large the sample size would need to be for the three values of K1 to have the power at least 0.95 with an alpha level of 0.05.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Kappa Test for Agreement Between Two Raters** procedure window by clicking on **Correlation**, then **Kappa Test for Rater Agreement**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Power	0.95
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
K1	0.5 0.6 0.7
K0	0.4
P (Category Frequencies)	0.4 0.5 0.1
Alternative Hypothesis (H1)	Two-Sided

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results

Test Type = Two-sided Z test.

H0: Kappa = K0 vs. H1: Kappa <> K0.

	Sample Size	Kappa H0	Kappa H1			- Rating Categories -	
Power	N	K0	K1	Alpha	Beta	k	Frequencies
0.95003	983	0.40	0.50	0.05000	0.04997	3	0.40, 0.50, 0.10
0.95031	228	0.40	0.60	0.05000	0.04969	3	0.40, 0.50, 0.10
0.95078	92	0.40	0.70	0.05000	0.04922	3	0.40, 0.50, 0.10

The required sample sizes are 983, 228, and 92 for alternative kappa values of 0.5, 0.6, and 0.7, respectively.

Example 3 – Finding the Minimum Detectable Kappa

Continuing with the last example, we will now determine what is the minimum value of kappa that can be detected with 100 subjects and power of 0.95.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Kappa Test for Agreement Between Two Raters** procedure window by clicking on **Correlation**, then **Kappa Test for Rater Agreement**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	K1 (K1 > K0)
Power	0.95
Alpha	0.05
N (Sample Size)	200
K1	<i>Ignored since this is the Find setting</i>
K0	0.4
P (Category Frequencies)	0.4 0.5 0.1
Alternative Hypothesis (H1)	Two-Sided
Reports Tab	
Decimal Places - Kappa	4

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results

Test Type = Two-sided Z test.

H0: Kappa = K0 vs. H1: Kappa <> K0.

	Sample					- Rating Categories -	
Power	Size	Kappa H0	Kappa H1	Alpha	Beta	k	Frequencies
0.95000	N 200	K0 0.4000	K1 0.6122	0.05000	0.05000	3	0.40, 0.50, 0.10

The test detects a kappa value of 0.6122 with 95% power.

Example 4 – Validation using Flack, Afifi, Lachenbruch, and Schouten (1988)

Flack, Afifi, Lachenbruch, and Schouten (1988) page 324 presents a table (Table 2) of calculated sample sizes required for 80% power in a one-sided test of $H_1: \text{Kappa} > 0.4$ vs. $H_1: \text{Kappa} = 0.4$ computed at $K_1 = 0.6$ and $\alpha = 0.05$. The sample sizes are computed for various sets of category frequencies.

Table 2

PD	Frequencies		Sample Size for 80% Power
	N	PS	
0.50	0.26	0.24	93
0.50	0.30	0.20	99
0.55	0.30	0.15	109
0.60	0.30	0.10	119
0.60	0.21	0.19	107

This example will replicate these results.

Setup

This section presents the values of each of the parameters needed to run this example. First, you will need to open the Flack Validation Proportions.S0 dataset by selecting **Tools**, then **Spreadsheet** from the PASS Home window menus. On the Spreadsheet window, select **File**, then **Open** from the menus. Navigate to the **Data** folder that is located in your documents folder and select **Flack Validation Proportions.S0**. Once the dataset is loaded, go back to the PASS Home window, load the **Kappa Test for Agreement Between Two Raters** procedure window by clicking on **Correlation**, then **Kappa Test for Rater Agreement**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Power	0.80 0.90
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
K1	0.6
K0	0.4
P (Category Frequencies)	=C1-C5
Alternative Hypothesis (H1)	One-Sided (H1: K1 > K0)

811-12 Kappa Test for Agreement Between Two Raters

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results

Test Type = One-sided Z test.

H0: Kappa \leq K0 vs. H1: Kappa $>$ K0.

Power	Sample Size N	Kappa H0 K0	Kappa H1 K1	Alpha	Beta	- Rating Categories - k Frequencies	
0.80218	93	0.40	0.60	0.05000	0.19782	3	0.50, 0.26, 0.24
0.80143	99	0.40	0.60	0.05000	0.19857	3	0.50, 0.30, 0.20
0.80253	109	0.40	0.60	0.05000	0.19747	3	0.55, 0.30, 0.15
0.80286	120	0.40	0.60	0.05000	0.19714	3	0.60, 0.30, 0.10
0.80259	106	0.40	0.60	0.05000	0.19741	3	0.60, 0.21, 0.19

The sample sizes computed by *PASS* match those in Flack, Afifi, Lachenbruch, and Schouten (1988). Slight differences are due to rounding.

Chapter 815

Inequality Tests for One Coefficient Alpha

Introduction

Coefficient alpha, or *Cronbach's alpha*, is a measure of the reliability of a scale consisting of k parts. The k parts usually often represent k items on a questionnaire or k raters. This module calculates power and sample size for testing whether coefficient alpha, ρ , is different from a given value such as zero.

Technical Details

Feldt et al. (1987) has shown that if $\hat{\rho}$ is the estimated value of coefficient alpha computed from a sample of size N questionnaires with k items, the statistic W is distributed as an F ratio with degrees of freedom $N-1$ and $(k-1)(N-1)$, where

$$W = \frac{1 - \rho_0}{1 - \hat{\rho}}$$

and ρ_0 is the value of ρ assumed by the null hypothesis, H_0 .

Calculating the Power

Using the above definition of W , the power of the significance test of $\rho > \rho_0$ is calculated as follows:

1. Find F_α such that $\text{Prob}\left(F_{1-\alpha, N-1, (k-1)(N-1)}\right) = 1 - \alpha$
2. Compute $\rho_C = \frac{F_\alpha + \rho_0 - 1}{F_\alpha}$

815-2 Inequality Tests for One Coefficient Alpha

3. Compute $W_1 = \frac{1 - \rho_1}{1 - \rho_c}$, where ρ_1 is the value of ρ at which the power is calculated.
4. Compute the power = $1 - \Pr(W_1 > F_{N-1, (k-1)(N-1)})$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be calculated from the values of the other parameters. Under most conditions, you would either select *Power and Beta* or *N*.

Select *N* when you want to determine the sample size needed to achieve a given power and alpha error level.

Select *Power and Beta* when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis when in fact it is true.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

Specify the number of observations in the sample. You may enter a range such as *10 to 100 by 10* or a list of values separated by commas or blanks such as *20 50 100*.

K (Number of Items or Raters)

K is the number of items or raters in the study. Since it is a count, it must be an integer greater than one. You may enter a list of values separated by blanks.

Effect Size

CA0 (Coefficient Alpha|H0)

Specify the value of ρ_0 , the value of coefficient alpha under the null hypothesis. Usually, this value will be zero, but any value between -1 and 1 is valid as long as it is not equal to CA1.

You may enter a list of values separated by blanks such as *0 0.1 0.2*.

CA1 (Coefficient Alpha|H1)

Specify the value of ρ_1 , the value of coefficient alpha at which the power is computed. Usually, this value is positive, but any value between -1 and 1 is valid as long as it is not equal to CA0.

You may enter a list of values separated by blanks such as *0.1 0.2 0.3*.

Test

Alternative Hypothesis

This option specifies whether the alternative hypothesis is one-sided or two-sided. It also specifies the direction of the hypothesis test. The null hypothesis is $H_0: \rho_0 = \rho$. The alternative hypothesis enters into power calculations by specifying the rejection region of the hypothesis test. Its accuracy is critical.

Possible selections are:

- **H1: CA0 <> CA1**

This is the most common selection. It yields the *two-tailed* test. Use this option when you are testing whether values are different, but you do not want to specify beforehand which is larger.

- **H1: CA0 < CA1**

This option yields a *one-tailed* test.

- **H1: CA0 > CA1**

This option also yields a *one-tailed* test.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Finding the Power

Suppose a study is being designed to test whether the coefficient alpha is 0.6 against the two-sided alternative. Find the power when $K = 20$, $\alpha = 0.05$, $CA1 = 0.65$ 0.70 0.75, and $N = 50$ 100 200 300 500 700 1000 and 1400.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Coefficient Alpha** procedure window by clicking on **Correlation**, then **One Coefficient Alpha**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	50 100 200 300 500 700 1000 1400
K (Number of Items or Raters)	20
CA0 (Coefficient Alpha H0)	0.6
CA1 (Coefficient Alpha H1)	0.65 0.70 0.75
Alternative Hypothesis	H1: CA0 <> CA1

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H1: CA0<>CA1

Power	Sample Size (N)	Number of Items (K)	Coefficient Alpha H1 (CA1)	Coefficient Alpha H0 (CA0)	Signif. Level (Alpha)	Beta
0.11084	50	20	0.65000	0.60000	0.05000	0.88916
0.16444	100	20	0.65000	0.60000	0.05000	0.83556
0.27111	200	20	0.65000	0.60000	0.05000	0.72889
0.37314	300	20	0.65000	0.60000	0.05000	0.62686
0.55224	500	20	0.65000	0.60000	0.05000	0.44776
0.69191	700	20	0.65000	0.60000	0.05000	0.30809
.
.
.

Report Definitions

Power is the probability of rejecting a false null hypothesis.

N is the total sample size.

K is the number of items or raters.

CA1 is the value of coefficient alpha at which the power is computed.

CA0 is the value of coefficient alpha under the null hypothesis.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

Beta is the probability of accepting a false null hypothesis. It should be small.

H0 is the null hypothesis that coefficient alpha equals CA0.


H1 is the alternative hypothesis that coefficient alpha does not equal CA0.

Summary Statements

A sample of 50 subjects each responding to 20 items achieves 11% power to detect the difference between the coefficient alpha under the null hypothesis of 0.60000 and the coefficient alpha under the alternative hypothesis of 0.65000 using a two-sided F-test with a significance level of 0.05000.

This report shows the values of each of the parameters, one scenario per row. The values from this table are plotted in the chart below.

Plots Section



This plot shows the relationship between CA1, N, and power.

Example 2 – Finding the Sample Size

Continuing with the last example, find the sample size necessary to achieve a power of 90% with a 0.05 significance level.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Coefficient Alpha** procedure window by clicking on **Correlation**, then **One Coefficient Alpha**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
K (Number of Items or Raters)	20
CA0 (Coefficient Alpha H0)	0.6
CA1 (Coefficient Alpha H1)	0.65 0.70 0.75
Alternative Hypothesis	H1: CA0 <> CA1

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H1: CA0<>CA1						
Power	Sample Size (N)	Number of Items (K)	Coefficient Alpha H1 (CA1)	Coefficient Alpha H0 (CA0)	Signif. Level (Alpha)	Beta
0.90022	1233	20	0.65000	0.60000	0.05000	0.09978
0.90073	265	20	0.70000	0.60000	0.05000	0.09927
0.90261	100	20	0.75000	0.60000	0.05000	0.09739

This report shows the dramatic increase in sample size that is needed to achieve the desired sample power as CA1 gets closer to CA0.

Example 3 – Validation using Bonett

Bonett (2002) page 337 presents a table in which the sample sizes were calculated for several parameter configurations. When $CA_0 = 0$, $CA_1 = 0.50$, $\alpha = 0.10$, $\beta = 0.05$, and $k = 2, 5, 10$, and 100, he finds N to be 93, 59, 52, and 48, respectively.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for One Coefficient Alpha** procedure window by clicking on **Correlation**, then **One Coefficient Alpha**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.95
Alpha	0.1
N (Sample Size)	<i>Ignored since this is the Find setting</i>
K (Number of Items or Raters)	2 5 10 100
CA0 (Coefficient Alpha H0)	0
CA1 (Coefficient Alpha H1)	0.5
Alternative Hypothesis	H1: CA0 <> CA1

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H1: CA0<>CA1						
Power	Sample Size (N)	Number of Items (K)	Coefficient Alpha H1 (CA1)	Coefficient Alpha H0 (CA0)	Signif. Level (Alpha)	Beta
0.95176	93	2	0.50000	0.00000	0.10000	0.04824
0.95253	59	5	0.50000	0.00000	0.10000	0.04747
0.95047	52	10	0.50000	0.00000	0.10000	0.04953
0.95213	48	100	0.50000	0.00000	0.10000	0.04787

The sample sizes match Bonett's results exactly.

815-8 Inequality Tests for One Coefficient Alpha

Chapter 820

Inequality Tests for Two Coefficient Alphas

Introduction

Coefficient alpha, or *Cronbach's alpha*, is a popular measure of the reliability of a scale consisting of k parts. The k parts often represent k items on a questionnaire (scale) or k raters. This module calculates power and sample size for testing whether two coefficient alphas are different when the two samples are either dependent or independent.

Technical Details

Feldt et al. (1999) presents methods for testing one-, or two-, sided hypotheses about two coefficient alphas, which we label ρ_1 and ρ_2 . The results assume that N_1 observations for each of k_1 items are available for one scale and N_2 observations for each of k_2 items are available for another scale. These sets of observations may either be from two independent groups of subjects (independent case) or two sets of observations on each subject (dependent case). In the dependent case, $N_1 = N_2$ and the correlation coefficient between the overall scores of each scale is represented by ϕ . For the independent case $\phi = 0$.

Suppose $\hat{\rho}_1$ and $\hat{\rho}_2$ are the sample estimates of ρ_1 and ρ_2 , respectively. Hypothesis tests are based on the result that the test statistic,

$$W = \left(\frac{1 - \hat{\rho}_2}{1 - \hat{\rho}_1} \right) \left(\frac{1 - \rho_1}{1 - \rho_2} \right) = \hat{\delta} \left(\frac{1 - \rho_1}{1 - \rho_2} \right),$$

is approximately distributed as a central F variable with degrees of freedom ν_1 and ν_2 . The values of ν_1 and ν_2 depend on N_1 , N_2 , k_1 , k_2 , and ϕ .

820-2 Inequality Tests for Two Coefficient Alphas

Also define

$$c_i = (N_i - 1)(k_i - 1), \quad i = 1, 2$$

Independent Case

When the two scales are independent, there are two situations that must be considered separately.

If $c_i > 1000$ and $k_i > 25$, the values of ν_1 and ν_2 are computed using

$$\nu_1 = N_1 - 1$$

$$\nu_2 = N_2 - 1$$

otherwise, they are computed using

$$\nu_1 = \frac{2A^2}{2B - AB - A^2}$$

$$\nu_2 = \frac{2A}{A - 1}$$

where

$$A = \frac{c_1(N_2 - 1)}{(c_1 - 2)(N_2 - 3)}$$

$$B = \frac{(N_1 + 1)(N_2 - 1)^2(c_2 + 2)c_1^2}{(N_2 - 3)(N_2 - 5)(N_1 - 1)(c_1 - 2)(c_1 - 4)c_2}$$

Dependent Case

When the two scales are dependent, it follows that $N_1 = N_2 = N$. There are two situations that must be considered separately.

If $c_i > 1000$ and $k_i > 25$, the values of ν_1 and ν_2 are computed using

$$\nu_1 = \nu_2 = \frac{N - 1 - 7\phi^2}{1 - \phi^2}$$

otherwise, they are computed using

$$\nu_1 = \frac{2M^2}{V(2 - M) - M^2(M - 1)}$$

$$\nu_2 = \frac{2M}{M - 1}$$

where

$$M = A - \frac{2\phi^2}{N - 1}$$

$$V = B - A^2 - \frac{4\phi^2}{N-1}$$

Calculating the Power

Let ρ_{20} be the value of coefficient alpha in the second set under H_0 , ρ_{21} be the value of coefficient alpha in the second set at which the power is calculated, and ρ_1 be the value of coefficient alpha in the first set. The power of the one-sided hypothesis that $H_0: \rho_{20} \leq \rho_1$ versus the alternative that $H_1: \rho_{20} > \rho_1$ is calculated as follows:

1. Find F_α such that $\text{Prob}(F < F_{\alpha, v_1, v_2}) = \alpha$
2. Compute $\delta' = \frac{1}{F_\alpha} \left(\frac{1 - \rho_1}{1 - \rho_{20}} \right)$
3. Compute $W_1 = \left(\frac{1 - \rho_1}{1 - \rho_{21}} \right) \delta'$
4. Compute the power = $1 - \text{Pr}(W_1 > F_{v_1, v_2})$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options of interest for this procedure.

Solve For

Find (Solve For)

This option specifies the parameter to be calculated from the values of the other parameters. Under most conditions, you would either select *Power and Beta* or *NI*.

Select *NI* when you want to determine the sample size needed to achieve a given power and alpha error level.

Select *Power and Beta* when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when you reject the null hypothesis when in fact it is true.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size – Set 1

N1 (Sample Size in Set 1)

N1 is the sample size (number of observations or subjects) in dataset one. Feldt (1999) states that it is ill-advised to use sample sizes less than 30.

You may enter a list of values separated with blanks or a range of values.

K1 (Items/Scale in Set 1)

K1 is the number of items or raters in dataset one. Since it is a count, it must be an integer. K1 must be greater than or equal to two.

You may enter a list of values separated with blanks or a range of values.

Sample Size – Set 1

N2 (Sample Size in Set 2)

N2 is the sample size (number of observations or subjects) in dataset two. Feldt (1999) states that it is ill-advised to use sample sizes less than 30. Note that if phi is non-zero, this value will be forced equal to N1 regardless of what is entered here.

For ease of input, when phi is non-zero, you can enter multiples of N1. For example, all of the following entries are allowed:

N1 2N1 0.5N1

You may enter a list of values separated with blanks or a range of values.

K2 (Items/Scale in Set 2)

K2 is the number of items or raters in dataset two. Since it is a count, it must be an integer greater than or equal to two.

You may enter a list of values separated with blanks or a range of values.

For ease of input, you can enter multiples of K1. For example, all of the following entries are allowed:

K1 2K1 0.5K1

Effect Size – Coefficient Alpha Set 1**CA1 (Coefficient Alpha Set 1)**

Specify the value of ρ_1 , the value of coefficient alpha, for dataset one. Often, this value will be zero, but any value between -1 and 1 is valid as long as it is not equal to CA21.

You may enter a list of values separated with blanks or a range of values.

Effect Size – Coefficient Alpha Set 2**CA20 (Coefficient Alpha Set 2|H0)**

Enter the value of coefficient alpha in dataset two under the null hypothesis. The null hypothesis is H0: CA1=CA20, so often you will set this value equal to CA1. Any value between -1 and 1 (non-inclusive) is valid as long as it is not equal to CA21.

You may enter *CA1* to indicate that you want the value of CA1 copied here. You may also enter a multiple of CA1 such as *1.5CA1*.

You may enter a list of values separated with blanks or a range of values.

CA21 (Coefficient Alpha Set 2|H1)

Enter the value of coefficient alpha in dataset two at which the power is computed. Any value between -1 and 1 (non-inclusive) is valid as long as it is not equal to CA20. The values of CA20 and CA21 should match the direction set by the 'Alternative Hypothesis' option.

You may enter a list of values separated with blanks or a range of values.

Effect Size – Correlation**Phi (Correlation Between Sets)**

This option implicitly specifies the two datasets as being either independent (datasets on different subjects) or dependent (both datasets on the same subjects). Suppose you calculate the average score for each subject for both dataset one and dataset two. This parameter is the correlation between those two averages over all subjects. If the correlation is zero, the two datasets are assumed to be independent. That is, it is assumed that they come from different sets of subjects.

If the correlation is non-zero, the two datasets are assumed to be dependent. That is, it is assumed that both sets of items were measured on the same subjects. Typical values in this case are between 0.2 and 0.7. When the datasets are dependent, it is assumed that $N1=N2$.

Since this is a correlation, the theoretical range is from -1 to 1. Typical values are 0.0 for independent designs and between 0.2 and 0.7 for dependent designs.

Test

Alternative Hypothesis

This option specifies whether the alternative hypothesis is one-sided or two-sided. It also specifies the direction of the hypothesis test. The null hypothesis is $H_0 : \rho_1 = \rho_2$. The alternative hypothesis enters into power calculations by specifying the rejection region of the hypothesis test. Its accuracy is critical.

Possible selections are:

- **H1: CA1 <> CA2**

This is the most common selection. It yields the *two-tailed* test. Use this option when you are testing whether values are different, but you do not want to specify beforehand which is larger.

- **H1: CA1 < CA2**

This option yields a *one-tailed* test.

- **H1: CA1 > CA2**

This option also yields a *one-tailed* test.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Finding the Power

Suppose a study is being designed to compare the coefficient alphas of two scales. The researchers are going to use a two-sided F-test at a significance level of 0.05. Past experience has shown that CA1 is approximately 0.4. The researchers will use different subjects in each dataset. Find the power when $K1 = K2 = 10$, $CA20 = CA1$, $N1 = 50, 100, 150, 200, 250, \text{ and } 300$, $N2 = N1$, and $CA21 = 0.6 \text{ and } 0.7$.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Coefficient Alphas** procedure window by clicking on **Correlation**, then **Two Coefficient Alphas**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

Option**Value****Data Tab**

Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N1 (Sample Size in Set 1).....	50 to 300 by 50
K1 (Items/Scale in Set 1)	10
N2 (Sample Size in Set 2).....	N1
K2 (Items/Scale in Set 2)	K1
CA1 (Coefficient Alpha Set 1)	0.4
CA20 (Coefficient Alpha Set 2 H0)	CA1
CA21 (Coefficient Alpha Set 2 H1)	0.60 0.70
Phi (Correlation Between Sets).....	0
Alternative Hypothesis (H1)	H1: CA1 <> CA2

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results
Numeric Results for Comparing Two Coefficient Alphas
H1: CA1 <> CA2

	Sample Sizes	Number of Items	Coef. Alpha Set 2 H1 (CA21)	Coef. Alpha Set 2 H0 (CA20)	Coef. Alpha Set 1 (CA1)	Signif. Level (Alpha)	Corr. Between Datasets (Phi)	Beta
Power	N1/N2	K1/K2						
0.2642	50/50	10/10	0.6000	0.4000	0.4000	0.0500	0.0000	0.7358
0.4775	100/100	10/10	0.6000	0.4000	0.4000	0.0500	0.0000	0.5225
0.6481	150/150	10/10	0.6000	0.4000	0.4000	0.0500	0.0000	0.3519
0.7725	200/200	10/10	0.6000	0.4000	0.4000	0.0500	0.0000	0.2275
0.8576	250/250	10/10	0.6000	0.4000	0.4000	0.0500	0.0000	0.1424
0.9132	300/300	10/10	0.6000	0.4000	0.4000	0.0500	0.0000	0.0868
0.6253	50/50	10/10	0.7000	0.4000	0.4000	0.0500	0.0000	0.3747
0.9026	100/100	10/10	0.7000	0.4000	0.4000	0.0500	0.0000	0.0974
0.9793	150/150	10/10	0.7000	0.4000	0.4000	0.0500	0.0000	0.0207
0.9961	200/200	10/10	0.7000	0.4000	0.4000	0.0500	0.0000	0.0039
0.9993	250/250	10/10	0.7000	0.4000	0.4000	0.0500	0.0000	0.0007
0.9999	300/300	10/10	0.7000	0.4000	0.4000	0.0500	0.0000	0.0001

Report Definitions

H0, H1 abbreviate the null and alternative hypotheses, respectively.

Power is the probability of rejecting H0 when it is false. It should be close to one.

N1 & N2 are the sample sizes of datasets one and two, respectively.

K1 & K2 are the number of items in datasets one and two, respectively.

CA21 is the coefficient alpha in dataset two at which the power is calculated.

CA20 is the coefficient alpha in dataset two under H0.

CA1 is the coefficient alpha in dataset one.

Phi is the correlation between the average scores of each of the two datasets.

Alpha is the probability of a type-I error: rejecting H0 when it is true.

Beta is the probability of a type-II error: accepting H0 when it is false. Power = 1 - Beta.


820-8 Inequality Tests for Two Coefficient Alphas

Summary Statements

Samples of 10 items on 50 subjects in dataset one and 10 items on 50 subjects in dataset two achieve 26% power to detect the difference between the coefficient alphas in the two datasets. Under the null hypothesis, the coefficient alphas in datasets one and two are 0.4000 and 0.4000, respectively. The power is computed assuming that the coefficient alpha of dataset two is actually 0.6000. The test statistic used is the two-sided F-test. The significance level of the test was 0.0500.

This report shows the values of each of the parameters, one scenario per row. The values from this table are plotted in the chart below.

Plots Section



This plot shows the relationship between CA21, N1, and power.

Example 2 – Finding the Sample Size

Continuing with the last example, find the sample size necessary to achieve a power of 90% at the 0.05 significance level.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Coefficient Alphas** procedure window by clicking on **Correlation**, then **Two Coefficient Alphas**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N1
Power	0.90
Alpha	0.05
N1 (Sample Size in Set 1)	<i>Ignored since this is the Find setting</i>
K1 (Items/Scale in Set 1)	10
N2 (Sample Size in Set 2)	N1
K2 (Items/Scale in Set 2)	K1
CA1 (Coefficient Alpha Set 1)	0.4
CA20 (Coefficient Alpha Set 2 H0)	CA1
CA21 (Coefficient Alpha Set 2 H1)	0.60 0.70
Phi (Correlation Between Sets)	0
Alternative Hypothesis (H1)	H1: CA1 <> CA2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results for Comparing Two Coefficient Alphas H1: CA1 <> CA2

Power	Sample Sizes N1/N2	Number of Items K1/K2	Coef. Alpha Set 2 H1 (CA21)	Coef. Alpha Set 2 H0 (CA20)	Coef. Alpha Set 1 (CA1)	Signif. Level (Alpha)	Corr. Between Datasets (Phi)	Beta
0.9000	286/286	10/10	0.6000	0.4000	0.4000	0.0500	0.0000	0.1000
0.9026	100/100	10/10	0.7000	0.4000	0.4000	0.0500	0.0000	0.0974

This report shows that 286 subjects per dataset are needed when CA21 is 0.60 and 100 subjects per dataset are needed when CA21 is 0.70.

Example 3 – Validation using Feldt

Feldt et al. (1999) presents an example in which $CA1 = 0$, $CA20 = 0$, $CA21 = 0.5$, $\alpha = 0.05$, $\Phi = 0$, $N1 = N2 = 60$, and $k = 5$. They find the power of a one-sided test to be 0.761.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Inequality Tests for Two Coefficient Alphas** procedure window by clicking on **Correlation**, then **Two Coefficient Alphas**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N1 (Sample Size in Set 1)	60
K1 (Items/Scale in Set 1)	5
N2 (Sample Size in Set 2)	N1
K2 (Items/Scale in Set 2)	K1
CA1 (Coefficient Alpha Set 1)	0.0
CA20 (Coefficient Alpha Set 2 H0)	CA1
CA21 (Coefficient Alpha Set 2 H1)	0.5
Phi (Correlation Between Sets)	0
Alternative Hypothesis (H1)	H1: CA1 < CA2

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when H1: CA0<CA1								
Power	Sample Sizes N1/N2	Number of Items K1/K2	Coef. Alpha Set 2 H1 (CA21)	Coef. Alpha Set 2 H0 (CA20)	Coef. Alpha Set 1 (CA1)	Signif. Level (Alpha)	Corr. Between Datasets (Phi)	Beta
0.7655	60/60	5/5	0.5000	0.0000	0.0000	0.0500	0.0000	0.2345

Note that *PASS*'s result is slightly different from Feldt's because *PASS* uses fractional degrees of freedom and Feldt rounds to the closest integer. Although the difference in power is small, allowing fractional degrees of freedom is more accurate.

Chapter 850

Cox Regression

Introduction

Cox proportional hazards regression models the relationship between the hazard function $\lambda(t|X)$ of survival time and k covariates using the following formula

$$\log\left(\frac{\lambda(t|X)}{\lambda_0(t)}\right) = \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_k X_k$$

where $\lambda_0(t)$ is the baseline hazard. Note that the covariates may be discrete or continuous.

This procedure calculates power and sample size for testing the hypothesis that $\beta_1 = 0$ versus the alternative that $\beta_1 = B$. Note that β_1 is the change in log hazards for a one-unit change in X_1 when the rest of the covariates are held constant. The procedure assumes that this hypothesis will be tested using the Wald (or score) statistic

$$z = \frac{\hat{\beta}_1}{\sqrt{\text{Var}(\hat{\beta}_1)}}$$

Power Calculations

Suppose you want to test the null hypothesis that $\beta_1 = 0$ versus the alternative that $\beta_1 = B$.

Hsieh and Lavori (2000) gave a formula relating sample size, α , β , and B when X_1 is normally distributed. The sample size formula is

$$D = \frac{(z_{1-\alpha/2} + z_{1-\beta})^2}{(1 - R^2)\sigma^2 B^2}$$

where D is the number of events, σ^2 is the variance of X_1 , and R^2 is the proportion of variance explained by the multiple regression of X_1 on the remaining covariates. It is interesting to note that the number of censored observations does not enter in to the power calculations. To obtain a formula for the sample size, N , we inflate D by dividing by P , the proportion of subjects that fail.

850-2 Cox Regression

Thus, the formula for N is

$$N = \frac{(z_{1-\alpha/2} + z_{1-\beta})^2}{P(1 - R^2)\sigma^2 B^2}$$

This formula is an extension of an earlier formula for the case of a single, binary covariate derived by Schoenfeld (1983). Thus, it may be used with discrete or continuous covariates.

Assumptions

It is important to note that this formulation assumes that proportional hazards model with k covariates is valid. However, it does not assume exponential survival times.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. Under most situations, you will select either *Power and Beta* for a power analysis or N for sample size determination.

Select N when you want to calculate the sample size needed to achieve a given power and alpha level.

Select *Power and Beta* when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal probabilities of the event of interest when in fact they are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error (alpha). A type-I error occurs when you reject the null hypothesis of equal probabilities when in fact they are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

This option specifies the total number of observations in the sample. You may enter a single value or a list of values.

Note that when the Overall Event Rate is set to 1.0, the sample size becomes the number of events.

P (Overall Event Rate)

Enter one or more values for the event rate. The event rate is the proportion of subjects in which the event of interest occurs during the duration of the study. This is the proportion of non-censored subjects. Since the values entered here are proportions, they must be in the range $0 < P \leq 1$.

Note that when this value is set to 1.0, the sample size is the number of events (deaths).

Effect Size – Hazard Ratio

B (Log Hazard Ratio)

This procedure calculates power or sample size for testing the hypothesis that $\beta_1 = 0$ versus the alternative that $\beta_1 = B$ in a Cox regression. Enter one or more values of B here.

B is the predicted change in log (base e) hazards corresponding to a one unit change in X_1 when the other covariates are held constant. Thus, if you want to detect a hazard ratio of 1.5, enter $\ln(1.5) = 0.4055$. Although any non-zero value may be entered, common values are between -3 and 3.

Effect Size – Covariates (X1 is the Variable of Interest)

R-Squared of X1 with Other X's

This is the R-Squared that is obtained when X_1 is regressed on the other X 's (covariates) in the model. Use this to account for the influence on power and sample size of adding other covariates. Note that the number of additional variables does not matter in this formulation. Only their overall relationship with X_1 through this R-Squared value is used.

Of course, this value is restricted to being greater than or equal to zero and less than one. Use zero when there are no other covariates.

S (Standard Deviation of X1)

Enter an estimate of the standard deviation of X1, the predictor variable of interest. The formulation used here assumes that X1 follows the normal distribution. However, you can obtain approximate results for non-normal variables by putting in the correct value here. For example, if X1 is binary, the standard deviation is given by $\sqrt{p(1-p)}$ where p is the proportion of either of the binary values in the population of X1.

If you don't have an estimate, you can press the SD button to obtain a window that will help you determine a rough estimate of the standard deviation.

Test

Alternative Hypothesis

Specify whether the test is one-sided or two-sided. When a two-sided hypothesis is selected, the value of alpha is halved by **PASS**. Everything else remains the same.

Note that the accepted procedure is to use the Two Sided option unless you can justify using a one-sided test.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Power for Several Sample Sizes

Cox regression will be used to analyze the power of a survival time study. From past experience, the researchers want to evaluate the sample size needs for detecting regression coefficients of 0.2 and 0.3 for the independent variable of interest. The variable has a standard deviation of 1.20. The R -squared of this variable with seven other covariates is 0.18.

The event rate is thought to be 70% over the 3-year duration of the study. The researchers will test their hypothesis using a 5% significance level with a two-sided Wald test. They decide to calculate the power at sample sizes between 5 and 250.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Cox Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Cox Regression (Y is Elapsed Time)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	5 to 250 by 40
P (Overall Event Rate)	0.70
B (Log Hazard Ratio)	0.2 0.3
R-Squared of X1 with Other X's	0.18
S (Standard Deviation of X1)	1.2
Alternative Hypothesis	Two-Sided

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

	Sample Size (N)	Reg. Coef. (B)	S.D. of X1 (SD)	Event Rate (P)	R-Squared X1 vs Other X's (R2)	Two- Sided Alpha	Beta
Power	5	0.2000	1.2000	0.7000	0.1800	0.05000	0.93983
	45	0.2000	1.2000	0.7000	0.1800	0.05000	0.77041
	85	0.2000	1.2000	0.7000	0.1800	0.05000	0.61163
	125	0.2000	1.2000	0.7000	0.1800	0.05000	0.47092
	165	0.2000	1.2000	0.7000	0.1800	0.05000	0.35357
	205	0.2000	1.2000	0.7000	0.1800	0.05000	0.25996
	245	0.2000	1.2000	0.7000	0.1800	0.05000	0.18777
	5	0.3000	1.2000	0.7000	0.1800	0.05000	0.91151
	45	0.3000	1.2000	0.7000	0.1800	0.05000	0.55185
	85	0.3000	1.2000	0.7000	0.1800	0.05000	0.28957
	125	0.3000	1.2000	0.7000	0.1800	0.05000	0.13798
	165	0.3000	1.2000	0.7000	0.1800	0.05000	0.06135
	205	0.3000	1.2000	0.7000	0.1800	0.05000	0.02588
	245	0.3000	1.2000	0.7000	0.1800	0.05000	0.01047

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N is the size of the sample drawn from the population.

B is the size of the regression coefficient to be detected.

SD is the standard deviation of X1.

P is the event rate.

R2 is the R-squared achieved when X1 is regressed on the other covariates.

Alpha is the probability of rejecting a true null hypothesis.


Beta is the probability of accepting a false null hypothesis.

Summary Statements

A Cox regression of the log hazard ratio on a covariate with a standard deviation of 1.2000 based on a sample of 5 observations achieves 6% power at a 0.93983 significance level to detect a regression coefficient equal to 0.2000. The sample size was adjusted since a multiple regression of the variable of interest on the other covariates in the Cox regression is expected to have an R-Squared of 0.1800. The sample size was adjusted for an anticipated event rate of 0.7000.

This report shows the power for each of the scenarios.

Plots Section



Example 2 – Validation using Hsieh

Hsieh and Lavori (2000) present an example which we will use to validate this program. In this example, $B = 1.0$, $S = 0.3126$, $R^2 = 0.1837$, $P = 0.738$, one-sided alpha = 0.05, and power = 0.80. They calculated $N = 107$.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Cox Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Cox Regression (Y is Elapsed Time)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.80
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
P (Overall Event Rate)	0.738
B (Log Hazard Ratio)	1.0
R-Squared of X1 with Other X's	0.1837
S (Standard Deviation of X1)	0.3126
Alternative Hypothesis	One-Sided

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	Sample Size (N)	Reg. Coef. (B)	S.D. of X1 (SD)	Event Rate (P)	R-Squared X1 vs Other X's (R2)	One-Sided Alpha	Beta
0.80321	106	1.000	0.313	0.738	0.184	0.05000	0.19679

Note that *PASS* calculated 106 rather than the 107 calculated by Hsieh and Lavori (2000). The discrepancy is due to the intermediate rounding that they did. To show this, we will run a second example from Hsieh and Lavori in which $R^2 = 0$ and $P = 1.0$. In this case, $N = 64$.

Numeric Results with $R^2 = 0$ and $P = 1.0$

Power	Sample Size (N)	Reg. Coef. (B)	S.D. of X1 (SD)	Event Rate (P)	R-Squared X1 vs Other X's (R2)	One-Sided Alpha	Beta
0.80399	64	1.000	0.313	1.000	0.000	0.05000	0.19601

Note that *PASS* also calculated 64. Hsieh and Lavori obtained the 107 by adjusting this 64 for P first and then for R^2 . *PASS* does both adjustments at once, obtaining the 106. Thus, the difference is due to intermediate rounding.

Example 3 – Validation for Binary X1 using Schoenfeld

Schoenfeld (1983), page 502, presents an example for the case when X1 is binary. In this example, $B = \ln(1.5) = 0.4055$, $S = 0.5$, $R^2 = 0.0$, $P = 0.71$, one-sided alpha = 0.05, and power = 0.80. Schoenfeld calculated $N = 212$.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the *PASS* Home window, load the **Cox Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Cox Regression (Y is Elapsed Time)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

Option	Value
Data Tab	
Find (Solve For)	N
Power	0.80
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
P (Overall Event Rate)	0.71
B (Log Hazard Ratio)	0.4055
R-Squared of X1 with Other X's	0.0

850-8 Cox Regression

Data Tab (continued)

S (Standard Deviation of X1) **0.5**

Alternative Hypothesis **One-Sided**

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	Sample Size (N)	Reg. Coef. (B)	S.D. of X1 (SD)	Event Rate (P)	R-Squared X1 vs Other X's (R2)	One- Sided Alpha	Beta
0.80028	212	0.406	0.500	0.710	0.000	0.05000	0.19972

Note that *PASS* also obtains $N = 212$.

Chapter 855

Linear Regression

Introduction

Linear regression is a commonly used procedure in statistical analysis. One of the main objectives in linear regression analysis is to test hypotheses about the slope (sometimes called the regression coefficient) of the regression equation. This module calculates power and sample size for testing whether the slope is a value other than the value specified by the null hypothesis.

Difference between Linear Regression and Correlation

The correlation coefficient is used when both X and Y are from the normal distribution (in fact, the assumption actually is that X and Y follow a bivariate normal distribution). That is, X is assumed to be a random variable whose distribution is normal. In the linear regression context, no statement is made about the distribution of X . In fact, X is not even a random variable. Instead, it is a set of fixed values such as 10, 20, 30 or -1, 0, 1. Because of this difference in definition, we have included both Linear Regression and Correlation algorithms. They gave different results. This module deals with the Linear Regression (fixed X) case.

Technical Details

Suppose that the dependence of a variable Y on another variable X can be modeled using the simple linear equation

$$Y = A + BX$$

In this equation, A is the Y -intercept, B is the slope, Y is the dependent variable, and X is the independent variable.

The nature of the relationship between Y and X is studied using a sample of N observations. Each observation consists of a data pair: the X value and the Y value. The values of A and B are estimated from these observations. Since the linear equation will not fit the observations exactly, estimated values of A and B must be used. These estimates are found using the method of least squares. Using these estimated values, each data pair may be modeled using the equation

$$Y_i = a + bX_i + e_i$$

Note that a and b are the estimates of the population parameters A and B . The e values represent the discrepancies between the estimated values ($a + bX$) and the actual values Y . They are called the errors or residuals.

If it is assumed that these e values are normally distributed, tests of hypotheses about A and B can be constructed. Specifically, we can employ an F ratio to test the null hypothesis that the slope is B_0 versus the alternative hypothesis that the slope is B (B not B_0). The power function of this F test can be written

855-2 Linear Regression

$$Power = \Pr(F > F_{\alpha})$$

where F_{α} is the critical value based on the central- F distribution with 1 and $N - 2$ degrees of freedom and the significant level α and F is distributed as a non-central F with degrees of freedom 1 and $N - 2$ and non-centrality parameter λ . The value of λ is

$$\lambda = N \left(\frac{SX(B - B_0)}{\sigma} \right)^2$$

where σ^2 is the variance of the e 's and

$$SX = \sqrt{\frac{\sum_{i=1}^N (X_i - \bar{X})^2}{N}}$$

The values of each of the parameters α , β , σ^2 , N , SX , and B can be determined from the others using the above formulation.

Note that the power for a one-sided test may be found by using 2α for α in the above formulation.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. Under most situations, you will select either *Power and Beta* for a power analysis or *N* for sample size determination.

Select *N* when you want to calculate the sample size needed to achieve a given power and alpha level.

Select *Power and Beta* when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error (alpha). A type-I error occurs when you reject the null hypothesis when in fact it is true.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

This is the number of observations in the study.

Test

Alternative Hypothesis

Specify whether the test is one-sided or two-sided. When a two-sided hypothesis is selected, the value of alpha is halved by *PASS*. Everything else remains the same.

Note that the accepted procedure is to use the Two Sided option unless you can justify using a one-sided test.

Effect Size – Slope

B0 (Slope|H0)

This is the value of the slope assumed by the null hypothesis. Often this value is set to zero, but this is not necessary. The alternative hypothesis is that the slope is not equal to this value.

B1 (Slope|H1)

This is the value of the slope at which the power is computed. The hypothesis being tested is that the slope is some value other than B0.

Effect Size – Standard Deviation of X's

SX (Standard Deviation of X's)

This is the standard deviation of the *X* values in the sample. It is not necessarily the standard deviation of *X* in the population. For example, suppose the *X* values are five -1's and five 1's. The computed standard deviation of these values (dividing by *N* rather than *N* - 1) is 1.0.

855-4 Linear Regression

You will often want to compute the power or sample size for a specific set of X values. Instead of computing SX by hand, you can use the keyword XS (short for X 's) followed by the list of X values. For example, the phrase

$XS -1,1$

is translated into a 1.0 (which is the standard deviation of these two values). This calculation assumes that the sample is allocated equally to the two values. Hence, an N of 10 implies that five are assigned to -1 and five to 1.

If you are planning a study involving two random variables, X and Y , that come from a bivariate normal population, you should enter the actual standard deviation of X here.

Effect Size – Residual Variance Calculation

Residual Variance Method

The standard deviation of the residuals is needed for the power and sample size calculations. These residuals are the e_i in the regression model

$$Y_i = a + bX_i + e_i$$

However, their standard deviation is not available until after the study is complete. *PASS* provides three methods for specifying the standard deviation of the residuals: *Std Deviation of Y*, *Correlation*, and specifying it directly.

SY (Standard Deviation of Y)

Enter an estimate of the standard deviation of Y . This standard deviation ignores X . An estimate of this value must be found from previous studies, pilot studies, or using your best guess. This option is used when 'Residual Variance Method' is set to 'SY'.

When this value is used, the standard deviation of the residuals is computed using the relationship

$$\sigma = \sqrt{\sigma_Y^2 - B^2 SX^2}$$

This value must be greater than zero.

R (Correlation)

Enter an estimate of the correlation between Y and the X values. An estimate of this correlation must be found from previous studies, pilot studies, or using your best guess. This value must be greater than zero and less than one—negative values are allowed. This option is used when 'Residual Variance Method' is set to 'R'.

When this method is used, the standard deviation of the residuals is computed using the relationship

$$\sigma = B(SX)\sqrt{1/R^2 - 1}$$

where R is the correlation.

S (Standard Deviation of Residuals)

Enter an estimate for the value of the standard deviation of the residuals. This option is used when 'Residual Variance Method' is set to 'S'.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating the Power

Suppose a power analysis must be conducted for a linear regression study that will test the relationship between two variables, Y and X . The test will look at the power using two significance levels, 0.01 and 0.05 and several sample sizes between 5 and 85. Based on previous studies, the standard deviation of Y will be assumed to be 1.0. The standard deviation of the X 's in the sample will also be assumed as 1.0. The experimenter decides that unless the slope is at least 0.5, the relationship between X and Y is too weak to be considered.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Linear Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Linear Regression (Y is Continuous)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.01, 0.05
N (Sample Size)	5 to 85 by 10
Alternative Hypothesis	Two-Sided
B0 (Slope H0).....	0.0
B1 (Slope H1).....	0.5
SX (Standard Deviation of X's)	1
Residual Variance Method.....	SY (Std. Dev. of Y)
SY (Standard Deviation of Y).....	1

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Testing of $B=B_0$ where $B_0 = 0.00$

Power	Sample Size (N)	Slope (B)	Standard Deviation of X (SX)	Standard Deviation of Y (SY)	Standard Deviation of Residuals (S)	Alpha	Beta
0.03533	5	0.50	1.00	1.00	0.87	0.01000	0.96467
0.15194	5	0.50	1.00	1.00	0.87	0.05000	0.84806
0.26836	15	0.50	1.00	1.00	0.87	0.01000	0.73164
0.54369	15	0.50	1.00	1.00	0.87	0.05000	0.45631
0.54097	25	0.50	1.00	1.00	0.87	0.01000	0.45903
0.78944	25	0.50	1.00	1.00	0.87	0.05000	0.21056
0.74759	35	0.50	1.00	1.00	0.87	0.01000	0.25241
0.91225	35	0.50	1.00	1.00	0.87	0.05000	0.08775
0.87415	45	0.50	1.00	1.00	0.87	0.01000	0.12585
0.96601	45	0.50	1.00	1.00	0.87	0.05000	0.03399
0.94183	55	0.50	1.00	1.00	0.87	0.01000	0.05817
0.98755	55	0.50	1.00	1.00	0.87	0.05000	0.01245
0.97470	65	0.50	1.00	1.00	0.87	0.01000	0.02530
0.99564	65	0.50	1.00	1.00	0.87	0.05000	0.00436
0.98953	75	0.50	1.00	1.00	0.87	0.01000	0.01047
0.99853	75	0.50	1.00	1.00	0.87	0.05000	0.00147
0.99585	85	0.50	1.00	1.00	0.87	0.01000	0.00415
0.99952	85	0.50	1.00	1.00	0.87	0.05000	0.00048

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N is the size of the sample drawn from the population. To conserve resources, it should be small.

B_0 is the slope under the null hypothesis.

B is the slope at which the power is calculated.

SX is the standard deviation of the X values.

SY is the standard deviation of Y.

Alpha is the probability of rejecting a true null hypothesis. It should be small.


Beta is the probability of accepting a false null hypothesis. It should be small.

Summary Statements

A sample size of 5 achieves 4% power to detect a change in slope from 0.00 under the null hypothesis to 0.50 under the alternative hypothesis when the standard deviation of the X's is 1.00, the standard deviation of Y is 1.00, and the two-sided significance level is 0.01000.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the power versus the sample size for the two values of alpha.

Example 2 – Validation using Neter, Wasserman, and Kutner

Neter, Wasserman, and Kutner (1983) pages 71 and 72 present a power analysis when $N = 10$, $Slope = 0.25$, $\alpha = 0.05$, $SX = \sqrt{(3400 / 10)} = 18.439$, and

$SY = \sqrt{10 + (0.25)^2 (3400 / 10)} = 5.59015$. They found the power to be approximately 0.97.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Linear Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Linear Regression (Y is Continuous)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	10
Alternative Hypothesis	Two-Sided
B0 (Slope H0).....	0.00
B1 (Slope H1).....	0.25

855-8 Linear Regression

Data Tab (continued)

SX (Standard Deviation of X's) **18.439**
Residual Variance Method **SY (Std. Dev. of Y)**
SY (Standard Deviation of Y) **5.59015**

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Testing of $B=B_0$ where $B_0 = 0.00$

Power	Sample Size (N)	Slope (B)	Standard Deviation of X (SX)	Standard Deviation of Y (SY)	Standard Deviation of Residuals (S)	Alpha	Beta
0.97975	10	0.25	18.44	5.59	3.16	0.05000	0.02025

The power of 0.97975 matches their result to two decimals. Note that they used interpolation from a table to obtain their answer.

Chapter 856

Confidence Intervals for Linear Regression Slope

Introduction

This routine calculates the sample size necessary to achieve a specified distance from the slope to the confidence limit at a stated confidence level for a confidence interval about the slope in simple linear regression.

Caution: This procedure assumes that the slope and standard deviation/correlation estimates of the future sample will be the same as the slope and standard deviation/correlation estimates that are specified. If the slope and standard deviation/correlation estimates are different from those specified when running this procedure, the interval width may be narrower or wider than specified.

Technical Details

For a single slope in simple linear regression analysis, a two-sided, $100(1 - \alpha)\%$ confidence interval is calculated by

$$b_1 \pm t_{1-\alpha/2, n-2} s_{b_1}$$

where b_1 is the calculated slope and s_{b_1} is the estimated standard deviation of b_1 , or

$$s_{b_1} = \frac{s}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2}}$$

where

$$s = \sqrt{\frac{\sum_{i=1}^n (Y_i - \hat{Y})^2}{n - 2}}$$

The value s^2 is often obtained from regression tables as MSE.

856-2 Confidence Intervals for Linear Regression Slope

A one-sided $100(1 - \alpha)\%$ upper confidence limit is calculated by

$$b_1 + t_{1-\alpha, n-2} s_{b_1}$$

Similarly, the one-sided $100(1 - \alpha)\%$ lower confidence limit is

$$b_1 - t_{1-\alpha, n-2} s_{b_1}$$

Each confidence interval is calculated using an estimate of the slope plus and/or minus a quantity that represents the distance from the mean to the edge of the interval. For two-sided confidence intervals, this distance is sometimes called the precision, margin of error, or half-width. We will label this distance, D .

The basic equation for determining sample size when D has been specified is

$$D = t_{1-\alpha/2, n-2} s_{b_1}$$

This equation can be solved for any of the unknown quantities in terms of the others. The value $\alpha/2$ is replaced by α when a one-sided interval is used.

In this procedure, the slope and the standard deviation of the X 's are entered as input, where

$$b_1 = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sum_{i=1}^n (X_i - \bar{X})^2}$$

and

$$s_X = \sqrt{\frac{\sum_{i=1}^n (X_i - \bar{X})^2}{n-1}}$$

One of three different additional inputs can be used to calculate s .

1. Standard deviation of the Y 's:

$$s_Y = \sqrt{\frac{\sum_{i=1}^n (Y_i - \bar{Y})^2}{n-1}}$$

In this case s is generated using

$$s = \sqrt{\left(s_Y^2 - b_1^2 s_X^2\right) \frac{n-1}{n-2}}$$

2. Correlation:

$$r = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2 \sum_{i=1}^n (Y_i - \bar{Y})^2}}$$

In this case s is generated using

$$s = b_1 s_X \sqrt{\left(\frac{1}{r^2} - 1\right) \frac{n-1}{n-2}}$$

3. Directly using s :

$$s = \sqrt{\frac{\sum_{i=1}^n (Y_i - \hat{Y})^2}{n-2}}$$

which is often obtained from regression tables as the square root of MSE.

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population slope is $1 - \alpha$.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters.

Confidence

Confidence Level

The confidence level, $1 - \alpha$, has the following interpretation. If thousands of samples of n items are drawn from a population using simple random sampling and a confidence interval is calculated for each sample, the proportion of those intervals that will include the true population slope is $1 - \alpha$.

Often, the values 0.95 or 0.99 are used. You can enter single values or a range of values such as 0.90, 0.95 or 0.90 to 0.99 by 0.01.

856-4 Confidence Intervals for Linear Regression Slope

Sample Size

N (Sample Size)

Enter one or more values for the sample size. This is the number of individuals selected at random from the population to be in the study.

N must be greater than or equal to 3.

You can enter a single value or a range of values.

One-Sided or Two-Sided Interval

Interval Type

Specify whether the interval to be used will be a two-sided confidence interval, an interval that has only an upper limit, or an interval that has only a lower limit.

Precision

Distance from Slope to Limit(s)

This is the distance from the confidence limit(s) to the sample slope. For two-sided intervals, it is also known as the precision, half-width, or margin of error.

You can enter a single value or a list of values. The value(s) must be greater than zero.

Regression – Slope

B (Slope)

Enter the magnitude of the sample slope.

In the linear equation, $Y=A+BX$, B represents the slope of the line relating the dependent variable Y to the independent variable X.

Caution: The sample size estimates for this procedure assume that the slope that is achieved when the confidence interval is produced is the same as the slope entered here. If the sample slope is different from the one specified here, the width of the interval may be narrower or wider than specified.

You can enter a single value or a range of values such as 1,2,3.

This value must be greater than zero.

Regression – Standard Deviation of X's

SX (Standard Deviation of X's)

This is the standard deviation of the X values *in the sample*.

Caution: The sample size estimates for this procedure assume that the SX that is achieved when the confidence interval is produced is the same as the SX entered here. If the sample SX is different from the one specified here, the width of the interval may be narrower or wider than specified.

It is assumed that the standard deviation formula used in calculating SX uses the divisor N-1.

The value must be greater than zero.

Regression – Residual Variance Calculation

Residual Variance Method

Three methods can be used to calculate the standard deviation of the residuals. This option specifies which method should be used.

- **SY**

This option specifies that the formula

$$S = \text{SQRT}[SY^2 - (B \cdot SX)^2]$$

should be used to calculate the standard deviation of the residuals. Notice that this formula requires only SY, SX, and B (the correlation is not used).

- **R**

This option specifies that the formula

$$S = B(SX) \text{SQRT}[1/R^2 - 1]$$

should be used to calculate the standard deviation of the residuals. Notice that this formula requires only SX, B, and R (the value of SY is not used).

- **S**

This option specifies the value of S directly. The values of SY and R are not used.

SY (Standard Deviation of Y)

This is the standard deviation of the Y values *in the sample*.

This value is only used when the Residual Variance Method is set to 'SY (Std. Dev. of Y)'.

Caution: The sample size estimates for this procedure assume that the SY that is achieved when the confidence interval is produced is the same as the SY entered here. If the sample SY is different from the one specified here, the width of the interval may be narrower or wider than specified.

It is assumed that the standard deviation formula used in calculating SY uses the divisor N-1.

The value must be greater than zero.

856-6 Confidence Intervals for Linear Regression Slope

R (Correlation)

This is an estimate of the correlation between Y and X.

This value is only used when the Residual Variance Method is set to 'R (Correlation)'.

Caution: The sample size estimates for this procedure assume that the correlation that is achieved when the confidence interval is produced is the same as the correlation entered here. If the sample correlation is different from the one specified here, the width of the interval may be narrower or wider than specified.

Range: $0 < R < 1$

Note that R cannot be 0 or negative.

S (Standard Deviation of Residuals)

This is the standard deviation of the residuals from the regression of Y on X.

This value is only used when the Residual Variance Method is set to 'S (Std. Dev. of Residuals)'.

Caution: The sample size estimates for this procedure assume that the S that is achieved when the confidence interval is produced is the same as the S entered here. If the sample S is different from the one specified here, the width of the interval may be narrower or wider than specified.

It is assumed that the standard deviation formula used in calculating S is $\sqrt{\text{MSE}}$, which uses the divisor $N-2$.

The value (or values) must be greater than zero.

Iterations Tab

This tab sets an option used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Calculating Sample Size

Suppose a study is planned in which the researcher wishes to construct a two-sided 95% confidence interval for the slope such that the distance from the slope to the limits is no more than 1 unit. The confidence level is set at 0.95, but 0.99 is included for comparative purposes. The estimated slope is 1.7 and the standard deviation of the X's is 11.2. The standard deviation of the residuals estimate, based on the MSE from a similar study, is 48.6. Instead of examining only the interval width of 1, a series of widths from 0.5 to 1.5 will also be considered.

The goal is to determine the necessary sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for Linear Regression Slope** procedure window by clicking on **Confidence Intervals**, then **Correlation and Regression**, then **Linear Regression Slope**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95 0.99
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Distance from Slope to Limits	0.5 to 1.5 by 0.1
B (Slope)	1.7
SX (Standard Deviation of X's)	11.2
Residual Variance Method	S (Std. Dev. of Residuals)
S (Standard Deviation of Residuals)	48.6

856-8 Confidence Intervals for Linear Regression Slope

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results for Two-Sided Confidence Intervals for the Slope in Simple Linear Regression

Confidence Level	Sample Size (N)	Target Distance from Slope to Limits	Actual Distance from Slope to Limits	Slope (B)	Standard Deviation of X (SX)	Standard Deviation of Residuals (S)
0.950	293	0.500	0.500	1.70	11.20	48.60
0.950	205	0.600	0.599	1.70	11.20	48.60
0.950	152	0.700	0.698	1.70	11.20	48.60
0.950	117	0.800	0.798	1.70	11.20	48.60
0.950	93	0.900	0.899	1.70	11.20	48.60
0.950	76	1.000	0.998	1.70	11.20	48.60
0.950	64	1.100	1.093	1.70	11.20	48.60
0.950	54	1.200	1.196	1.70	11.20	48.60
0.950	47	1.300	1.289	1.70	11.20	48.60
0.950	41	1.400	1.388	1.70	11.20	48.60
0.950	36	1.500	1.491	1.70	11.20	48.60
0.990	505	0.500	0.500	1.70	11.20	48.60
0.990	352	0.600	0.600	1.70	11.20	48.60
0.990	260	0.700	0.700	1.70	11.20	48.60
0.990	201	0.800	0.798	1.70	11.20	48.60
0.990	160	0.900	0.897	1.70	11.20	48.60
0.990	130	1.000	0.999	1.70	11.20	48.60
0.990	109	1.100	1.095	1.70	11.20	48.60
0.990	92	1.200	1.197	1.70	11.20	48.60
0.990	79	1.300	1.298	1.70	11.20	48.60
0.990	69	1.400	1.395	1.70	11.20	48.60
0.990	61	1.500	1.491	1.70	11.20	48.60

References

Ostle, B. and Malone, L.C. 1988. Statistics in Research. Iowa State University Press. Ames, Iowa.

Report Definitions

Confidence level is the proportion of confidence intervals (constructed with this same confidence level, sample size, etc.) that would contain the true slope.

N is the size of the sample drawn from the population.

Distance from Slope to Limit is the distance from the confidence limit(s) to the sample slope. For two-sided intervals, it is also known as the precision, half-width, or margin of error.

Target Distance from Slope to Limit is the value of the distance that is entered into the procedure.

Actual Distance from Slope to Limit is the value of the distance that is obtained from the procedure.

B is the sample slope.

SX is the sample standard deviation of the X values.


S is the standard deviation of the sample residuals.

Summary Statements

A sample size of 293 produces a two-sided 95% confidence interval with a distance from the sample slope to the limits that is equal to 0.500 when the sample slope is 1.70, the standard deviation of the X's is 11.20, and the standard deviation of the residuals is 48.60.

This report shows the calculated sample size for each of the scenarios.

Plots Section



This plot shows the sample size versus the distance from the sample slope to the limits for the two confidence levels.

Example 2 – Validation using Ostle and Malone

Ostle and Malone (1988) page 234 give an example of a calculation for a confidence interval for the slope when the confidence level is 95%, the slope is 7.478, the standard deviation of the X 's is 3.8944, the standard deviation of the residuals is 5.369792, and the distance from the slope to the limits is 0.87608. The sample size is 13.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Confidence Intervals for Linear Regression Slope** procedure window by clicking on **Confidence Intervals**, then **Correlation and Regression**, then **Linear Regression Slope**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N (Sample Size)
Confidence Level	0.95
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Interval Type	Two-Sided
Distance from Slope to Limits	0.87608
B (Slope)	7.478
SX (Standard Deviation of X 's)	3.8944
Residual Variance Method	S (Std. Dev. of Residuals)
S (Standard Deviation of Residuals)	5.369792

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Confidence Level	Sample Size (N)	Target Distance from Slope to Limits	Actual Distance from Slope to Limits	Slope (B)	Standard Deviation of X (SX)	Standard Deviation of Residuals (S)
0.95000	13	0.87608	0.87608	7.48	3.89	5.37

PASS also calculated the sample size to be 13.


Chapter 860

Logistic Regression

Introduction

Logistic regression expresses the relationship between a binary response variable and one or more independent variables called *covariates*. A covariate can be discrete or continuous.

Consider a study of death from disease at various ages. This can be put in a logistic regression format as follows. Let a binary response variable Y be one if death has occurred and zero if not. Let X be the individual's age. Suppose a large group of various ages is followed for ten years and then both Y and X are recorded for each person. In order to study the pattern of death versus age, the age values are grouped into intervals and the proportions that have died in each age group are calculated. The results are displayed in the following plot.



As you would expect, as age increases, the proportion dying of disease increases. However, since the proportion dying is bounded below by zero and above by one, the relationship is approximated by an “S” shaped curve. Although a straight-line might be used to summarize the relationship between ages 40 and 60, it certainly could not be used for the young or the elderly.

Under the logistic model, the proportion dying, P , at a given age can be calculated using the formula


$$P = \frac{e^{\beta_0 + \beta_1 X}}{1 + e^{\beta_0 + \beta_1 X}}$$

This formula can be rearranged so that it is linear in X as follows

$$\text{Log}\left(\frac{P}{1-P}\right) = \beta_0 + \beta_1 X$$

860-2 Logistic Regression

Note that the left side is the logarithm of the odds of death versus non-death and the right side is a linear equation for X . This is sometimes called the *logit* transformation of P . When the scale of the vertical axis of the plot is modified using the logit transformation, the following straight-line plot results.



In the logistic regression model, the influence of X on Y is measured by the value of the slope of X which we have called β_1 . The hypothesis that $\beta_1 = 0$ versus the alternative that $\beta_1 = B \neq 0$ is of interest since if $\beta_1 = 0$, X is not related to Y .

Under the alternative hypothesis that $\beta_1 = B$, the logistic model becomes

$$\log\left(\frac{P_1}{1 - P_1}\right) = \beta_0 + BX$$

Under the null hypothesis, this reduces to

$$\log\left(\frac{P_0}{1 - P_0}\right) = \beta_0$$

To test whether the slope is zero at a given value of X , the difference between these two quantities is formed giving

$$\beta_0 + BX - \beta_0 = \log\left(\frac{P_1}{1 - P_1}\right) - \log\left(\frac{P_0}{1 - P_0}\right)$$

which reduces to

$$\begin{aligned} BX &= \log\left(\frac{P_1}{1 - P_1}\right) - \log\left(\frac{P_0}{1 - P_0}\right) \\ &= \log\left(\frac{P_1 / (1 - P_1)}{P_0 / (1 - P_0)}\right) \\ &= \log(OR) \end{aligned}$$

where OR is odds ratio of P_1 and P_0 . This relationship may be solved for OR giving

$$OR = e^{BX}$$

This shows that the odds ratio of P_1 and P_0 is directly related to the slope of the logistic regression equation. It also shows that the value of the odds ratio depends on the value of X . For a given value of X , testing that B is zero is equivalent to testing OR is one. Since OR is commonly used and well understood, it is used as a measure of effect size in power analysis and sample size calculations.

Power Calculations

Suppose you want to test the null hypothesis that $\beta_1 = 0$ versus the alternative that $\beta_1 = B$. Hsieh, Block, and Larsen (1998) have presented formulae relating sample size, α , power, and B for two situations: when X_1 is normally distributed and when X_1 is binomially distributed.

When X_1 is normally distributed, the sample size formula is

$$N = \frac{(z_{1-\alpha/2} + z_{1-\beta})^2}{P^*(1-P^*)B^2}$$

where P^* is the event rate (probability that $Y = 1$) at the mean of X_1 . Note that B is defined in terms of an increase of one standard deviation of X_1 above the mean.

When X_1 is binomially distributed and $X_1 = 0$ or 1 , the sample size formula is

$$N = \frac{\left(z_{1-\alpha/2} \sqrt{\frac{\bar{P}(1-\bar{P})}{R}} + z_{1-\beta} \sqrt{P_0(1-P_0) + \frac{P_1(1-P_1)(1-R)}{R}} \right)^2}{(P_0 - P_1)^2(1-R)}$$

where P_0 is the event rate at $X_1 = 0$ and P_1 is the event rate at $X_1 = 1$, R is the proportion of the sample with $X_1 = 1$, and \bar{P} is the overall event rate given by

$$\bar{P} = (1-R)P_0 + R(P_1).$$

Multiple Logistic Regression

The multiple logistic regression model relates the probability distribution of Y to two or more covariates X_1, X_2, \dots, X_k by the formula

$$\log\left(\frac{P}{1-P}\right) = \beta_0 + \beta_1 X_1 + \dots + \beta_k X_k$$

where P is the probability that $Y = 1$ given the values of the covariates. It is a simple extension of the simple logistic regression model that was just presented. In power analysis and sample size work, attention is placed on a single covariate while the influence of the other covariates is statistically removed by placing them at their mean values.

860-4 Logistic Regression

When there are multiple covariates, the following adjustment was given by Hsieh (1998) to give the total sample size, N_m

$$N_m = \frac{N}{1 - \rho^2}$$

where ρ is the multiple correlation coefficient between X_1 (the variable of interest) and the remaining covariates. Notice that the number of extra covariates does not matter in this approximation.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *PI*, *N*, *Alpha*, and *Power and Beta*. Under most situations, you will select either *Power and Beta* for a power analysis or *N* for sample size determination.

Select *N* when you want to calculate the sample size needed to achieve a given power and alpha level.

Select *Power and Beta* when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. A type-II error occurs when you fail to reject the null hypothesis of equal probabilities of the event of interest when in fact they are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error (alpha). A type-I error occurs when you reject the null hypothesis of equal probabilities when in fact they are equal.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

This option specifies the total number of observations in the sample. You may enter a single value or a list of values.

Test

Alternative Hypothesis

Specify whether the test is one-sided or two-sided. When a two-sided hypothesis is selected, the value of alpha is halved by *PASS*. Everything else remains the same.

Note that the accepted procedure is to use the Two Sided option unless you can justify using a one-sided test.

Effect Size – Baseline Probability

P0 (Baseline Probability that Y=1)

This option specifies one or more P_0 values. The interpretation of P_0 depends on whether X_1 is binary or continuous.

Binomial Covariate

When X_1 is binary, P_0 is the probability that $Y = 1$ when $X_1 = 0$. All other covariates are assumed to be equal to their mean values. In this case, the logistic equation reduces to

$$\log\left(\frac{P_0}{1 - P_0}\right) = \beta_0$$

so that

$$P_0 = \frac{e^{\beta_0}}{1 + e^{\beta_0}}$$

Normal Covariate

When X_1 is normally distributed, P_0 is the probability that $Y = 1$ when $X_1 = \mu_{X_1}$, where μ_{X_1} is the mean of X_1 . That is, P_0 is the baseline probability that $Y = 1$ when X_1 is ignored. All other

860-6 Logistic Regression

covariates are assumed to be equal to their mean values. In this case, the logistic equation reduces to

$$\log\left(\frac{P_0}{1 - P_0}\right) = \beta_0 + \beta_1\mu_{X_1}$$

so that

$$P_0 = \frac{e^{\beta_0 + \beta_1\mu_{X_1}}}{1 + e^{\beta_0 + \beta_1\mu_{X_1}}}$$

Effect Size – Alternative Probability

Use P1 or Odds Ratio

This option specifies the whether to specify PI directly or to specify it by specifying the odds ratio. Since the relationship between the odds ration, PI , and $P0$ is given by

$$OR = \frac{P_1 / (1 - P_1)}{P_0 / (1 - P_0)}$$

specifying OR and $P0$ implicitly specifies PI .

This options lets you specify whether you want to state the alternative hypothesis in terms of PI or the odds ratio.

P1 (Alternative Probability that Y=1)

This option specifies the effect size to be detected by specifying P_1 . As was shown earlier, the slope of the logistic regression can be expressed in terms of P_0 and P_1 . Hence, by specifying P_1 , you are also specifying the slope.

This option is only used when the User P1 or Odds Ratio option is set to PI . Its interpretation depends on whether X_1 is binomial or normal.

Binomial Covariate

When X_1 is binary, PI is the probability that $Y = 1$ when $X_1 = 1$. All other covariates are assumed to be equal to their mean values. In this case, the logistic equation reduces to

$$\log\left(\frac{P_1}{1 - P_1}\right) = \beta_0 + \beta_1$$

since $X_1 = 1$.

Normal Covariate

When X_1 is normally distributed, PI is the probability that $Y = 1$ when $X_1 = \mu_{x_1} + \sigma_{x_1}$. That is, when x_1 is one standard deviation above the mean. All other covariates are assumed to be equal to their mean values. In this case, the logistic equation reduces to

$$\log\left(\frac{P_1}{1 - P_1}\right) = \beta_0 + \beta_1 x_1$$

so that

$$P_1 = \frac{e^{\beta_0 + \beta_1 x_1}}{1 + e^{\beta_0 + \beta_1 x_1}}$$

Odds Ratio (Odds1/Odds0)

This option specifies the odds ratio to be detected by the study. As was shown earlier, the slope of the logistic regression can be expressed in terms of P_0 and the odds ratio. Hence, by specifying OR , you are also specifying the slope. Using the formula

$$P_1 = \frac{OR(P_0)}{1 - P_0 + OR(P_0)}$$

specifying OR and P_0 implicitly specifies P_1 .

This option is only used when the User P1 or Odds Ratio option is set to *Odds Ratio*. Its interpretation depends on whether X_1 is binomial or normal.

Binomial Covariate

When X_1 is binary, this option gives the odds ratio of P_1 and P_0 . All other covariates are assumed to be equal to their mean values. In this case, the logistic equation reduces to

$$\log\left(\frac{P_1}{1 - P_1}\right) = \beta_0 + \beta_1$$

since $X_1 = 1$.

This odds ratio compares the odds of obtaining $Y = 1$ when $X_1 = 1$ to the odds of obtaining $Y = 1$ when $X_1 = 0$.

Normal Covariate

When X_1 is normally distributed, this option gives the odds ratio of P_1 and P_0 , where P_1 is the probability that $Y = 1$ when $X_1 = x_1$, where x_1 is a value other than μ_{x_1} . All other covariates are assumed to be equal to their mean values. In this case, the logistic equation reduces to

$$\log\left(\frac{P_1}{1 - P_1}\right) = \beta_0 + \beta_1 x_1$$

so that

$$P_1 = \frac{e^{\beta_0 + \beta_1 x_1}}{1 + e^{\beta_0 + \beta_1 x_1}}$$

This odds ratio compares the odds of obtaining $Y = 1$ when $X_1 = x_1$ to the odds of obtaining Y when $X_1 = \mu_{x_1}$.

Effect Size – Covariates (X_1 is the Variable of Interest)

R-Squared of X_1 with Other X 's

This is the R-Squared that is obtained when X_1 is regressed on the other X 's (covariates) in the model. Use this to study the influence on power and sample size of adding other covariates. Note that the number of additional variables does not matter in this formulation. Only their overall relationship with X_1 through this R-Squared value is used.

Of course, this value is restricted to being greater than or equal to zero and less than one. Use zero when there are no other covariates.

X_1 (Independent Variable of Interest)

This option specifies whether the covariate is binary (binomial) or continuous (normal). This is a very important distinction since the sample size required for a particular power level is much larger for a binary covariate than for a continuous covariate.

This selection also changes the meaning of $P0$ and $P1$.

Percent of N with $X_1 = 1$

When X_1 is binary, this option specifies the proportion, R , of the sample in which $X_1 = 1$. Note that the value is specified as a percentage.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Power for a Continuous Covariate

A study is to be undertaken to study the relationship between post-traumatic stress disorder and heart rate after viewing video tapes containing violent sequences. Heart rate is assumed to be normally distributed. The event rate is thought to be 7% among soldiers. The researchers want a sample size large enough to detect an odds ratios of 1.5 or 2.0 with 90% power at the 0.05 significance level with a two-sided test. They decide to calculate the power at level sample sizes between 20 and 1200.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logistic Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Logistic Regression (Y is Binary)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	20 50 100 200 300 500 700 1000 1200
Alternative Hypothesis	Two-Sided
P0 (Baseline Probability that Y=1)	0.07
Use P1 or Odds Ratio	Odds Ratio
Odds Ratio (Odds1/Odds0)	1.5 2.0
R-Squared of X1 with Other X's	0
X1 (Independent Variable of Interest)	Continuous (Normal)
Percent of N with X1=1	<i>Ignored since X is continuous</i>

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	N	P0	P1	Odds Ratio	R Squared	Alpha	Beta
0.06716	20	0.070	0.101	1.500	0.000	0.05000	0.93284
0.10964	50	0.070	0.101	1.500	0.000	0.05000	0.89036
0.17737	100	0.070	0.101	1.500	0.000	0.05000	0.82263
0.30962	200	0.070	0.101	1.500	0.000	0.05000	0.69038
0.43325	300	0.070	0.101	1.500	0.000	0.05000	0.56675
0.63808	500	0.070	0.101	1.500	0.000	0.05000	0.36192
0.78147	700	0.070	0.101	1.500	0.000	0.05000	0.21853
0.90516	1000	0.070	0.101	1.500	0.000	0.05000	0.09484
0.94779	1200	0.070	0.101	1.500	0.000	0.05000	0.05221
0.12119	20	0.070	0.131	2.000	0.000	0.05000	0.87881
0.23903	50	0.070	0.131	2.000	0.000	0.05000	0.76097
0.42410	100	0.070	0.131	2.000	0.000	0.05000	0.57590
0.70579	200	0.070	0.131	2.000	0.000	0.05000	0.29421

860-10 Logistic Regression

0.86504	300	0.070	0.131	2.000	0.000	0.05000	0.13496
0.97696	500	0.070	0.131	2.000	0.000	0.05000	0.02304
0.99673	700	0.070	0.131	2.000	0.000	0.05000	0.00327
0.99986	1000	0.070	0.131	2.000	0.000	0.05000	0.00014
0.99998	1200	0.070	0.131	2.000	0.000	0.05000	0.00002

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N is the size of the sample drawn from the population.

P0 is the response probability at the mean of the covariate, X.

P1 is the response probability when X is increased to one standard deviation above the mean.

Odds Ratio is the odds ratio when P1 is on top. That is, it is $[P1/(1-P1)]/[P0/(1-P0)]$.

R-Squared is the R2 achieved when X is regressed on the other independent variables in the regression.

Alpha is the probability of rejecting a true null hypothesis.


Beta is the probability of accepting a false null hypothesis.

Summary Statements

A logistic regression of a binary response variable (Y) on a continuous, normally distributed, independent variable (X) with a sample size of 20 observations achieves 7% power at a 0.050 significance level to detect a change in Prob(Y=1) from the value of 0.070 at the mean of X to 0.101 when X is increased to one standard deviation above the mean. This change corresponds to an odds ratio of 1.500.

This report shows the power for each of the scenarios. The report shows that a power of 90% is reached at a sample size of about 300 for an odds ratio of 2.0 and 1000 for an odds ratio of 1.5.

Plot Section



This plot shows the power versus the sample size for the two values of the odds ratio.

Example 2 – Sample Size for a Continuous Covariate

Continuing with the previous study, determine the exact sample size necessary to attain a power of 90%.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logistic Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Logistic Regression (Y is Binary)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Alternative Hypothesis	Two-Sided
P0 (Baseline Probability that Y=1)	0.07
Use P1 or Odds Ratio	Odds Ratio
Odds Ratio (Odds1/Odds0)	1.5 2.0
R-Squared of X1 with Other X's	0
X1 (Independent Variable of Interest)	Continuous (Normal)
Percent of N with X1=1	<i>Ignored since X is continuous</i>

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	N	P0	P1	Odds Ratio	R Squared	Alpha	Beta
0.89978	981	0.070	0.101	1.500	0.000	0.05000	0.10022
0.89920	335	0.070	0.131	2.000	0.000	0.05000	0.10080

This report shows the power for each of the scenarios. The report shows that a power of 90% is achieved at a sample size of 335 for an odds ratio of 2.0 and 981 for an odds ratio of 1.5.

Example 3 – Effect Size for a Continuous Covariate

Continuing the previous study, suppose the researchers can only afford a sample size of 500 individuals. They want to determine if a meaningful odds ratio can be detected with this sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logistic Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Logistic Regression (Y is Binary)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	P1>P0 or Odds Ratio>1
Power	0.90
Alpha	0.05
N (Sample Size)	500
Alternative Hypothesis	Two-Sided
P0 (Baseline Probability that Y=1)	0.07
Use P1 or Odds Ratio	<i>Ignored since this is the Find setting</i>
Odds Ratio (Odds1/Odds0).....	<i>Ignored since this is the Find setting</i>
R-Squared of X1 with Other X's	0
X1 (Independent Variable of Interest)	Continuous (Normal)
Percent of N with X1=1	<i>Ignored since X is continuous</i>

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	N	P0	P1	Odds Ratio	R Squared	Alpha	Beta
0.90000	500	0.070	0.117	1.765	0.000	0.05000	0.10000

This report shows that this experimental design can detect an odds ratio of 1.765. That is, it can detect a shift in the event rate from 0.070 to 0.117.

Example 4 – Sample Size for a Binary Covariate

A study is to be undertaken to study the relationship between post-traumatic stress disorder and gender. The event rate is thought to be 7% among males. The researchers want a sample size large enough to detect an odds ratio of 1.5 with 90% power at the 0.05 significance level with a two-sided test.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logistic Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Logistic Regression (Y is Binary)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	N
Power	0.90
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Alternative Hypothesis	Two-Sided
P0 (Baseline Probability that Y=1)	0.07
Use P1 or Odds Ratio	Odds Ratio
Odds Ratio (Odds1/Odds0)	1.5
R-Squared of X1 with Other X's	0
X1 (Independent Variable of Interest)	Binary (X=0 or 1)
Percent of N with X1=1	50

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	N	Pcnt N X=1	P0	P1	Odds Ratio	R Squared	Alpha	Beta
0.89997	3326	50.000	0.070	0.101	1.500	0.000	0.05000	0.10003

The sample size is estimated at 3326. This should be evenly divided among males and females.

Example 5 – Validation for a Continuous Covariate

Hsieh (1998) page 1628 gives the power as 95% when $N = 317$, $\alpha = 0.05$ (two-sided), $P0 = 0.5$, and the odds ratio is 1.5. The covariate is assumed to be continuous.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logistic Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Logistic Regression (Y is Binary)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example5** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	317
Alternative Hypothesis	Two-Sided
P0 (Baseline Probability that Y=1)	0.50
Use P1 or Odds Ratio	Odds Ratio
Odds Ratio (Odds1/Odds0).....	1.5
R-Squared of X1 with Other X's	0
X1 (Independent Variable of Interest)	Continuous (Normal)

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	N	P0	P1	Odds Ratio	R Squared	Alpha	Beta
0.95049	317	0.500	0.600	1.500	0.000	0.05000	0.04951

PASS calculates a power of 0.95049 which matches Hsieh.

Example 6 – Validation for a Binary Covariate

Hsieh (1998) page 1626 gives the power as 95% when $N = 1282$ (equal sample sizes for both groups), $\alpha = 0.05$ (two-sided), $P0 = 0.4$, and the $P1 = 0.5$. The covariate is assumed to be binary.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Logistic Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Logistic Regression (Y is Binary)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example6** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	1282
Alternative Hypothesis	Two-Sided
P0 (Baseline Probability that Y=1)	0.4
Use P1 or Odds Ratio	P1
P1 (Alternative Probability that Y=1)	0.5
R-Squared of X1 with Other X's	0
X1 (Independent Variable of Interest)	Binary (X=0 or 1)
Percent of N with X1=1	50

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Power	N	Pcnt N X=1	P0	P1	Odds Ratio	R Squared	Alpha	Beta
0.95021	1282	50.000	0.400	0.500	1.500	0.000	0.05000	0.04979

PASS calculates a power of 0.95021 which matches Hsieh.

Chapter 865

Multiple Regression

Introduction

The Linear Regression chapter dealt with the regression of a dependent variable on a single independent variable. This module deals with the general case of multiple regression in which there are two or more independent variables.

Instead of using the correlation coefficient as the index of interest, the square of the correlation coefficient, R^2 , has been adopted, mainly because of its ease of interpretation and its ability to be partitioned among various subgroups of the independent variables. R^2 ranges from zero to one.

When performing a regression analysis, a typical hypothesis involves testing the significance of a subgroup of the independent variables after considering a second, nonoverlapping, group of independent variables. For example, suppose you have five independent variables. One common hypothesis tests whether a certain variable is influential (has a nonzero coefficient in the regression equation) after considering the other four variables. To perform this test, you partition the R^2 from fitting all five variables into the R^2 of the first four variables and the R^2 added by the fifth variable. An F -test is constructed that tests whether this second R^2 value is significantly different from zero. If it is, the fifth variable is significant after adjusting for the other four variables.

Definition of Terms

Cohen (1988) provides a comprehensive discussion of this topic. We strongly suggest that you refer to the examples given there for a thorough understanding of this subject. We will provide a brief introduction, concentrating on defining the terms and statistics that are used.

Consider a sample of N rows of data. Suppose each row consists of the value of a dependent variable, Y , followed by the values of k independent variables, X_1, X_2, \dots, X_k . The multiple regression equation of Y on the X 's is

$$Y_j = \beta_0 + \beta_1 X_{1j} + \beta_2 X_{2j} + \dots + \beta_k X_{kj} + \varepsilon_j$$

The β 's are called the population regression coefficients or beta weights. They represent the true values in the population from which the sample was taken. Their sample estimates, $b_0, b_1, b_2, \dots, b_k$, are calculated from the data using least squares methodology. The estimated regression equation becomes

865-2 Multiple Regression

$$Y_j = b_0 + b_1 X_{1j} + b_2 X_{2j} + \cdots + b_k X_{kj} + e_j$$

Upon removing the residuals, e_j , the predicted values are calculated as

$$\hat{Y}_j = b_0 + b_1 X_{1j} + b_2 X_{2j} + \cdots + b_k X_{kj}$$

R^2 is defined as the proportion of the variation of Y that is explained (accounted for) by the variation in the X 's. The equation for R^2 is

$$R^2 = \frac{\sum_{i=1}^N (\hat{Y}_i - \bar{Y})^2}{\sum_{i=1}^N (Y_i - \bar{Y})^2}$$

or

$$R^2 = 1 - \frac{\sum_{i=1}^N e_i^2}{\sum_{i=1}^N (Y_i - \bar{Y})^2}$$

Notice that the numerator sum of squares in the above equations depend on the X 's through \hat{Y}_i . Therefore, each time you change the X 's included in the regression equation, you will obtain a different value of R^2 . Hence, you need special notation to indicate on which X 's the particular R^2 is based.

Let C and T refer to subgroups of the X 's. Define $R_{T|C}^2 = R_{TC}^2 - R_C^2$ to be the R^2 added when Y is regressed on the variables in set T after the variables in set C . Define R_C^2 to be the R^2 achieved when Y is regressed on those in set C ignoring the variables in set T . Define R_{TC}^2 to be the R^2 achieved when Y is regressed on the variables in both sets C and T .

Using the above notation, you can construct F -tests that will test whether the β 's corresponding to certain sets of X 's are simultaneously zero while controlling for other variables. For example, to test the significance of the X 's in set T while removing the influence of the X 's in set C from experimental error, you would use:

$$F_{u,v} = \frac{(R_{T|C}^2) / u}{(1 - R_C^2 - R_{T|C}^2) / v}$$

where u is the number of variables in T , $v = N - k - 1$, and k is the total number of variables in C and T . Note that this formula includes the effect size, f^2 .

$$f^2 = \left(\frac{R_{T|C}^2}{1 - R_C^2 - R_{T|C}^2} \right)$$

Most significance tests in regression analysis, correlation analysis, analysis of variance, and analysis of covariance may be constructed using these F -ratios.

For example, suppose the seven X 's available are dividing into two sets: set C includes variables X_1, X_3, X_6 and set T includes X_2, X_4, X_5, X_7 . The above F -ratio tests the null hypothesis that $\beta_2 = \beta_4 = \beta_5 = \beta_7 = 0$ after adjusting for X_1, X_3, X_6 .

Calculating the Power

The calculation of the power of a particular test proceeds as follows:

1. Determine the critical value, $F_{u,v,\alpha}$ where u is the numerator degrees of freedom, v is the denominator degrees of freedom, and α is the probability of a type-I error.
2. Calculate the noncentrality parameter λ using the formula:

$$\lambda = N \frac{R_{T|C}^2}{(1 - R_C^2 - R_{T|C}^2)}$$

3. Compute the power as the probability of being greater than $F_{u,v,\alpha}$ on a noncentral- F distribution with noncentrality parameter λ .

Note that the formula for λ is different from that used in **PASS 6.0**. The algorithm used in **PASS 6.0** was based on formula (9.3.1) in Cohen (1988) which gives approximate answers. This version of **PASS** using an algorithm that gives exact answers.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are R^2 , N , $Alpha$, and $Power$ and $Beta$. Under most situations, you will select either $Power$ and $Beta$ or N .

Select N when you want to calculate the sample size needed to achieve a given power and alpha level.

Select $Power$ and $Beta$ when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error (alpha). A type-I error occurs when you reject the null hypothesis when in fact it is true.

Values of alpha must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

This option specifies the value(s) for N , the sample size. Note that this value must be at least one greater than the total number of variables specified in sets C and T .

Effect Size – C: Variables Controlled

These options refer to the independent variables (X 's) to be controlled (adjusted) for (these are the variables in set C). That is, the influence of these variables is statistically removed. These are variables that are included in the regression equation but are not being tested for statistical significance.

Number of Independent Variables Controlled

This option specifies the number of variables in C . This number must be greater than or equal to zero. Note that the total number of variables specified in the two Number options must be less than $N-1$.

R²(C)

This box specifies the R^2 achieved by the variables in set C when they are fit alone in the regression equation. Note that this amount must be between zero and one and that the total of the two R^2 values must be less than one.

Effect Size – T: Variables Tested

These options refer to the set of independent variables (X 's) whose coefficients are to be tested to determine if they all are zero. That is, this is the set of variables to be tested for statistical significance.

Number of Independent Variables Tested

This option specifies the number of variables in T . This is u . This number must be greater than or equal to one. Note that the total number of variables specified in the two Number options must be less than $N - 1$.

$R^2(T|C)$

This box specifies the increase in R^2 due to the variables in set T after including the variables in set C in the regression equation. Note that this amount must be between zero and one and that the total of the two R^2 values must be less than one.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Testing the Addition or Deletion of a Single Variable

This example calculates the power of an F test constructed to test a fifth variable which adds 0.05 to R^2 after considering four other variables whose combined R^2 value is 0.50. Sample sizes from 10 to 150 will be investigated. The significance level is 0.05.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Multiple Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Multiple Regression (Y is Continuous)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting.</i>
Alpha	0.05
N (Sample Size)	10 to 150 by 20
Number of Indep Vars Controlled	4
R2(C).....	0.50
Number of Indep Vars Tested	1
R2(T C).....	0.05

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results							
Power	N	Alpha	Beta	Ind. Variables Tested		Ind. Variables Controlled	
				Cnt	R2	Cnt	R2
0.13037	10	0.05000	0.86963	1	0.05000	4	0.50000
0.41799	30	0.05000	0.58201	1	0.05000	4	0.50000
0.63511	50	0.05000	0.36489	1	0.05000	4	0.50000
0.78431	70	0.05000	0.21569	1	0.05000	4	0.50000
0.87818	90	0.05000	0.12182	1	0.05000	4	0.50000
0.93366	110	0.05000	0.06634	1	0.05000	4	0.50000
0.96493	130	0.05000	0.03507	1	0.05000	4	0.50000
0.98192	150	0.05000	0.01808	1	0.05000	4	0.50000

Report Definitions

Power is the probability of rejecting a false null hypothesis.

N is the number of observations on which the multiple regression is computed.

Alpha is the probability of rejecting a true null hypothesis. It should be small.

Beta is the probability of accepting a false null hypothesis. It should be small.

Cnt refers to the number of independent variables in that category.

R2 is the amount that is added to the overall R-Squared value by these variables.

Ind. Variables Tested are those variables whose regression coefficients are tested against zero.


Ind. Variables Controlled are those variables whose influence is removed from experimental error.

Summary Statements

A sample size of 10 achieves 13% power to detect an R-Squared of 0.05000 attributed to 1 independent variable(s) using an F-Test with a significance level (alpha) of 0.05000. The variables tested are adjusted for an additional 4 independent variable(s) with an R-Squared of 0.50000.

This report shows the values of each of the parameters, one scenario per row. The definitions of each of the columns is given in the Report Definitions section.

Note that in this particular example, a reasonable power of 0.90 is not reached until the sample size is 110.

Plots Section

This plot shows the relationship between sample size and power.

Example 2 – Minimum Detectable R-Squared

Suppose the researcher in Example1 can only afford a sample size of 30. He wants to know the minimum detectable R^2 that can be detected if the power is 80% and 90%.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Multiple Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Multiple Regression (Y is Continuous)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	R-Squared of IV's Tested
Power	0.80 0.90
Alpha	0.05
N (Sample Size)	30
Number of Indep Vars Controlled	4
R2(C).....	0.50
Number of Indep Vars Tested	1
R2(T C).....	<i>Ignored since this is the Find setting</i>

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results				Ind. Variables Tested		Ind. Variables Controlled	
Power	N	Alpha	Beta	Cnt	R2	Cnt	R2
0.90000	30	0.05000	0.10000	1	0.13787	4	0.50000
0.80000	30	0.05000	0.20000	1	0.11066	4	0.50000

This report shows that at 90% power, a sample size of 30 cannot detect an R^2 less than 0.13787.

Example 3 – Sample Size, Many X's

Suppose you have 25 observations on 19 independent variables. You run a regression analysis with all 19 variables and obtain an R^2 value of 0.90. You find that 4 of these variables will result in an R^2 of 0.60, leaving an R^2 increment of 0.30 for the other 15. You decide to test the significance of these 15 variables. The F ratio is

$$\begin{aligned} F_{15,5} &= \frac{(0.90 - 0.60) / 15}{(1.0 - 0.90) / 5} \\ &= 1.0 \end{aligned}$$

Obviously, an F of 1.0 is nonsignificant, so you might be tempted to discard these variables and proceed with the four that are significant. Before doing so, you decided to do a power analysis on this F test.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Multiple Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Multiple Regression (Y is Continuous)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting.</i>
Alpha	0.05
N (Sample Size)	25
Number of Indep Vars Controlled	4
R2(C).....	0.60
Number of Indep Vars Tested	15
R2(T C).....	0.3

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results				Ind. Variables Tested		Ind. Variables Controlled	
Power	N	Alpha	Beta	Cnt	R2	Cnt	R2
0.72022	25	0.05000	0.27978	15	0.30000	4	0.60000

The power is only 0.72022, which is a little low. Hence, you cannot be certain whether these 15 variables are unimportant or if you just do not have a large enough sample.

Example 4 – Validation

Ralph O’Brien, in a private communication to Jerry Hintze, gave the result that when $\alpha = 0.05$, $N = 15$, $K = 2$, and $R\text{-Squared} = 0.6$, the power is 0.96829.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Multiple Regression** procedure window by clicking on **Regression (Y is a Function of X’s)**, then **Multiple Regression (Y is Continuous)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting.</i>
Alpha	0.05
N (Sample Size)	15
Number of Indep Vars Controlled	0
R2(C).....	0.0
Number of Indep Vars Tested	2
R2(T C).....	0.6

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results				Ind. Variables Tested		Ind. Variables Controlled	
Power	N	Alpha	Beta	Cnt	R2	Cnt	R2
0.96829	15	0.05000	0.03171	2	0.60000	6	0.00000

The power of 0.96829 matches O’Brien’s result.

Chapter 870

Poisson Regression

Introduction

Poisson regression is used when the dependent variable is a count. Following the results of Signorini (1991), this procedure calculates power and sample size for testing the hypothesis that $\beta_1 = 0$ versus the alternative that $\beta_1 = B$. Note that e^{β_1} is the change in the rate for a one-unit change in X_1 when the rest of the covariates are held constant. The procedure assumes that this hypothesis will be tested using the score statistic

$$z = \frac{\hat{\beta}_1}{\sqrt{\text{Var}(\hat{\beta}_1)}}$$

The Poisson Distribution

The Poisson distribution models the probability of y events (i.e. failure, death, or existence) with the formula

$$\Pr(Y = y | \mu) = \frac{e^{-\mu} \mu^y}{y!} \quad (y = 0, 1, 2, \dots)$$

Notice that the Poisson distribution is specified with a single parameter μ . This is the mean incidence rate of a rare event per unit of *exposure*. Exposure may be time, space, distance, area, volume, or population size. Because exposure is often a period of time, we use the symbol t to represent the exposure. When no exposure value is given, it is assumed to be one.

The parameter μ may be interpreted as the risk of a new occurrence of the event during a specified exposure period, t . The probability of y events is then given by

$$\Pr(Y = y | \mu, t) = \frac{e^{-\mu t} (\mu t)^y}{y!} \quad (y = 0, 1, 2, \dots)$$

The Poisson distribution has the property that its mean and variance are equal.

The Poisson Regression Model

In Poisson regression, we suppose that the Poisson incidence rate μ is determined by a set of k regressor variables (the X 's). The expression relating these quantities is

$$\mu = \exp(\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_k X_k)$$

The regression coefficients $\beta_0, \beta_1, \beta_2, \dots, \beta_k$ are unknown parameters that are estimated from a set of data. Their estimates are labeled b_0, b_1, \dots, b_k .

Using this notation, the fundamental Poisson regression model for an observation i is written as

$$\Pr(Y_i = y_i | \mu_i, t_i) = \frac{e^{-\mu_i t_i} (\mu_i t_i)^{y_i}}{y_i!}$$

where

$$\mu_i = \lambda(\mathbf{x}_i' \boldsymbol{\beta})$$

$$\lambda(\mathbf{x}_i' \boldsymbol{\beta}) = \exp(\beta_0 + \beta_1 X_{1i} + \beta_2 X_{2i} + \cdots + \beta_k X_{ki})$$

That is, for a given set of values of the regressor variables, the outcome follows the Poisson distribution.

Power Calculations

Suppose you want to test the null hypothesis that $\beta_1 = 0$ versus the alternative that $\beta_1 = B1$. Signorini (1991) gives the formula relating sample size, α , β , and $B1$ when X_1 is the only covariate of interest as

$$N = \phi \frac{\left(z_{1-\alpha/2} \sqrt{V(b_1 | \beta_1 = 0)} + z_{1-\beta} \sqrt{V(b_1 | \beta_1 = B1)} \right)^2}{\mu_T e^{\beta_0} B1^2}$$

where N is the sample size, ϕ is a measure of over-dispersion, μ_T is the mean exposure time, and z is the standard normal deviate. Following the extension used in Hsieh, Block, and Larsen (1998) and Hsieh and Lavori (2000), when there are other covariates, *PASS* uses the approximation

$$N = \phi \frac{\left(z_{1-\alpha/2} \sqrt{V(b_1 | \beta_1 = 0)} + z_{1-\beta} \sqrt{V(b_1 | \beta_1 = B1)} \right)^2}{\mu_T e^{\beta_0} B1^2 (1 - R^2)}$$

where R^2 is the square of the multiple correlation coefficient when the covariate of interest is regressed on the other covariates. The variance in the null case is given by

$$V(b_1 | \beta_1 = 0) = \frac{1}{\text{Var}(X_1)}$$

The variance for the non-null case depends on the underlying distribution of X . Common choices are given next.

Normal

$$V(b_1|\beta_1 = B1) = e^{-\left(B1\mu_{X1} + \frac{B1^2\sigma_{X1}^2}{2}\right)}$$

$$V(X1) = \sigma_{X1}^2$$

Exponential

$$V(b_1|\beta_1 = B1) = \frac{(\lambda_{X1} - B)^3}{\lambda_{X1}}$$

$$V(X1) = \lambda_{X1}^{-2}$$

Uniform, on the Interval [C,D]

$$V(b_1|\beta_1 = B1) = \frac{m}{m(m_{11}) - m_1^2}$$

$$V(X1) = \frac{(D - C)^2}{12}$$

where

$$m = \frac{e^{B1D} - e^{B1C}}{(D - C)B1}$$

$$m_1 = \frac{e^{B1D}(B1D - 1) - e^{B1C}(B1C - 1)}{(D - C)B1^2}$$

$$m_{11} = \frac{e^{B1D}(2 - 2B1D + B1^2 D^2) - e^{B1C}(2 - 2B1C + B1^2 C^2)}{(D - C)B1^3}$$

Binomial, Parameter π_{X1}

$$V(b_1|\beta_1 = B1) = \frac{1}{1 - \pi_{X1}} + \frac{1}{\pi_{X1}e^{B1}}$$

$$V(X1) = \pi_{X1}(1 - \pi_{X1})$$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Axes/Legend/Grid, Plot Text, or Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains most of the parameters and options that you will be concerned with.

Solve For

Find (Solve For)

This option specifies the parameter to be solved for from the other parameters. Under most situations, you will select either *Power and Beta* for a power analysis or *N* for sample size determination.

Select *N* when you want to calculate the sample size needed to achieve a given power and alpha level. Select *Power and Beta* when you want to calculate the power of an experiment.

Error Rates

Power or Beta

This option specifies one or more values for power or for beta (depending on the chosen setting). Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta is the probability of a type-II error, which occurs when a false null hypothesis is not rejected.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also very popular.

A single value (e.g., 0.90) may be entered here or a range of values (e.g., 0.8 to 0.95 by 0.05) may be entered.

Alpha (Significance Level)

This option specifies one or more values for the probability of a type-I error. A type-I error occurs when a true null hypothesis is rejected.

Values must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

A single value (e.g., 0.05) may be entered here or a range of values (e.g., 0.05 to 0.2 by 0.05) may be entered.

Sample Size

N (Sample Size)

This option specifies the total number of observations in the sample. You may enter a single value (e.g., 25) or a list of values (e.g., 10 to 100 by 10).

Test

Alternative Hypothesis

Specify whether the alternative hypothesis is one-sided or two-sided. When a two-sided alternative hypothesis is selected, the specified value of alpha is divided by 2. When a one-sided hypothesis is selected, the value of alpha is left as specified.

In general, the accepted procedure is to use the Two-Sided option unless a one-sided test can be justified.

Effect Size – Response Rates

Exp(B0) [Baseline Response Rate]

Exp(B0) is the baseline response rate. This is the response rate that occurs when all covariates are equal to zero. Depending on the dataset being analyzed, the response rate might be the hazard rate, death rate, survival rate, or accident rate.

Exp(B1)/Exp(B0) (Rate Ratio)

$B1$ is the value of the regression coefficient under the alternative hypothesis. In Poisson regression, it is more natural to specify a value for $\exp(B1)/\exp(B0)$ than for $B1$ because $\exp(B1)/\exp(B0)$ represents the ratio of the response rate when $X1$ is increased one unit and all other covariates are constant to its baseline rate. Depending on the dataset being analyzed, the response rate might be the hazard rate, death rate, survival rate, or accident rate.

For example, suppose the baseline response rate (the value of $\exp(\beta_0)$) is 0.70 and you want the study to be large enough to detect a 30% increase in the response rate. The value entered here would be 1.3.

Effect Size – Exposure Time and Over-dispersion

MuT (Mean Exposure Time)

This is the mean exposure time, μ_T , over which the response rate is calculated. If the response rates are for one year, enter 1 here. If they are for 30 days, enter 30 here.

Phi

Phi is the over-dispersion parameter for Poisson regression. When there is no over-dispersion, set this value to one.

Effect Size – Covariates (X1 is the Variable of Interest)

R-Squared of X1 with Other X's

This is the R -Squared that is obtained when X_1 is regressed on the other X 's (covariates) in the model. Use this to account for the influence on power and sample size of adding other covariates. Note that the number of additional variables does not matter in this formulation. Only their overall relationship with X_1 through this R -Squared value is used.

870-6 Poisson Regression

This value is restricted to being greater than or equal to zero and less than one. Use zero when there are no other covariates.

Distribution of X1

This option specifies the distribution of the covariate being analyzed. You can choose from the normal, exponential, uniform, and binomial distributions. Once you have selected the anticipated distribution of the covariate, you must also specify values for the corresponding parameter(s) of the distribution.

Normal(M)

Specify the mean of the normal covariate distribution. This option is used when Distribution of X1 is set to Normal(M,S).

Normal(S)

Specify the standard deviation of the normal covariate distribution. This option is used when Distribution of X1 is set to Normal(M,S).

Uniform(C)

Specify the minimum of the uniform covariate distribution. This option is used when Distribution of X1 is set to Uniform(C,D). The uniform distribution of the covariate is assumed to range from C to D .

Uniform(D)

Specify the maximum of the uniform covariate distribution. This option is used when Distribution of X1 is set to Uniform(C,D). The uniform distribution of the covariate is assumed to range from C to D .

Exponential(L)

Specify the value of λ associated with the exponential covariate distribution. This option is used when Distribution of X1 is set to Exponential(L). Note that the pdf of the exponential distribution is assumed to be $(\lambda) \exp(-\lambda X)$.

Binomial(P)

Specify the value of the proportion for the binomial covariate distribution. This option is used when the Distribution of X1 is set to Binomial(P). For example, if the covariate being studied is an indicator variable of treatment and control and you anticipated having an equal number of treatments and controls, this value should be set to 0.50.

Iterations Tab

This tab sets a couple of options used in the iterative procedures.

Maximum Iterations

Maximum Iterations Before Search Termination

Specify the maximum number of iterations allowed before the search for the criterion of interest is aborted. When the maximum number of iterations is reached without convergence, the criterion is left blank. A value of 500 is recommended.

Example 1 – Power for Several Sample Sizes

Poisson regression will be used to analyze the power for a study of the relationship between the number of flaws on a manufactured article and the experience (measured in years) of the operator. The researchers want to evaluate the sample size needs for detecting ratios in response rates of 1.3 and 1.5. The experience of an operator is assumed to be normally distributed with mean 3.2 and standard deviation 2.1. No other covariates will be included in the analysis. The researchers will test their hypothesis using a 5% significance level with a two-sided Wald test. They decide to calculate the power at sample sizes between 5 and 50.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Poisson Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Poisson Regression (Y is a Count)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find (Solve For)	Power and Beta
Power	<i>Ignored since this is the Find setting</i>
Alpha	0.05
N (Sample Size)	5 to 50 by 5
Alternative Hypothesis	Two-Sided
Exp(B0)	1.0
Exp(B1)	1.3 1.5
MuT (Mean Exposure Time)	1.0
Phi	1.0
R-Squared of X1 with Other X's	0.0
Distribution of X1	Normal(M,S)
Normal(M)	3.2
Normal(S)	2.1

Annotated Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

**Numeric Results when X1 is Normal with Mean = 3.2 and Sigma = 2.1
And Phi (Over-Dispersion Parameter) = 1.0000**

	Sample	Response	Baseline	Mean	R-Squared	Two-	
Power	Size	Rate	Rate	Exposure	X1 vs	Sided	Beta
	(N)	Ratio	Exp(B0)	Time	Other X's	Alpha	
				(MuT)	(R2)		
0.28466	5	1.3000	1.0000	1.0000	0.0000	0.05000	0.71534
0.43245	10	1.3000	1.0000	1.0000	0.0000	0.05000	0.56755
0.55407	15	1.3000	1.0000	1.0000	0.0000	0.05000	0.44593
0.65321	20	1.3000	1.0000	1.0000	0.0000	0.05000	0.34679
0.73282	25	1.3000	1.0000	1.0000	0.0000	0.05000	0.26718

870-8 Poisson Regression

0.79585	30	1.3000	1.0000	1.0000	0.0000	0.05000	0.20415
0.84516	35	1.3000	1.0000	1.0000	0.0000	0.05000	0.15484
0.88334	40	1.3000	1.0000	1.0000	0.0000	0.05000	0.11666
0.91262	45	1.3000	1.0000	1.0000	0.0000	0.05000	0.08738
0.93491	50	1.3000	1.0000	1.0000	0.0000	0.05000	0.06509
0.47562	5	1.5000	1.0000	1.0000	0.0000	0.05000	0.52438
0.78817	10	1.5000	1.0000	1.0000	0.0000	0.05000	0.21183
0.92798	15	1.5000	1.0000	1.0000	0.0000	0.05000	0.07202
0.97821	20	1.5000	1.0000	1.0000	0.0000	0.05000	0.02179
0.99394	25	1.5000	1.0000	1.0000	0.0000	0.05000	0.00606
0.99842	30	1.5000	1.0000	1.0000	0.0000	0.05000	0.00158
0.99961	35	1.5000	1.0000	1.0000	0.0000	0.05000	0.00039
0.99991	40	1.5000	1.0000	1.0000	0.0000	0.05000	0.00009
0.99998	45	1.5000	1.0000	1.0000	0.0000	0.05000	0.00002
1.00000	50	1.5000	1.0000	1.0000	0.0000	0.05000	0.00000

Report Definitions

Power is the probability of rejecting a false null hypothesis. It should be close to one.

N is the size of the sample drawn from the population.

$\text{Exp}(B1)/\text{Exp}(B0)$ is the response rate ratio due to a one-unit change in $X1$.

$\text{Exp}(B0)$ is the response rate when all covariates have a value of zero.

Phi is the over-dispersion parameter used when the Poisson model does not fit.

R2 is the R-squared achieved when $X1$ is regressed on the other covariates.

Alpha is the probability of rejecting $\text{Exp}(B1)/\text{Exp}(B0)$ is one.


Beta is the probability of accepting a false null hypothesis.

Summary Statements

A Poisson regression of a dependent variable of counts on normally distributed independent variable with mean = 3.2 and standard deviation = 2.1 using a sample of 5 observations achieves 28% power at a 0.71534 significance level to detect a response rate ratio of at least 1.3000 due a one-unit change in the IV. The baseline rate is 1.0000 and the mean exposure time is 1.0000.

This report shows the power for each of the scenarios. Note that if you were interested in $B1$ instead of $\text{Exp}(B1)$, you would simply take the natural logarithm of the value of $\text{Exp}(B1)$.

Plots Section



Example 2 – Validation using Signorini

Signorini (1991), page 449, presents an example which we will use to validate this program. In the example, $\text{Exp}(B1)/\text{Exp}(B0) = 1.3$, $\text{Exp}(B0) = 0.85$, $R^2 = 0.0$, $\text{Mu } T = 1.0$, and $\text{Phi} = 1.0$. The independent variable is assumed to be binomial with proportion 0.5. A one-sided test with $\alpha = 0.05$ will be used. Sample sizes for power = 0.80, 0.90, and 0.95 are calculated to be 406, 555, and 697.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Poisson Regression** procedure window by clicking on **Regression (Y is a Function of X's)**, then **Poisson Regression (Y is a Count)**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Find	N
Power	0.80 0.90 0.95
Alpha	0.05
N (Sample Size)	<i>Ignored since this is the Find setting</i>
Alternative Hypothesis	One-Sided
Exp(B0)	0.85
Exp(B1)	1.3
MuT (Mean Exposure Time)	1.0
Phi	1.0
R-Squared of X1 with Other X's	0.0
Distribution of X1	Binomial(P)
Binomial(P)	0.5

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results

Numeric Results when X1 is Binomial with Proportion = 0.5 And Phi (Over-Dispersion Parameter) = 1.0000							
Power	Sample Size (N)	Response Rate Ratio	Baseline Rate Exp(B0)	Mean Exposure Time (MuT)	R-Squared X1 vs Other X's (R2)	Two- Sided Alpha	Beta
0.95000	697	1.3000	0.8500	1.0000	0.0000	0.05000	0.05000
0.90000	556	1.3000	0.8500	1.0000	0.0000	0.05000	0.10000
0.80000	406	1.3000	0.8500	1.0000	0.0000	0.05000	0.20000

Note that **PASS** calculated 556 rather than the 555 calculated by Signorini (1991). The discrepancy is due to rounding.

Chapter 880

Randomization Lists

Introduction

This module is used to create a randomization list for assigning subjects to one of up to eight groups or treatments. Six randomization algorithms are available. Four of the algorithms (Efron's biased coin randomization, Smith's randomization, Wei's urn randomization, and random sorting using maximum allowable % deviation) are designed to generate balanced random samples throughout the course of an experiment. The other two (complete randomization and random sorting) are less complex but have higher probabilities of imbalance. Technical details for each of these algorithms (except random sorting) as well as a discussion of randomization issues in general are given in Rosenberger and Lachin (2002), Pocock (1983), and Piantadosi (2005).

Why Randomize?

Random allocation has been an important part of scientific discovery since the time of R.A. Fisher. He and other scientists recognized the potential for bias in nonrandomized experiments that could compromise treatment comparisons. Randomization controls the influence of unknown factors, reducing potential bias in an experiment. Rosenberger and Lachin (2002) state on page 18 that "...the randomization of subjects between the treatment groups is the paramount statistical element that allows one to claim that a study is unbiased." The elimination of bias, however, is not the only reason for randomization. Pocock (1983) states on page 50 that randomization is used "to provide a basis for the standard methods of statistical analysis such as significance tests." Clearly, randomization is necessary in order to produce accurate and generalizable statistical conclusions.

Randomization in Clinical Trials

In non-clinical (or non-human) experiments, researchers often have tight control over the environment and conditions surrounding an experiment, so randomization can usually be implemented with minor difficulty. Clinical experiments, however, are quite different. Because human patients are used in clinical trials, various ethical issues must be addressed. These ethical considerations complicate the experimental design and often require adjustments in the way subjects are randomly assigned to treatments. Other factors that influence randomization in clinical trials are purely logistical. How can the investigators ensure that treatments are administered the same for all patients without the patient or the doctor knowing what treatment is being given? These and other design factors influence how randomization is administered in clinical trials.

One of the issues that arise in clinical experiments is treatment imbalance. Clinical trials are usually administered over time with patients enrolling at different time points throughout the study. It is often desirable to maintain balance in the number of patients assigned to each treatment throughout the course of the experiment. This is particularly true when time-dependent covariates influence the response or when sequential testing will be used to analyze results. Several randomization algorithms have been developed to produce lists that balance the number of patients assigned to each treatment throughout the experiment while still maintaining the randomness of the assignments. These include Efron's biased coin randomization, Smith's randomization, Wei's urn randomization, and random sorting using maximum allowable % deviation. Each of these algorithms will be discussed in detail in the section that follows.

Randomization Algorithms

Six different randomization algorithms are available in *PASS*. These can be roughly divided into two categories: those that aim to produce balanced randomization lists and those that do not. The following table outlines the goal of each algorithm by the number of groups each algorithm will allow.

	Non-Balancing Algorithms	Balancing Algorithms
2 Groups	Complete Randomization [†] , Random Sorting	Efron's Biased Coin ^{*†} , Smith ^{*†} , Wei's Urn ^{*†} , Random Sorting using Maximum Allowable % Deviation
k Groups	Complete Randomization [†] , Random Sorting	Wei's Urn [*] , Random Sorting using Maximum Allowable % Deviation

*These randomization algorithms have the additional restriction that unequal treatment allocation is not allowed, i.e. all groups must have the same target sample size.

†These randomization algorithms produce randomization lists in which the actual group sample size may not equal the target group sample size.

The discussion of each algorithm that follows will be based on the following scenario and notation. Suppose we have k treatments and that n_i subjects (not necessarily all equal) are to be assigned to each treatment, $i = 1, 2, \dots, k$. The value n_i will be referred to in discussion as the “target” sample size for each group. Let the actual sample size for each group be a_i . For some algorithms, the actual group sample size may not always equal the target sample size for all groups. The total sample size is

$$N = \sum_{i=1}^k n_i = \sum_{i=1}^k a_i ,$$

and the target allocation ratio for each group is

$$R_i = \frac{n_i}{N} .$$

Define the probability of assignment of subject j to treatment i as p_{ij} .

Define the number of subjects in each group after the j^{th} subject is assigned as $n_i[j]$.

Non-Balancing Algorithms

Complete Randomization

The complete randomization algorithm is commonly referred to as a “coin flip”. For the case to two treatments, a coin is flipped each time a subject is to be randomized, determining the assignment. The probability of assignment to either treatment is 0.5 for all subjects.

The algorithm generalizes to more than two groups and allows for unequal allocation. The probability of assignment of subject j to group i is

$$p_{ij} = R_i$$

for all j . If the target sample size is the same for all k groups, then

$$p_{ij} = 1/k$$

The complete randomization algorithm proceeds by randomly assigning subjects to treatments using the above assignment probabilities. The algorithm may result in imbalances between groups even when the target group sample sizes are equal, i.e. the actual sample sizes may not always equal the target sample sizes for all groups. This algorithm is not recommended when balance is important throughout the course of an experiment.

Random Sorting

The random sorting algorithm can be used for any number of treatment groups and any allocation ratios. The random sorting algorithm always results in randomized assignment lists in which the actual group sample sizes match the target group sample sizes, i.e. $a_i = n_i$ for all i . The algorithm begins by creating a virtual database containing group names. Each row corresponds to one group, and the i^{th} group is represented by n_i rows. For example, if we were randomly assigning three groups (A, B, and C) with ten subjects in each group, then the database would consist of ten rows containing A's, followed by ten rows containing B's and ten rows containing C's, for a total of 30 rows. Next, a random number is assigned to each row in the database, and then the database is sorted based on the column of random numbers to place the group names in random order. The first subject is then assigned to the group in row one of the randomly sorted database, the second subject is assigned to the group in row two, and so forth. Although this algorithm always results in groups containing the target sample sizes, longitudinal imbalances among groups may still occur, therefore, this algorithm is not recommended when balance is important throughout the course of an experiment.

Balancing Algorithms

Efron's Biased Coin Randomization

This algorithm may only be used for random assignment of subjects to two treatments. The target sample sizes must also be the same for both groups. In order to achieve longitudinal balance between groups, the algorithm dynamically changes the group assignment probabilities. The algorithm is outlined in Efron (1971).

First we define a constant probability p (called “Efron's p ” in *PASS*), where $0.5 < p \leq 1$. A common value for p is $2/3$. Also define a difference function $D_j = n_1[j] - n_2[j]$. The probability of assignment of subject j to group 1 is

$$p_{1j} = \begin{cases} 1/2 & \text{if } D_{j-1} = 0 \\ p & \text{if } D_{j-1} < 0 \\ 1-p & \text{if } D_{j-1} > 0 \end{cases}$$

Efron's biased coin randomization proceeds by randomly assigning subjects to treatments using the above assignment probabilities. When group 1 has more subjects assigned than group 2, the assignment probability changes to make group 2 more probable for assignment. When group 2 has more, then group 1 becomes more probable. The algorithm may result in final imbalances between groups, but the degree of imbalance throughout the randomization process is greatly reduced.

Smith's Randomization

This algorithm may only be used for random assignment of subjects to two treatments. The target sample sizes must also be the same for both groups. Like Efron's biased coin randomization, Smith's algorithm dynamically changes the group assignment probabilities based on the degree of imbalance to achieve longitudinal balance between groups. The algorithm is outlined in Smith (1984).

We define a positive exponent parameter ρ (called "Smith's Exponent" in *PASS*). The probability of assignment of subject j to group 1 is

$$p_{1j} = \frac{n_2[j-1]^\rho}{n_1[j-1]^\rho + n_2[j-1]^\rho}$$

Smith's randomization proceeds by randomly assigning subjects to treatments using the above assignment probabilities. When group 1 has more subjects assigned than group 2, the assignment probability changes to make group 2 more probable for assignment. When group 2 has more, then group 1 becomes more probable. The algorithm may result in final imbalances between groups, but the degree of imbalance throughout the randomization process is greatly reduced.

Wei's Urn Randomization

This algorithm may be used for random assignment of subjects to two or more treatments. The target sample sizes must also be the same for all groups. Like Smith's randomization, Wei's urn randomization algorithm dynamically changes the group assignment probabilities based on the degree of imbalance to achieve longitudinal balance between groups. Urn randomization is reviewed in Wei and Lachin (1988).

Define positive parameters A and B (called "Wei's A" and "Wei's B" in *PASS*, respectively). We start the algorithm by placing A balls representing each group in an urn. A single ball is then randomly chosen from the urn, recorded, and *replaced*, and then B new balls corresponding to each of the other groups are added to the urn. Therefore, when a ball from one group is chosen, the probability shifts to make the other groups more probable on the next draw. The probability of the first assignment is $1/k$. After that, the probability of assignment of subject j to group i is

$$p_{ij} = \frac{A + B(j-1) - Bn_i[j-1]}{kA + B(j-1)(k-1)}.$$

The algorithm continues until all subjects have been assigned to one of the groups. The algorithm may result in final imbalances between groups, but the degree of imbalance throughout the

randomization process is diminished due to the shifting of probabilities toward the underrepresented groups.

Random Sorting using Maximum Allowable % Deviation

This algorithm is equivalent to the random sorting algorithm described earlier except that a search is conducted to find a randomization list that satisfies the Maximum Allowable % Deviation criterion. The % Deviation for group i after subject j has been assigned is defined as

$$\begin{aligned}\% \text{ Deviation}_{ij} &= \left| \frac{n_i[j] - E(n_i[j])}{n_i} \right| \times 100 \\ &= \left| \frac{n_i[j] - jR_i}{n_i} \right| \times 100.\end{aligned}$$

The % Deviation measures how far the actual sample size for group i is from the expected sample size after subject j is randomly assigned. The Maximum Allowable % Deviation represents the upper bound for this measure. The search is conducted by creating an assignment list based on random sorting and then running through the assignments and calculating the maximum % Deviation for all groups after each assignment. If the maximum % Deviation is greater than the Maximum Allowable % Deviation value specified, then the list is rejected, the number of iterations is incremented, and the random sorting algorithm is started again with a new set of random numbers. The search continues until a randomization list is generated for which the criterion is satisfied for all individual assignments. Conducting a search in this manner assures a degree of balance throughout the course of the experiment.

For example, for 40 subjects to be assigned to two groups A and B with equal allocation ratios (0.5), suppose that there are 7 assigned A's and 3 assigned B's after 10 random assignments and the Maximum Allowable % Deviation is 10%. With the allocation ratio at 0.5, we would expect to have 5 A's and 5 B's ($10 \times 0.5 = 5$) after 10 assignments. Therefore, the % Deviation for group A is $|7 - 5|/20 = 10\%$ and the % Deviation for group B is $|3 - 5|/20 = 10\%$. Both of these are equal to the Maximum Allowable % Deviation so the next assignment would be tested. If the next assignment were to group A then the randomization list would be rejected because the % Deviation for group A is $|8 - 5|/20 = 15\% > 10\%$. A new randomization list based on random sorting would be generated and the search would continue.

Comparison of Balancing Properties

Rosenberger and Lachin (2002) provides a simulation comparison of the balancing properties of complete randomization, Efron's biased coin ($p = 2/3$), Wei's Urn ($A = 0, B = 1$), and Smith (exponent = 5). The simulation was carried out for two treatment groups with target sample sizes of 25 in each group. They found that complete randomization did not balance as well as the other three "restricted randomization" procedures. Efron's biased coin and Smith's randomization algorithms were very close in terms of bias and variability. Wei's urn was found to be slightly more variable.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Design, Report, and Storage tabs. To find out more about using the Reports and Template tabs, go to the Procedure Window chapter.

Design Tab

The Design tab contains most of the parameters and options that you will be concerned with.

Specify the Groups

Number of Groups

This option specifies the number of groups or treatments to which the subjects will be randomly assigned. Up to eight groups are allowed for random assignment.

Specify the Groups – Group Labels and Target Group Sizes

All Group Sizes are Equal

Check this box if all groups have the same target sample size. The target sample size is then specified in the Equal Group Size box to the right. This is used as the target sample size for every group. If this box is not checked, then a target sample size must be entered for each group individually.

Equal Group Size

Specify a single target sample size to be used for all of the groups. This box is only used if the All Group Sizes are Equal option is checked. If the target sample sizes are not all equal, then they must be entered individually for each group.

Group Label

Specify a label or treatment name for each group. This label will be used to distinguish groups in the output and reports. Each group must have a different label.

Group Size

Specify a target sample size for the group designated to the left of this box. This option is only used if the All Group Sizes are Equal box is unchecked.

If all of the target group sizes are equal, check the All Group Sizes are Equal box and enter a single number in the Equal Group Size box.

Randomization Algorithm

Randomization Algorithm

Choose one of six randomization algorithms. These algorithms are each described in detail in the Randomization Algorithms section above. Some of the algorithms may result in actual (final) sample sizes that are not equal to the target sample sizes. To search for a randomization list in which the final group sizes are equal to the target group sample sizes for every group, check the “Search for a randomization list...” box.

The following algorithms are available.

- **Complete Randomization**

This is the simplest of all randomization algorithms and is commonly known as “coin flipping”. The algorithm may result in imbalances between groups even when the target group sample sizes are equal, i.e. the actual sample sizes may not always equal the target sample sizes for all groups. This algorithm is not recommended when balance is important throughout the course of an experiment.

- **Efron’s Biased Coin (2 Treatments, Equal Sample Sizes)**

This algorithm may only be used for random assignment of subjects to two treatments. The target sample sizes must also be the same for both groups. In order to achieve longitudinal balance between groups, the algorithm dynamically changes the group assignment probabilities. A value for Efron’s p must also be specified.

- **Random Sorting**

The random sorting algorithm can be used for any number of treatment groups and any allocation ratio set. The random sorting algorithm always results in randomized assignment lists in which the actual group sample sizes match the target group sample sizes, i.e. $a_i = n_i$ for all i . The random assignment is conducted by sorting a database of group labels using a column of random numbers. This algorithm is not recommended when balance is important throughout the course of an experiment.

- **Random Sorting using Max Allowable % Deviation (Search)**

This algorithm is equivalent to the random sorting algorithm except that a search is conducted to find a randomized list for which all % Deviations are less than or equal to the Maximum Allowable % Deviation specified. The search is conducted until a list is found or the maximum iterations are reached. A value for Maximum Allowable % Deviation must also be specified.

- **Smith (2 Treatments, Equal Sample Sizes)**

This algorithm may only be used for random assignment of subjects to two treatments. The target sample sizes must also be the same for both groups. Smith’s algorithm dynamically changes the group assignment probabilities based on the degree of imbalance to achieve longitudinal balance between groups. A value for Smith’s Exponent must also be specified.

- **Wei’s Urn (Equal Sample Sizes)**

This algorithm may be used for random assignment of subjects to two or more treatments. The target sample sizes must also be the same for all groups. Wei’s urn randomization algorithm dynamically changes the group assignment probabilities based on the degree of imbalance to achieve longitudinal balance between groups. Values for Wei’s A and Wei’s B must also be specified.

Search for a randomization list in which all final group sizes exactly match the target group sizes

Some of the randomization algorithms may result in actual group sample sizes that do not exactly match the target group sample sizes specified. This option forces a search to find a randomization list in which the final sample sizes match the target sample sizes for all groups.

880-8 Randomization Lists

The search is conducted by creating a randomization list using the user-specified randomization algorithm and then looking at the final sample sizes. If the sample sizes do not match the target sample sizes for all groups, then the list is discarded and the algorithm is restarted. This continues until a list with the exact sample sizes is found.

Maximum Iterations in Search

Specify the number of iterations before the randomization list search is terminated. This option is used for the Random Sorting using Max Allowable % Deviation (Search) algorithm and in searching for a randomization list where final group sizes match the target group sizes.

Algorithm Parameters

Efron's p

Specify the probability parameter p for Efron's bias coin algorithm. This value must be greater than 0.5 and less than or equal to 1. A probability of 0.67 is often used.

Maximum Allowable % Deviation

Specify a single value for the Maximum Allowable % Deviation for the random sorting using maximum allowable % deviation algorithm. You can enter the value as a percentage (e.g. 10%) or as a decimal (e.g. 0.10).

Smith's Exponent

Specify the exponent parameter ρ for Smith's randomization algorithm. Smith (1984) favors the design with $\rho = 5$.

Wei's A

Specify the initialization parameter A for Wei's urn algorithm. This is the number of balls in the urn for each group when the algorithm starts. A value of 0 is often used.

Wei's B

Specify the parameter B for Wei's urn algorithm. This is the number of balls added to the other urns when a ball is selected. A value of 1 is often used. If B is 0 then Wei's algorithm is the same as complete randomization.

Reports Tab

The Reports tab contains options influencing the reports.

Output Options

Subject ID Prefix

Enter a text value to be added to the beginning of each subject ID. For example, if this value were "sub_" then the subjects would be labeled sub_1, sub_2, sub_3, etc. in the reports.

Storage Tab

The storage tab contains options for specifying spreadsheet storage variables.

Spreadsheet Storage

Write List to Spreadsheet

If this box is checked, then the randomization list will be written to the spreadsheet using the variables specified. Writing the list to the spreadsheet allows you to easily store the randomization list for further analysis in *NCSS* or other software. The spreadsheet can be saved as a data file.

Store Subject ID's In

Specify a spreadsheet variable in which to store the subject ID values.

Store Group Assignments In

Specify a spreadsheet variable in which to store the randomized group assignments.

Example 1 – Randomization with Equal Allocation Ratios

A clinical researcher wishes to randomly assign 30 subjects to three treatments. The requirement is that each group has exactly 10 subjects and that the maximum % deviation is no larger than 20% of the target sample size.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Randomization Lists** procedure window by clicking on **Design of Experiments**, then **Randomization Lists**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Design Tab	
Number of Groups.....	3
All Group Sizes are Equal.....	Checked
Equal Group Size	10
Grp 1 Label	A
Grp 2 Label	B
Grp 3 Label	C
Randomization Algorithm.....	Random Sorting using Max Allowable % Deviation
Maximum Iterations in Search	1000
Maximum Allowable % Deviation.....	20%

Annotated Output

Click the Run button to perform the calculations and generate the following output. Your output will be different because of random assignment.

Summary and References Sections

Summary

Randomization Algorithm	Random Sorting (Maximum Allowable % Deviation = 20%)	
Search Iterations	4	
Number of Groups	3	
Total Sample Size	30	
Group Sample Sizes	Actual	Target
-- A	10	10
-- B	10	10
-- C	10	10

References

Piantadosi, S. 2005. Clinical Trials - A Methodological Perspective. John Wiley & Sons. New Jersey.
 Pocock, S.J. 1983. Clinical Trials - A Practical Approach. John Wiley & Sons. New York.
 Rosenberger, W.F., and Lachin, J.M. 2002. Randomization in Clinical Trials - Theory and Practice. John Wiley & Sons. New York.

This report displays the summary of the randomization algorithm and references.

Randomization List Section

Randomization List

Subject ID	Group Assignment	Largest % Deviation from Target	Cumulative Sample Size (A, B, C)
1	C	6.7%	(0, 0, 1)
2	A	6.7%	(1, 0, 1)
3	B	0.0%	(1, 1, 1)
4	C	6.7%	(1, 1, 2)
5	C	13.3%	(1, 1, 3)
6	A	10.0%	(2, 1, 3)
7	C	16.7%	(2, 1, 4)
8	B	13.3%	(2, 2, 4)
9	A	10.0%	(3, 2, 4)
10	B	6.7%	(3, 3, 4)
11	A	6.7%	(4, 3, 4)
12	A	10.0%	(5, 3, 4)
13	A	16.7%	(6, 3, 4)
14	C	16.7%	(6, 3, 5)
15	B	10.0%	(6, 4, 5)
16	B	6.7%	(6, 5, 5)
17	C	6.7%	(6, 5, 6)
18	C	10.0%	(6, 5, 7)
19	B	6.7%	(6, 6, 7)
20	A	6.7%	(7, 6, 7)
21	C	10.0%	(7, 6, 8)
22	A	13.3%	(8, 6, 8)
23	A	16.7%	(9, 6, 8)
24	B	10.0%	(9, 7, 8)
25	C	13.3%	(9, 7, 9)
26	B	6.7%	(9, 8, 9)
27	C	10.0%	(9, 8, 10)
28	B	6.7%	(9, 9, 10)
29	A	6.7%	(10, 9, 10)
30	B	0.0%	(10, 10, 10)

This report shows the complete randomization list with details after each assignment. The largest observed % deviation from the target is 16.7%.

Subject ID

The identification value of the current subject.

Group Assignment

The group to which the current subject was randomly assigned.

Largest % Deviation from Target

The largest observed % deviation after the current assignment was made. This measures how far away the group sample sizes are from the expected sample size based on the targets.

Cumulative Sample Size (Grp 1, Grp 2, ..., Grp 8)

The cumulative sample size total for each group after the current assignment was made.

Example 2 – Randomization with Unequal Allocation Ratios

A researcher wishes to randomly assign 40 subjects to three treatments using complete randomization. The requirement is that first group has exactly 20 subjects and the other two have exactly 10 subjects each.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Randomization Lists** procedure window by clicking on **Design of Experiments**, then **Randomization Lists**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Design Tab	
Number of Groups.....	3
All Group Sizes are Equal	Unchecked
Grp 1 Label	Control
Grp 1 Size	20
Grp 2 Label	A
Grp 2 Size	10
Grp 3 Label	B
Grp 3 Size	10
Randomization Algorithm.....	Complete Randomization
Search for randomization list.....	Checked
Maximum Iterations in Search	1000

Output

Click the Run button to perform the calculations and generate the following output. Your output will be different because of random assignment.

Results

Summary

Randomization Algorithm	Complete Randomization	
Search Iterations	66	
Number of Groups	3	
Total Sample Size	40	
Group Sample Sizes	Actual	Target
-- Control	20	20
-- A	10	10
-- B	10	10

Randomization List

Subject ID	Group Assignment	Largest % Deviation from Target	Cumulative Sample Size (Control, A, B)
1	A	7.5%	(0, 1, 0)
2	B	5.0%	(0, 1, 1)
3	B	12.5%	(0, 1, 2)
4	Control	10.0%	(1, 1, 2)
5	Control	7.5%	(2, 1, 2)
6	Control	5.0%	(3, 1, 2)
7	Control	7.5%	(4, 1, 2)
8	A	0.0%	(4, 2, 2)
9	B	7.5%	(4, 2, 3)
10	B	15.0%	(4, 2, 4)
11	B	22.5%	(4, 2, 5)
12	B	30.0%	(4, 2, 6)
13	Control	27.5%	(5, 2, 6)
14	Control	25.0%	(6, 2, 6)
15	B	32.5%	(6, 2, 7)
16	Control	30.0%	(7, 2, 7)
17	A	27.5%	(7, 3, 7)
18	B	35.0%	(7, 3, 8)
19	A	32.5%	(7, 4, 8)
20	Control	30.0%	(8, 4, 8)
21	Control	27.5%	(9, 4, 8)
22	A	25.0%	(9, 5, 8)
23	A	22.5%	(9, 6, 8)
24	A	20.0%	(9, 7, 8)
25	Control	17.5%	(10, 7, 8)
26	Control	15.0%	(11, 7, 8)
27	B	22.5%	(11, 7, 9)
28	Control	20.0%	(12, 7, 9)
29	Control	17.5%	(13, 7, 9)
30	Control	15.0%	(14, 7, 9)
31	Control	12.5%	(15, 7, 9)
32	Control	10.0%	(16, 7, 9)
33	A	7.5%	(16, 8, 9)
34	B	15.0%	(16, 8, 10)
35	Control	12.5%	(17, 8, 10)
36	Control	10.0%	(18, 8, 10)
37	Control	12.5%	(19, 8, 10)
38	A	5.0%	(19, 9, 10)
39	A	2.5%	(19, 10, 10)
40	Control	0.0%	(20, 10, 10)

The algorithm required 66 iterations to find a list using complete randomization in which the target and actual sample sizes are equal.

Example 3 – Saving the Randomization List to the Spreadsheet

A researcher would like to randomize 20 subjects to two groups using Efron's biased coin randomization and save the randomization list to the spreadsheet. The researcher is targeting 10 per group, but will allow imbalance if it results from the algorithm.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Randomization Lists** procedure window by clicking on **Design of Experiments**, then **Randomization Lists**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Design Tab	
Number of Groups.....	2
All Group Sizes are Equal.....	Checked
Equal Group Size	10
Grp 1 Label	High
Grp 2 Label	Low
Randomization Algorithm.....	Efron's Biased Coin
Search for randomization list.....	Unchecked
Efron's p	0.67
Storage Tab	
Write List to Spreadsheet.....	Checked
Store Subject ID's In	1
Store Group Assignments In.....	2

Output

Click the Run button to perform the calculations and generate the following output. Your output will be different because of random assignment.

Results

Summary		
Randomization Algorithm	Efron's Biased Coin (p = 0.67)	
Number of Groups	2	
Total Sample Size	20	
Group Sample Sizes	Actual	Target
-- High	9	10
-- Low	11	10

880-14 Randomization Lists

Randomization List

Subject ID	Group Assignment	Largest % Deviation from Target	Cumulative Sample Size (High, Low)
1	High	5.0%	(1, 0)
2	High	10.0%	(2, 0)
3	Low	5.0%	(2, 1)
4	High	10.0%	(3, 1)
5	Low	5.0%	(3, 2)
6	Low	0.0%	(3, 3)
7	Low	5.0%	(3, 4)
8	High	0.0%	(4, 4)
9	Low	5.0%	(4, 5)
10	High	0.0%	(5, 5)
11	Low	5.0%	(5, 6)
12	High	0.0%	(6, 6)
13	Low	5.0%	(6, 7)
14	High	0.0%	(7, 7)
15	Low	5.0%	(7, 8)
16	Low	10.0%	(7, 9)
17	High	5.0%	(8, 9)
18	High	0.0%	(9, 9)
19	Low	5.0%	(9, 10)
20	Low	10.0%	(9, 11)

In this case, no search was conducted. The actual and target group sample sizes turned out to be slightly different. Go to the **PASS** Spreadsheet to view the results stored there.

The screenshot shows the 'PASS Spreadsheet - [Untitled]' window. The menu bar includes File, Edit, Window, and Help. The toolbar contains icons for New, Open, Last, Save, Cut, Copy, Paste, Find, Undo, Font, Bold, Italic, Underline, Left, Right, Back, Help, and Print. The spreadsheet grid has columns labeled C1 through C10 and rows numbered 1 through 21. The data in the grid is as follows:

	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10
1		1 High								
2		2 High								
3		3 Low								
4		4 High								
5		5 Low								
6		6 Low								
7		7 Low								
8		8 High								
9		9 Low								
10		10 High								
11		11 Low								
12		12 High								
13		13 Low								
14		14 High								
15		15 Low								
16		16 Low								
17		17 High								
18		18 High								
19		19 Low								
20		20 Low								
21										

The status bar at the bottom shows 'Column Names' and 'Sheet1'. The font settings are displayed as '9 40 FONT BUTTON:'.

Chapter 881

Two-Level Designs

Introduction

This program generates a 2^k factorial design for up to seven factors. It allows the design to be blocked and replicated. The design rows may be output in standard or random order. The design data generated by this procedure can be produced in a spreadsheet as well as the output window.

When blocking is specified, this procedure determines whether the design is listed on page 408 of Box, Hunter, and Hunter (1978). If it is one of the designs specified there, the indicated confounding pattern is used. If not, the blocks are confounded using the standard procedure in which highest-order interactions are confounded first, so long as they do not cause main effects to be confounded with blocks.

Experimental Design

Experimental design is the planning of an efficient, reliable, and accurate technical study. The range of application of experimental design principles is as broad as science and industry. One person may be planning a long-term agricultural experiment, while another may have eight hours to rectify a production problem.

Through the years, researchers and statisticians working together have outlined the basic steps necessary to conduct an effective investigation. These steps form an experimental strategy that seems to work well in many settings.

The experimental design modules lend you, the investigator, a hand with the planning and analysis of your investigation. Once you have determined the scope of your investigation, the design modules will provide a data collection plan that will minimize the amount of data collected and maximize the amount of conclusive information available.

The experimental design chapters will not attempt to teach you the principles of experimental design. There are many excellent books and pamphlets on this subject. The focus of the manual will be to remind you of the basic principles of experimental design and then explain where and how the program can help in your study. We suggest that you consult one or two of the following texts for detailed coverage of experimental design: Box, Hunter, and Hunter (1978), Davies (1971), Lawson (1987), or Montgomery (1984).

Experimental Design Definitions

Alias

Two terms are aliased if their levels are identical throughout the design (except possibly for a difference in sign). Aliasing occurs in designs that are less than one full replication. The two terms are completely confounded with one another. It is impossible to determine from the data if an effect is due to the first, second, or both terms.

Blocking

A block refers to a batch of runs conducted together. For example, a block may be the experiments run on a particular day, or the experiments conducted on a particular batch of material.

Confounding

Two terms are confounded when their influences on the response variable cannot be separated. Confounding usually occurs when blocks are equated to high-order interactions.

Experiment (Run)

An action to at least one of the items being studied which has an observable outcome. Each run produces one observation (value) of the response variable.

Experimental Design

The collection of experiments to be completed during an investigation or study.

Experimental Error

The influence on the response of all independent variables not included in the study. This *error* is a fact of life, since it is usually impossible to control every independent variable that might influence the response.

Factorial Designs

A factorial design consists of all combinations of factor levels of two or more factors. Many designs are two-level designs. Since the total number of factor-level combinations is the product of the number of levels of each factor, these two-level designs are known as 2^k factorial designs (where k is the number of factors).

The two levels of each factor are often referred to as the high and the low levels. For example, if one of the factors were agitation at 100 rpm and 200 rpm, then 100 would be the low level and 200 would be the high level.

The designs produced by this procedure are orthogonal. This means that an equal amount of information is provided about the influence of each factor. It also means that there is no overlapping of information. A study using these designs clearly shows the unique influence of each factor.

One of the greatest strengths of the factorial experiment is that it allows the study of several factors at once, rather than only one factor at a time. Since each factor is paired with all possible

combinations of the other factors, the researcher is confident that the measured effect of the factor is valid under a broad range of conditions.

Independent Variable (Factor)

A variable whose influence on the response variable is being studied by deliberately varying it from run to run.

Interaction

The interaction among factors refers to that part of the change in the response from run to run that may be accounted for by a specific combination of two or more factors. Another way of explaining interaction is that the average effect of one factor depends on the level of another factor.

The order of an interaction is the number of factors in the interaction. Hence AB is a second-order interaction and ABCD is a fourth-order interaction.

The Taylor's series expansion of a function is often used to justify the assumption that higher-order interactions are less significant (smaller influence on the response) than are main effects and low-order interactions.

Levels

A factor (independent variable) is set at different values or levels during an experiment.

Main Effect

The change in the average response as a factor is varied is called the main effect of that factor. In a factor with two levels, the main effect is the average of all runs at the high level of the factor minus the average of all runs at the low level of the factor.

Response or Dependent Variable

The variable whose value is observed at the completion of each run.

Replication

This is the number of times an experiment is repeated at identical factor levels. You must have some replication to determine the underlying (error) variability that occurs in the experiment. One *rep* refers to the running of every possible factor combination. Designs may be partially replicated (a few treatment settings are repeated), fractionally replicated (less than one complete replication), or completely replicated. It should be obvious that each time a run is repeated, the precision of the experimental results is increased.

Two-Level Factorial Designs

Many of the designs that can be produced in *PASS* are factorial designs. Two-level designs are those in which all factors have only two values. This may seem like a severe restriction, but in many studies, this is all that is needed.

Factorial designs allow you to fit linear (as opposed to quadratic) models with all possible interactions. The number of runs is often quite large, so the runs are often grouped together in blocks.

Fractional Factorial Designs

Fractional factorial designs are constructed by taking well-chosen subsets of a complete factorial design. Fractional factorials are useful because they require much fewer runs, although they do not allow the separation of main effects from high-order interactions.

This program gives two-level fractional factorial designs. These are usually defined as one-half rep, one-quarter rep, etc. They may be run all at once or in blocks.

Screening Designs

Screening designs are used in the initial phases of a study when you wish to investigate the main effects of several factors (up to 31) simultaneously. These designs allow you to determine which factors warrant closer investigation and which may be ignored.

Screening designs allow the investigation of main effects only. They use a small fraction of the total runs that would be needed for a complete factorial design.

Many of the Taguchi designs are really screening designs.

Response Surface Designs

These designs provide for factors with more than two levels. The options for that procedure are Central Composite and Box-Behnken response surface designs.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Replications

The number of replications (repeats) of the entire experiment.

Block Size

The number of experiments (runs) per block. This determines the number of blocks. This number must be a power of 2 (2, 4, 8, 16, etc.).

Sort Order

The order of the generated rows. The rows may be in random or standard order.

- **Random**

The rows are randomly ordered (random blocks and random rows within blocks). Use this option when the order of application to experimental units is governed by the row number.

- **Standard**

The rows are not randomly ordered. Instead, they are placed in standard order. Use this option when you want to quickly see the structure of the design.

Factor Values

Each factor has two possible values (levels) which are specified here. These are the values that will be written to the database. The first value is used to represent the low value. The second value represents the high value. You may use both text and numeric values.

The number of variables created depends on how many of these boxes have values in them.

Data Storage to Spreadsheet

Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

Block Column

The column to contain the block identification numbers. The blocks are numbered from one to B, where B is the number of blocks. This column is optional. If this option is left blank, no blocks will be generated.

First Factor Column

This is where the group of columns that is to contain your design begins. The K-1 columns after this column are also filled with data. The number of variables used is determined by the number of Factor Values boxes that contain data.

Warning: The program fills these columns with data, so any previous data will be lost.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Two-Level Design

This section presents an example of how to generate an experimental design using this program.

CAUTION: since the purpose of this routine is to generate design data, you should always begin with an empty spreadsheet.

In this example, we will show you how to generate a five-factor design in blocks of eight runs each.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Two-Level Designs** procedure window by clicking on **Design of Experiments**, then **Two-Level Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Two-Level Designs window, select the **Design tab**.
- Select **1** in the **Replications** box.
- Select **8** in the **Block Size** box.
- Select **Standard** in the **Sort Order** box.
- Set the **Factor 1** box to **1 2**.
- Set the **Factor 2** box to **10 20**.
- Set the **Factor 3** box to **Low High**.
- Set the **Factor 4** box to **-1 1**.
- Set the **Factor 5** box to **0 1**.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **Block Column** box.
- Enter **2** in the **First Factor Column** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Sample Design Data

Experimental Design

Row	Block	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5
1	1	1	10	Low	-1	0
2	1	2	20	Low	-1	0
3	1	2	10	High	1	0
4	1	1	20	High	1	0
5	1	2	10	High	-1	1
6	1	1	20	High	-1	1
7	1	1	10	Low	1	1
8	1	2	20	Low	1	1
9	2	2	10	Low	-1	0
10	2	1	20	Low	-1	0
11	2	1	10	High	1	0
12	2	2	20	High	1	0
13	2	1	10	High	-1	1
14	2	2	20	High	-1	1
15	2	2	10	Low	1	1
16	2	1	20	Low	1	1
17	3	2	10	High	-1	0
18	3	1	20	High	-1	0
19	3	1	10	Low	1	0
20	3	2	20	Low	1	0
21	3	1	10	Low	-1	1
22	3	2	20	Low	-1	1
23	3	2	10	High	1	1
24	3	1	20	High	1	1
25	4	1	10	High	-1	0
26	4	2	20	High	-1	0
27	4	2	10	Low	1	0
28	4	1	20	Low	1	0
29	4	2	10	Low	-1	1
30	4	1	20	Low	-1	1
31	4	1	10	High	1	1
32	4	2	20	High	1	1

The block and factor values were also produced on the spreadsheet.

These values are also generated on the spreadsheet.

Chapter 882

Fractional Factorial Designs

Introduction

This program generates two-level fractional-factorial designs of up to sixteen factors with blocking. Reports show the aliasing pattern that is used. The design rows may be output in standard or random order. The design data generated by this procedure can be produced in a spreadsheet as well as the output window.

When generating a design, the program first checks to see if the design is among those listed on page 410 of Box, Hunter, and Hunter (1978). If the requested design is not listed in the above book, the design pattern is determined using the standard procedure in which the highest-order interactions are confounded first, and so on. The procedure creates the design such that main effects are not aliased with each other.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Runs

The desired size (number of rows) of the experiment. This number must be a power of two. This number determines what fraction of a complete replicate is run. For example, suppose you are contemplating an experiment with seven factors and have budget for sixteen runs. A full replication would take $2^7 = 128$ runs. Hence, this design is a 1/8th rep (note that $16/128 = 1/8$).

Block Size

The number of experiments (runs) per block. This determines the number of blocks. This number must be a power of 2 (2, 4, 8, 16, etc.). Of course, the block size must be less than or equal to one half the number of runs.

882-2 Fractional Factorial Designs

Sort Order

The order of the generated rows. The rows may be in random or standard order.

- **Random**

The rows are randomly ordered (random blocks and random rows within blocks). Use this option when the order of application to experimental units is governed by the row number.

- **Standard**

The rows are not randomly ordered. Instead, they are placed in standard order. Use this option when you want to quickly see the structure of the design.

Factor Values

Each factor has two possible values (levels) which are specified here. These are the values that will be written to the database. The first value is used to represent the low value. The second value represents the high value. You may use both text and numeric values.

The number of variables created depends on how many of these boxes have values in them.

Data Storage to Spreadsheet

Block Column

The column to contain the block identification numbers. The blocks are numbered from one to B, where B is the number of blocks. This column is optional. If this option is left blank, no blocks will be generated.

First Factor Column

This is where the group of columns that is to contain your design begins. The K-1 columns after this column are also filled with data. The number of columns used is determined by the number of Factor Values boxes that contain data.

Warning: The program fills these columns with data, so any previous data will be lost.

Reports Tab

These options designate the variables to contain the design and the values that will be placed in those variables.

Select Reports

Design Info Report

Specifies whether to display this report.

Aliases Report

Specifies whether to display this report.

Report Options

Aliases

One of the reports shows the confounding pattern among the columns of the design. However, when several factors are confounded, the number of terms aliased with each other gets very large. This option lets you limit the amount of information that the program displays.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Fractional Factorial Design

This section presents an example of how to generate an experimental design using this program.

CAUTION: since the purpose of this routine is to generate data, any existing data will be replaced. For this reason, you should begin with an empty spreadsheet.

In this example, we will show you how to generate a six-factor design using sixteen runs separated in blocks of four runs each.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Fractional Factorial Designs** procedure window by clicking on **Design of Experiments**, then **Fractional Factorial Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Fractional Factorial Designs window, select the **Design tab**.
- Select **16** in the **Runs** box.
- Select **4** in the **Block Size** box.
- Select **Standard** in the **Sort Order** box.
- Set the first **six of the Factor Values boxes** equal to **-1, 1**.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **Block Column** box.
- Enter **2** in the **First Factor Column** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Design Information Section

Design Information Section

Design:

1/4 replication of 6 factors in 4 blocks of 4 experiments.

Defining Contrast:

$I = ABCE = BCDF = ADEF$

Design Construction:

Generate a reduced model of the factors [A B C D].
The remaining factors are aliased with interactions
of this reduced model as follows:

$E = ABC$

$F = BCD$

Blocking Section

Block:

Blocks were generated by confounding them with the
following interactions from the reduced model:
ABCD, CD

This report provides technical information about the design that was generated.

Aliases Section

One-Factor Aliases Section

A+BCE+ABCDF+DEF
B+ACE+CDF+ABDEF
C+ABE+BDF+ACDEF
D+ABCDE+BCF+AEF
E+ABC+BCDEF+ADF
F+ABCEF+BCD+ADE

Two-Factor Interaction Aliases Section

AB+CE+ACDF+BDEF
AC+BE+ABDF+CDEF
AD+BCDE+ABCF+EF
AE+BC+ABCDEF+DF
AF+BCEF+ABCD+DE
BC+AE+DF+ABCDEF
BD+ACDE+CF+ABEF
BE+AC+CDEF+ABDF
BF+ACEF+CD+ABDE
CD+ABDE+BF+ACEF
CE+AB+BDEF+ACDF
CF+ABEF+BD+ACDE
DE+ABCD+BCEF+AF
DF+ABCDEF+BC+AE
EF+ABCF+BCDE+AD

This report lists the aliases of the main effects and low-order interactions. The number of aliases listed is controlled by the Aliases Shown option.

From the first line of the report, we find that factor A (factor 1) is confounded with interactions BCE, DEF, and ABCDF. If any of the three-factor interactions are known to be real, this design would not be useful. Note that no two-factor interactions (like AB or CD) are aliased with the main effects.

1/4 Rep of a Six-Factor Design in Blocks of 4 Runs

Experimental Design

Row	Block	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6
1	1	-1	-1	-1	-1	-1	-1
2	1	1	1	1	1	1	1
3	1	1	1	-1	-1	-1	1
4	1	-1	-1	1	1	1	-1
5	2	-1	1	-1	-1	1	1
6	2	1	-1	-1	-1	1	-1
7	2	-1	1	1	1	-1	1
8	2	1	-1	1	1	-1	-1
9	3	1	-1	1	-1	-1	1
10	3	-1	1	-1	1	1	-1
11	3	-1	1	1	-1	-1	-1
12	3	1	-1	-1	1	1	1
13	4	-1	-1	1	-1	1	1
14	4	1	1	-1	1	-1	-1
15	4	-1	-1	-1	1	-1	1
16	4	1	1	1	-1	1	-1

The block and factor values were also produced on the spreadsheet.

These values are also generated on the spreadsheet.

Chapter 883

Balanced Incomplete Block Designs

Introduction

This module generates balanced incomplete block designs. Designs for up to ten treatments are available.

In order to make precise measurements of treatment means, uniform experimental conditions should be maintained when comparing a number of treatments. This insures that differences among the treatment means result from the application of the treatment and not from some extraneous factor. To achieve this, experimental trials are often grouped together into blocks. In such designs, conditions are kept constant within the blocks and allowed to vary between the blocks. The best known design of this type is the *randomized block* design. In this design, all treatments are present in each block.

Occasionally, the size of convenient blocks will not accommodate all the treatments of interest. For example, suppose you wanted to test four types of automobile tires for wear. An obvious choice for a block would be an automobile. You might select ten automobiles for the study. Assuming that the tires were rotated among the four positions, this experiment would control for differences in tire wear due to the type of automobile and the terrain that each traveled. However, blocking difficulties arise if you want to test six types of tires. You could redesign the automobile, or you could adopt a *balanced incomplete block* design.

In a balanced incomplete block design, the treatments are assigned to the blocks so that every pair of treatments occurs together in a block the same number of times. This achieves the *balance* that is described in the title of the procedure. The balance means that all differences between treatments are measured with equal precision.

Following is an example of how four treatments are assigned to blocks with a natural size of three experimental units. Four blocks are required for this balanced incomplete block design.

<u>Block</u>	<u>Treatment</u>
1	A B C
2	A B D
3	A C D
4	B C D

883-2 Balanced Incomplete Block Designs

Note that each treatment occurs three times in this experimental layout. Also note that each pair of treatments occurs twice. These are the basic properties of the balanced incomplete designs.

Box, Hunter, and Hunter (1978) point out the following rules when using such designs.

1. Randomly assign the numbers to the blocks.
2. Randomly assign the letters to the treatments.
3. Randomly assign the treatments within the blocks.
4. Randomly group blocks as replicates. A replicate is a complete set of all treatments.

If you take these steps, this design can be used effectively in those situations in which the block size and the number of treatments do not match.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

Design Limits

The designs used in this procedure were taken from Cochran and Cox (1992). We have included designs with up to ten treatments. The following table shows what block sizes are available for each number of treatments.

<u>Number of Treatments</u>	<u>Block Sizes Available</u>
4	2, 3
5	2, 3, 4
6	2, 3, 4, 5
7	2, 3, 4, 6
8	2, 4, 7
9	2, 4, 5, 6, 8
10	2, 3, 4, 5, 6, 9

Note that some block sizes are not available for certain numbers of treatments.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Block Size

This option contains the size of the blocks. That is, this is the number of experimental units that are contained in each block.

Treatment Values

The values used to represent the treatments are specified here. These values may be letters, digits, words, or numbers. The list is delimited by blanks or commas. The number of treatments is implied by the number of items in this list.

Data Storage to Spreadsheet
Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

Store First Factor In

The block identification numbers of each row of the design are stored in this variable. The treatment identification numbers (or letters) are stored in the variable immediately to the right.

Warning: The program fills these variables with data, so any previous data will be lost.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name
File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save
Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Balanced Incomplete Block Design

This section presents an example of how to generate a balanced incomplete block design using this program. **CAUTION: since the purpose of this routine is to generate data, you should always begin with an empty spreadsheet.**

In this example, we will show you how to generate a design with four treatments in blocks of two experimental units each.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Balanced Incomplete Block Designs** procedure window by clicking on **Design of Experiments**, then **Balanced Incomplete Block Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Balanced Incomplete Block Designs window, select the **Design tab**.
- Set **Block Size** to **2**.
- Set **Treatment Values** to **1 2 3 4**.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **Store First Factor In** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

BIBD with Four Treatments in Blocks of Two

Experimental Design

Row	Block	Treatment
1	1	1
2	1	2
3	2	3
4	2	4
5	3	1
6	3	3
7	4	2
8	4	4
9	5	1
10	5	4
11	6	2
12	6	3

These values were also produced on the spreadsheet.

These values are also generated on the spreadsheet.

Chapter 884

Latin Square Designs

Introduction

This module generates Latin Square and Graeco-Latin Square designs. Designs for three to ten treatments are available.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

Latin Square designs are similar to randomized block designs, except that instead of the removal of one blocking variable, these designs are carefully constructed to allow the removal of two blocking factors. They accomplish this while reducing the number of experimental units needed to conduct the experiment.

Following is an example of a four treatment Latin Square. The experimental layout is as follows:

Rows	Columns			
	Col1	Col2	Col3	Col4
Row 1	A	B	C	D
Row 2	B	C	D	A
Row 3	C	D	A	B
Row 4	D	A	B	C

In the above table, the four treatments are represented by the four letters: A, B, C, and D. The letters are arranged so that each letter occurs only once within each row and each column. Notice that a simple random design would require $4 \times 4 \times 4 = 64$ experimental units. This Latin Square needs only 16 experimental units—a reduction of 75%!

The influence of a fourth factor may also be removed from the design by introducing a second set of letters, this time lower case. This design is known as the *Graeco-Latin Square*.

Rows	Columns			
	Col1	Col2	Col3	Col4
Row 1	Aa	Bb	Cc	Dd
Row 2	Bd	Ca	Db	Ac
Row 3	Cb	Dc	Ad	Ba
Row 4	Dc	Ad	Ba	Cb

Four factors at four levels each would normally require 256 experimental units, but this design only requires 16—a reduction in experimental units of almost 94%!

The Graeco-Latin Square is formed by combining two orthogonal Latin Squares. Graeco-Latin Squares are available for all numbers of treatments except six.

Latin Square Assumptions

It is important to understand the assumptions that are made when using the Latin Square design. The large reduction in the number of experimental units needed by this design occurs because it assumes the magnitudes of the interaction terms are small enough that they may be ignored. That is, the Latin Square design is a main effects only design. Another way of saying this is that the treatments, the row factor, and the column factor affect the response independently of one another.

Assuming that there are no interactions is quite restrictive, so before you use this design you should be able to defend this assumption. In practice, the influence of the interactions is averaged into the experimental error of the analysis of variance table. We say that the experimental error is inflated. This results in a reduced F-ratio for testing the treatment factor, and a reduced F-ratio lessens the possibility of achieving statistical significance.

Randomization

Probability statements made during the analysis of the experimental data require strict attention to the randomization process. The randomization process is as follows:

1. Randomly select a design from the set of orthogonal designs available.
2. Randomly assign levels of the row factor to the rows.
3. Randomly assign levels of the column factor to the columns.
4. Randomly assign treatments to the treatment letters (or numbers as the case may be).

Orthogonal Sets

These designs were taken from Rao, Mitra, and Matthai (1966). We have included designs with up to ten treatments. The number of available squares depends on the number of treatments. The following table shows the number of orthogonal squares stored within this procedure.

<u>Number of Treatments</u>	<u>Number of Orthogonal Designs</u>
3	2
4	3
5	4
6	1
7	6
8	7
9	8
10	2

Graeco-Latin Squares are generated by combining two of the available orthogonal squares. Note that there are no six-level Graeco-Latin Squares.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Row Values

The values used to represent the rows are specified here. These values may be letters, digits, words, or numbers. The list is delimited by blanks or commas. The number of rows is implied by the number of items in this list. The number of row, column, and treatment values must be equal. From three to ten values are allowed.

Column Values

The values used to represent the columns are specified here. These values may be letters, digits, words, or numbers. The list is delimited by blanks or commas. The number of rows is implied by the number of items in this list. The number of row, column, and treatment values must be equal. From three to ten values are allowed.

Treatment 1 Values

The values used to represent the treatments are specified here. These values may be letters, digits, words, or numbers. The list is delimited by blanks or commas. The number of rows is implied by the number of items in this list. The number of row, column, and treatment values must be equal. From three to ten values are allowed.

Treatment 2 Values

The values used to represent the second set of treatments are specified here. These values may be letters, digits, words, or numbers. The list is delimited by commas. The number of rows is implied by the number of items in this list. The number of row, column, and treatment values must be equal. From three to ten values are allowed.

Note that this value is left blank unless you want to generate a Graeco-Latin Square.

Experimental Setup – Orthogonal Designs

Orthogonal Design Number I

Select one of the available orthogonal designs. The number of available orthogonal designs is given in the table in Orthogonal Sets section above. Good scientific protocol requires that you randomly choose which of these designs is used.

Orthogonal Design Number II

This option is only used when the Treatment 2 Values box is non-blank (when you are generating a Graeco-Latin Square). Select a second of the available orthogonal designs to be combined with the first in forming a Graeco-Latin Square. The value here must be different from the value specified in Orthogonal Design I. Good scientific protocol requires that you randomly choose which of these designs is used.

Data Storage Variables

Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

Store First Factor In

The row values are stored in this column of the spreadsheet. The column values are stored in the column immediately to the right. The treatment values are stored in the column immediately to the right of the column column. If specified, the values of the second treatment are stored in the column immediately to the right of the first treatment column.

Warning: The program fills these column with data, so any previous data will be replaced.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Latin Square Design

This section presents an example of how to generate a Latin Square design using this program.

CAUTION: since the purpose of this routine is to generate data, you should begin with an empty spreadsheet.

In this example, we will show you how to generate a design with four treatments.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Latin Square Designs** procedure window by clicking on **Design of Experiments**, then **Latin Square Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Latin Square Designs window, select the **Design tab**.
- Set **Row Values** to **R1 R2 R3 R4**.
- Set **Column Values** to **C1 C2 C3 C4**.
- Set **Treatment 1 Values** to **A B C D**.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **Store First Factor In** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Four-Level Latin Square Design

Experimental Design

ID	Row	Column	Treatment 1
1	R1	C1	A
2	R1	C2	B
3	R1	C3	C
4	R1	C4	D
5	R2	C1	B
6	R2	C2	A
7	R2	C3	D
8	R2	C4	C
9	R3	C1	C
10	R3	C2	D
11	R3	C3	A
12	R3	C4	B
13	R4	C1	D
14	R4	C2	C
15	R4	C3	B
16	R4	C4	A

The values were also produced on the spreadsheet.

These values are also generated on the spreadsheet.

Chapter 885

Response Surface Designs

Introduction

Response-surface designs are the only designs provided that allow for more than two levels. There are two general types of response-surface designs. The central-composite designs give five levels to each factor. The Box-Behnken designs give three levels to each factor.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

The Central-Composite designs build upon the two-level factorial designs by adding a few center points and star points. A factor's five values are: $-a$, -1 , 0 , 1 , and a . The value of a is determined by the number of factors in such a way that the resulting design is orthogonal. For example, if you are going to use either four or five factors, the value of a is 2.00.

The actual values of the levels are determined from these five values as follows:

1. The low-level value is assigned to -1 .
2. The high-level value is assigned to 1 .
3. The average of these two values is assigned to 0 .
4. The values of $-a$ and a are used to find the minimum and the maximum values.

For example, suppose we entered 50 for the low-level and 60 for the high level. Further, suppose there were four factors in the experiment. The levels would be

<u>Coded Level</u>	<u>Actual Level</u>
$-a$	45
-1	50
0	55
1	60
a	65

The values of a depend on the number of factors in the design:

<u>Factors</u>	<u>Value of a</u>
2	1.41
3	1.73
4	2.00
5	2.00
6	2.24

885-2 Response Surface Designs

The Box-Behnken designs have two differences from the central-composite designs. First, they usually use fewer runs. Second, they only use three levels while the central-composite designs use five.

The actual values of the levels are determined in the same manner as the central-composite designs, except that the value of a is ignored.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Design Type

Specify whether to generate a *central-composite* or a *Box-Behnken* design. This selection controls the number of runs generated as well as the block size (if a blocking variable is present).

Sort Order

The order of the generated rows. The rows may be in random or standard order.

- **Random**

The rows are randomly ordered (random blocks and random rows within blocks). Use this option when the order of application to experimental units is governed by the row number.

- **Standard**

The rows are not randomly ordered. Instead, they are placed in standard order. Use this option when you want to quickly see the structure of the design.

Experimental Setup – Factor Values

Factor Values

Each factor has three or five possible values (levels). The values associated with -1 and 1 are entered here.

If a Box-Behnken design was selected, the resulting three values will be -1,0,1. For example, if you entered 10 20 here, the resulting values would be 10, 15, and 20.

If a central-composite design was selected, the resulting five values will be $-a$, -1, 0, 1, a . For example, if you had four factors and entered 50 60 here, the resulting values would be 45, 50, 55, 60, and 65.

These are the values that will be written to the database. You can only use numeric values.

Data Storage to Spreadsheet

Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

Block Column

The variable to contain the block identification numbers. The blocks are numbered from one to B, where B is the number of blocks. This column is optional. If this option is left blank, no blocks will be generated.

First Factor Column

This is where the group of columns that is to contain your design begins. The K-1 columns after this column are also filled with data. The number of variables used is determined by the number of Factor Values boxes that contain data.

Warning: The program fills these columns with data, so any previous data will be lost.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Response Surface Design

This section presents an example of how to generate an experimental design using this program.

CAUTION: since the purpose of this routine is to generate data, you should always begin with an empty spreadsheet.

In this example, we will show you how to generate a three-factor central composite design with blocks.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Response Surface Designs** procedure window by clicking on **Design of Experiments**, then **Response Surface Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Response Surface Designs window, select the **Design tab**.
- Select **Central-Composite** in the **Design Type** list box.
- Select **Standard** in the **Sort Order** list box.
- Set three of the **Factor Values** boxes equal to **-1 1**.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **Block Column** box.
- Enter **2** in the **First Factor Column** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Three-Factor Response-Surface Design

Experimental Design: Central-Composite

Row	Block	Factor 1	Factor 2	Factor 3
1	1	-1	-1	-1
2	1	1	-1	-1
3	1	-1	1	-1
4	1	1	1	-1
5	1	-1	-1	1
6	1	1	-1	1
7	1	-1	1	1
8	1	1	1	1
9	1	0	0	0
10	1	0	0	0
11	1	0	0	0
12	2	-1.73	0	0
13	2	1.73	0	0
14	2	0	-1.73	0
15	2	0	1.73	0
16	2	0	0	-1.73
17	2	0	0	1.73
18	2	0	0	0
19	2	0	0	0
20	2	0	0	0

The block and factor values were also produced on the spreadsheet.

These values are also generated on the spreadsheet. Note that there are three replicates of the center points in each block. Note the star points represented by -1.73 and 1.73.

Chapter 886

Screening Designs

Introduction

Screening designs are used to find the important factors from a large number (up to 31) of two-level factors. When the number of runs is 4, 8, 16, or 32 (powers of 2), the design is a regular fractional replication. When the number of runs is 12, 20, 24, or 28, the design used is a Plackett-Burman design.

This program uses the screening designs given in Lawson (1987). These designs make it possible to evaluate each main effect, although these are aliased with several interactions.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Runs

The desired size (number of rows) of the experiment. This number must be 4, 8, 12, 16, 20, 24, 28, or 32. This number determines which design is generated.

- **Random**

The rows are randomly ordered (random blocks and random rows within blocks). Use this option when the order of application to experimental units is governed by the row number.

- **Standard**

The rows are not randomly ordered. Instead, they are placed in standard order. Use this option when you want to quickly see the structure of the design.

886-2 Screening Designs

Sort Order

The order of the generated rows. The rows may be in random or standard order.

Experimental Setup – Factor Values

Factor Values

Each factor has two possible values (levels), which are specified here. These are the values that will be written to the database. The first value is used to represent the low value. The second value represents the high value. You may use both text and numeric values, although we recommend that you stick with numeric values since these may be used in the regression program.

Enter a pair of values separated by a blank or comma, such as ‘-1 1’ or ‘0 1.’

Data Storage to Spreadsheet

Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

First Factor Column

This is where the group of columns that is to contain your design begins. The K-1 columns after this column are also filled with data. The number of variables used is determined by the number of Factor Values boxes that contain data.

Warning: The program fills these columns with data, so any previous data will be lost.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Screening Design

This section presents an example of how to generate an experimental design using this program.

CAUTION: since the purpose of this routine is to generate data, you should always begin with an empty spreadsheet.

In this example, we will show you how to generate a six-factor design using 16 runs.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Screening Designs** procedure window by clicking on **Design of Experiments**, then **Screening Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Screening Designs window, select the **Design** tab.
- Set the number of **Runs** to **16**.
- Select **Standard** in the **Sort Order** list box.
- Set six of the **Factor Values** boxes equal to **-1 1**.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **First Factor Column** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Six-Factor Screening Design in Sixteen Runs

Experimental Design

Row	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6
1	-1	-1	-1	-1	-1	-1
2	1	-1	-1	-1	1	1
3	-1	1	-1	-1	1	1
4	1	1	-1	-1	-1	-1
5	-1	-1	1	-1	1	-1
6	1	-1	1	-1	-1	1
7	-1	1	1	-1	-1	1
8	1	1	1	-1	1	-1
9	-1	-1	-1	1	-1	1
10	1	-1	-1	1	1	-1
11	-1	1	-1	1	1	-1
12	1	1	-1	1	-1	1
13	-1	-1	1	1	1	1
14	1	-1	1	1	-1	-1
15	-1	1	1	1	-1	-1
16	1	1	1	1	1	1

The values were also produced on the spreadsheet.

These values are also generated on the spreadsheet. Usually, you would specify the number of runs as close to the number of variables as possible, while still leaving some degrees of freedom for an estimate of error.

Chapter 887

Taguchi Designs

Introduction

Taguchi experimental designs, often called orthogonal arrays (OA's), consist of a set of fractional factorial designs which ignore interaction and concentrate on main effect estimation. This procedure generates the most popular set of Taguchi designs.

Taguchi uses the following convention for naming the orthogonal arrays: $L_a(b^c)$, where a is the number of experimental runs, b is the number of levels of each factor, and c is the number of variables. Designs can have factors with several levels, although two and three level designs are the most common. The L_{18} design is perhaps the most popular.

When a design is generated, the levels of each factor can be stored in the current spreadsheet--replacing any data that is already there.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Design Type

This option designates the particular design that is to be generated. The available choices are:

- **L4 2^3**
This design consists of up to 3 factors at 2 levels each. There are 4 rows.
- **L8 2^7**
This design consists of up to 7 factors at 2 levels each. There are 8 rows.
- **L12 2^{11}**
This design consists of up to 11 factors at 2 levels each. There are 12 rows.
- **L16 2^{15}**
This design consists of up to 15 factors at 2 levels each. There are 16 rows.

887-2 Taguchi Designs

- **L32 2^{31}**
This design consists of up to 31 factors at 2 levels each. There are 32 rows.
- **L64 2^{63}**
This design consists of up to 63 factors at 2 levels each. There are 64 rows.
- **L9 3^4**
This design consists of up to 4 factors at 3 levels each. There are 9 rows.
- **L27 3^{13}**
This design consists of up to 13 factors at 3 levels each. There are 27 rows.
- **L27' 3^{22}**
This design consists of up to 22 factors at 3 levels each. There are 27 rows.
- **L16' 4^5**
This design consists of up to 5 factors at 4 levels each. There are 16 rows.
- **L25 5^6**
This design consists of up to 6 factors at 5 levels each. There are 25 rows.
- **L18 $2^1 \times 3^7$**
This design consists of one factor at 2 levels and up to 7 factors at 3 levels each. There are 18 rows.
- **L36 $2^3 \times 3^{13}$**
This design consists of up to 3 factors at 2 levels and up to 13 factors at 3 levels each. There are 36 rows.
- **L36' $2^{11} \times 3^{12}$**
This design consists of up to 11 factors at 2 levels and up to 12 factors at 3 levels each. There are 36 rows.
- **L54 $2^1 \times 3^{25}$**
This design consists of one factor at 2 levels and up to 25 factors at 3 levels each. There are 54 rows.
- **L32' $2^1 \times 4^9$**
This design consists of one factor at 2 levels and up to 9 factors at 4 levels each. There are 32 rows.
- **L50 $2^1 \times 5^{11}$**
This design consists of one factor at 2 levels and up to 11 factors at 5 levels each. There are 50 rows.

Experimental Setup – Factor Specification

2 Level Factors...5 Level Factors

The number of columns of this type (number of levels) that are generated. For example, if you selected L36 $2^3 \times 3^{13}$ as the Design Type, you could specify up to three two-level factors and up to thirteen three-level factors. You would enter the number of two-level factors in the 2-Level Factors box and the number of three-level factors in the 3-Level Factors box. Entries in the unused boxes (such as 4-Level and 5-Level in this example) are ignored. If you enter more than the maximum allowed, the maximum will be used.

Data Storage to Spreadsheet

Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

First Factor Column

This is where the group of columns of the spreadsheet that is to contain your design begins. The K-1 columns after this column are also filled with data, where K is the number of factors specified.

Warning: The program fills these columns with data, so any previous data will be lost.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Taguchi Design

This section presents an example of how to generate an experimental design using this program.

CAUTION: since the purpose of this routine is to generate data, you should always begin with an empty spreadsheet.

In this example, we will show you how to generate an L18 design.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Taguchi Designs** procedure window by clicking on **Design of Experiments**, then **Taguchi Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Taguchi Designs window, select the **Design** tab.
- Select **L18 2¹ x 3⁷** in the **Design Type** list box.
- Enter **1** in the **2-Level Factors** box.
- Enter **7** in the **3-Level Factors** box.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **First Factor Column** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Taguchi L18 Design

Experimental Design

Row	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8
1	1	1	1	1	1	1	1	1
2	1	1	2	2	2	2	2	2
3	1	1	3	3	3	3	3	3
4	1	2	1	1	2	2	3	3
5	1	2	2	2	3	3	1	1
6	1	2	3	3	1	1	2	2
7	1	3	1	2	1	3	2	3
8	1	3	2	3	2	1	3	1
9	1	3	3	1	3	2	1	2
10	2	1	1	3	3	2	2	1
11	2	1	2	1	1	3	3	2
12	2	1	3	2	2	1	1	3
13	2	2	1	2	3	1	3	2
14	2	2	2	3	1	2	1	3
15	2	2	3	1	2	3	2	1
16	2	3	1	3	2	3	1	2
17	2	3	2	1	3	1	2	3
18	2	3	3	2	1	2	3	1

The values were also produced on the spreadsheet.

These values are also generated on the spreadsheet. You can use the Find/Replace facility of the spreadsheet if you want to change the values from 1, 2, 3 to something more meaningful.

Chapter 888

D-Optimal Designs

Introduction

This procedure generates D-optimal designs for multi-factor experiments with both quantitative and qualitative factors. The factors can have a mixed number of levels. For example, you could use this procedure to design an experiment with two quantitative factors having three levels each and a qualitative factor having seven levels.

D-optimal designs are constructed to minimize the generalized variance of the estimated regression coefficients. In the multiple regression setting, the matrix \mathbf{X} is often used to represent the data matrix of independent variables. D-optimal designs minimize the overall variance of the estimated regression coefficients by maximizing the determinant of $\mathbf{X}'\mathbf{X}$. Designs that are D-optimal have been shown to be nearly optimal for several other criteria that have been proposed as well.

D-optimal designs are used when there is a limited budget and a completely replicated factorial design cannot be run. For example, suppose you want to study the response to three factors: A with three levels, B with four levels, and C with eight levels. One complete replication of this experiment would require $3 \times 4 \times 8 = 96$ points (we use the word 'point' to mean an experimental unit). If only 20 points can be afforded, the D-optimal design algorithm provides a reasonable choice for the 20 points.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

D-Optimal Design Overview

This section provides a brief overview of how the D-optimal design algorithm works. It will provide a general understanding of what the algorithm is trying to accomplish so that you can make intelligent choices for the various options.

Suppose you are studying the influence of height and weight on blood pressure. If you believe that a linear (straight line) relationship exists, you will only need to look at two height values and two weight values. An experiment designed to study this relationship would require four treatment combinations. However, if you decide that the relationship may be curvilinear, you will have to include at least three levels for each factor which results in nine treatment combinations. Clearly, the appropriate experimental design depends on the anticipated functional relationship between the response variable and the factors of interest.

The D-optimal algorithm works as follows. First, specify an approximate mathematical model which defines the functional form of the relationship between the response (Y) and the independent variables (the factors). Next, generate a set of possible candidate points based on this model. Finally, from these candidates select the subset that maximizes the determinant of the $\mathbf{X}'\mathbf{X}$

matrix. This is the D-optimal design. The details of this algorithm are given in Atkinson and Donev (1992).

The number of possible designs grows rapidly as the complexity of the model increases. This number is usually so large that an exhaustive search of all possible designs for a given sample size is not feasible.

The D-optimal algorithm begins with a randomly selected set of points. Points in and out of the current design are exchanged until no exchange can be found that increases the determinant of $\mathbf{X}'\mathbf{X}$. To cut down on the running time, the number of points considered during any one iteration may be limited.

Unfortunately, this method does not guarantee that the global maximum is found. To overcome this, the algorithm is repeated several times in hopes that at least one iteration leads to the global maximum. For this reason 50 or 100 random starting sets are needed. (During the testing of the algorithm, we found that some designs required 500 starts to obtain the global maximum.)

Factor Scaling

This algorithm deals with both *quantitative* (continuous) and *qualitative* (discrete) factors. The levels of quantitative factors are scaled so that the minimum value is -1 and the maximum value is 1. Qualitative factors are included as a set of variables. For example, suppose that a qualitative variable has four values. Three independent variables are created to represent this factor:

<u>Original</u>	<u>X1</u>	<u>X2</u>	<u>X3</u>
1	-10	0	0
2	0	-1	0
3	0	0	-1
4	1	1	1

As you can see, each of these variables compares a separate group with the last group. Also note that the number of generated variables is always one less than the number of levels.

Duplicates (Replicates)

The measurement of experimental error is extremely important in the analysis of an experiment. In most cases, if an estimate of experimental error is not available, the data from the experiment cannot be analyzed. One of the best estimates of experimental error comes from points that are duplicates (often called replicates) of each other. Since D-optimal designs are often used in situations with limited budgets, the experimenter is often tempted to ignore the need for duplicates and instead add points with additional treatment combinations. The tenth commandment for experimental design should be “Thou shalt have at least four duplicates in an experiment.”

Unfortunately, the D-optimal design algorithm ignores the need for duplicates. Instead, you have to add them after the experimental design has been found. You should set aside at least four points from the algorithm. For example, suppose you have budget for 20 design points. You would tell the program that you have only 16 points. The algorithm would find the best 16 point design. You would then duplicate four of the resulting design points to provide an estimate of experimental error. We recommend that you spread these duplicates out across the experiment so you can have some indication as to whether the magnitude of the experimental error is constant across all treatment settings.

Specifying a Model

Selecting an appropriate model is subjective by nature. Often, you will know very little about the true functional form of the relationship between the response and the factor variables. A common approach is to assume that a second-order Taylor-series approximation will work fairly well. You are assuming that the true function may be approximated by parabolic surface in the neighborhood of interest. Cutting down on the complexity of the model reduces the number of points that must be added to the experimental design.

When dealing with qualitative factors, you generally limit the model to first order interactions. Higher order interactions may be studied later when a complete experiment can be run.

Augmenting an Existing Design

Occasionally, you will want to add more points to an existing experimental design. This may be accomplished by forcing the algorithm to include points that are read from the spreadsheet. The D-optimal algorithm will pick the most useful additional points from the list of candidate points. One of the attractive features of the D-optimal design algorithm is that you can refine the model as your knowledge of it increases.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

N Per Block

This option specifies the required sample size. If you are not using blocks, enter a single number giving the total sample size. The sample size must be large enough to fit the designated model. If it is not large enough, you will be shown the minimum number of points necessary.

If you are using blocks, enter the sample size for each block, separated by blanks or commas. These sample sizes do not have to be equal, although they usually are. For example, if you have three blocks, you might enter 8,8,12 which would give an overall sample size of 28. The first block will have 8 points, the second 8 points, and the third 12 points.

You must be careful when specifying blocks when you also have forced design points. In this case, the first few blocks are matched with the forced design points. The size of the blocks must match the number of forced points. For example, suppose you have already run two blocks of four each and you want to augment this with three blocks of six each. You would have eight forced points. The entry in this field would be 4,4,6,6,6. If you entered 4,3,7,6,6 an error would occur because the forced points cannot be assigned exactly to one or more blocks. The bottom line is, you cannot force partial blocks into the design.

Input Columns (Candidate and Forced)

When specified, these columns contain either a set of points to be forced into the final design, a set of candidate points from which the design is to be selected, or both. The data must be arranged so that the forced points are located at the top of the spreadsheet followed by any candidate points. When candidate points are specified, no additional candidate points are generated. If you want to force points in the design and choose the rest from among those generated by the model statement, the total number of rows in these variables must equal the total number of forced rows specified below.

Note that these variables are matched with the factors specified in the model after those factors have been sorted.

Qualitative factors must be entered using positive integers (1, 2, 3, etc.). You cannot use any other identifiers. If you have data entered using some other scheme (such as A, B, C, etc.), you will have to recode the values so that they are positive integers.

Quantitative factors must be scaled so that the minimum value is -1 and the maximum value is 1. For example, suppose an existing design has a factor whose values are 10, 15, and 20. Here the minimum is 10 and the maximum is 20. You would transform these using the formula

$$\text{Scaled} = (\text{Original} - \text{Original} - \text{Max} - \text{Min}) / (\text{Max} - \text{Min})$$

Since, in this example, Max = 20 and Min = 10, the transformation reduces to $\text{New} = (\text{Original} + \text{Original} - 30)/10 = \text{Original} / 5 - 3$. You would create a new variable using the transformation $\text{Original} / 5 - 3$. This transformation would give $10/5 - 3 = -1$, $15/5 - 3 = 0$, and $20/5 - 3 = 1$. That is, the new variable would contain -1's, 0's, and 1's instead of 10's, 15's, and 20's.

Number Duplicates

It is very important to have duplicates of at least some of the design points to provide an estimate of experimental error. This option designates the number of duplicates to be generated. The first design point is duplicated, then the second, and so on. Even though this option is convenient, we recommend that you pick appropriate points for duplication by looking at scatter plots of the design.

If your design includes blocking, you should not create duplicates since that will give erroneous block sizes. Rather, you should manually create duplicates.

Input Data Type

If you have Input Variables specified, this option specifies the type of data contained in those variables. Two types of data are possible.

- **Factor Values**

Specifies that the input data contains indices of each factor. An expanded design matrix will be generated from these factor indices using the designated model. This is the more common data type.

- **Expanded Matrix**

Specifies that the input dataset contains the expanded design matrix. That is, the quadratic, cubic, and interaction terms have been created. The model statement is not used. You would use this option when you want to specify the candidate design set in more detail than is allowed by the program. The expanded matrix must include the intercept (a column of one's) if one is to be included in the model.

Forced Points

The number of rows in the Input Variables that should be forced into the final design. These rows must be located at the top of the database, before any candidate points. If the number of forced points is equal to the number of points read in, the generated design matrix is used. Otherwise, the additional rows are used as candidate points and no other rows are generated.

Optimize the Design for this Model

Your design is optimized for the model specified here. Specify main effects (factors) with names consisting of one or more letters, such as A B C. Specify interactions using an asterisk (*), such as A*B. You can use the bar (|) symbol (see examples below) as a shorthand method to specify a complete model. You can use parentheses. You can separate terms with blanks or the '+' (plus) sign. Duplicate terms are removed during the evaluation of the model. Note that the main effects are always sorted in alphabetical order.

Some examples will help to indicate how the model syntax works:

$C + B + A + B*A + C*A$	=	$A+B+C+A*B+A*C$ (Note the sorting!)
$A B$	=	$A+B+A*B$
$B A$	=	$A+B+A*B$
$A B \ A*A \ B*B$	=	$A+B+A*B+A*A+B*B$
$A A B B$ (Max Term Order=2)	=	$A+B+A*B+A*A+B*B$
$A B C$	=	$A+B+C+A*B+A*C+B*C+A*B*C$
$(A+B)*(C+D)$	=	$A*C+A*D+B*C+B*D$
$(A+B) C$	=	$A+B+C+(A+B)*C$
	=	$A+B+C+A*C+B*C$

You can experiment with various expressions by viewing the Model Terms report.

For quantitative factors, each term represents a single variable in the expanded design matrix. For qualitative variables, each term represents a set of variables in the expanded design matrix.

Note that qualitative terms should not be squared or cubed. That is, if A is a qualitative factor, you would not include A*A or an A*A*A in your model.

Max Term Order

This option specifies that maximum number of factors that can occur in an interaction term. For example, A*B*C is a third order interaction term and if this option were set to 2, the A*B*C would be removed from the model.

This option is particularly useful when used with the bar notation to remove unwanted terms.

Qualitative Factors and Levels

List any qualitative factors here followed by the number of levels given in parenthesis. Factors in the model which do not appear here are assumed to be quantitative (continuous). For example, you might enter A(5),B(4),C(7) to indicate three qualitative factors, one with five levels, the next with four levels, and the third with seven levels. Of course, the names used here must match the names used in the model statement.

Max Iterations

Specify the number of times the algorithm is started with a new random design. Often 50 or 100 iterations are necessary and 500 is not unheard of. As the number of Inclusion Points and Removal Points are increased (see below), the number of iterations may be decreased.

We suggest that you increase this value until the optimal design is found on several iterations as reported in the Determinant Analysis report.

Inclusion Points

This is the number of candidate points considered for addition during an iteration. Instead of considering all candidate points, only this many are used. A value between 1 and N_c-1 (where N_c is the number of candidate points) may be used. Usually, a value near $N_c/2$ is adequate.

Removal Points

This is the number of points currently in the design that will be considered for removal during a particular iteration. A value between 1 and N (the desired sample size) is used. Setting this value smaller than N speeds up the search, but reduces the possibility of finding the optimal design.

Include Intercept

This option specifies whether to include the intercept in the expanded design matrix. Usually, the intercept is left out of mixture designs. The intercept is automatically deleted in designs with more than one block.

Data Storage to Spreadsheet

Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

First Factor Column

If the Input Data Type is set to Factor Values, the final design is stored in a set of contiguous columns of the spreadsheet, beginning with this column. Be careful not to overwrite existing data. If you have four factors, the design will be stored in this variable and the next three to the right. Existing data will be lost!

If the Input Data Type is set to Expanded Matrix, an index is stored in this column that represents whether the row is used in the design. If the row is not in the optimum design, a zero is stored. If the row is in the optimum design, the number of times it occurs is stored here.

First Expanded Column

This option specifies the first column in which to store the expanded version of the selected design. The rest of the expanded design columns will be stored in the columns to the right. Use this option if you want to output the expanded design matrix for use in a multiple regression procedure.

Warning: The program fills these columns with data, so existing data will be replaced.

Data Storage to Spreadsheet – Storage Options

Rename Factor Columns with Factor Labels

The names of the factors that were used in the model statement are used to rename the columns in which the design is stored.

Clear Existing Data

Clear all existing data in the design columns before writing the new design data. This is especially useful if you are experimenting with several designs of different sizes. You will not be warned that data is being lost. The data will be cleared and the new design written automatically.

Reports Tab

This panel specifies the reports that will be generated.

Select Reports

Factor Report - Expanded Design Matrix Report

These options control which reports are displayed. Some of the reports may be fairly lengthy, so you will often want to omit them.

Report Options

Precision

Specify the precision of numbers in the report. A single-precision number will show seven-place accuracy, while a double-precision number will show thirteen-place accuracy. Note that the reports are formatted for single precision. If you select double precision, some numbers may run into others. Also note that all calculations are performed in double precision regardless of which option you select here. This is for reporting purposes only.

Decimal Places

Specify the number of decimal places shown when displaying the design.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – D-Optimal Design with 10 Points, 3 Factors

This section presents an example of how to generate a D-optimal design using this program.

CAUTION: since the purpose of this routine is to generate data, you should begin with an empty spreadsheet.

In this example, we will show you how to generate a 10-point design for a study involving three quantitative factors. We want the design optimized to estimate a second-order response surface model.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **D-Optimal Designs** procedure window by clicking on **Design of Experiments**, then **D-Optimal Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design and data storage.

- On the D-Optimal Designs window, select the **Design tab**.
- Set **N per Block** to **10**.
- Set **Optimize the Design for this Model** to **A|A|B|B|C|C**.
- Set **Max Term Order** to **2**.
- Set the **First Factor Column** to **1**.
- Set the **First Expanded Column** to **5**.
- Check **Rename Factor Column with Factor Labels**.
- Check **Clear Existing Data**.

2 Specify the reports.

- On the D-Optimal Designs window, select the **Reports tab**.
- Check all reports.
- Set **Decimal Places** to **0**.

3 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

10-Point, 3 Factor D-Optimal Design

A	B	C	C4	Int't	Ax	Bx	Cx	A_A	A_B	A_C	B_B	B_C	C_C
-1	-1	-1		1	-1	-1	-1	1	1	1	1	1	1
1	-1	-1		1	1	-1	-1	1	-1	-1	1	1	1
0	0	-1		1	0	0	-1	0	0	0	0	0	1
-1	1	-1		1	-1	1	-1	1	-1	1	1	-1	1
1	1	-1		1	1	1	-1	1	1	-1	1	-1	1
1	0	0		1	1	0	0	1	0	0	0	0	0
0	1	0		1	0	1	0	0	0	0	1	0	0
1	-1	1		1	1	-1	1	1	-1	1	1	-1	1
-1	0	1		1	-1	0	1	1	0	-1	0	0	1
1	1	1		1	1	1	1	1	1	1	1	1	1

Several columns in the spreadsheet are filled with data. The results may vary from the values seen here, as there may be equally optimal designs. The first, second, and third, variables (A, B, and C) contain the actual design. You would replace the -1's with the corresponding factor's minimum value, the 1's with the maximum value, and the 0's with the average of the two.

The variables from Intercept to C_C contain the expanded design matrix. Each variable is generated by multiplying the appropriate factor values. For example, in the first row, A_B is found by multiplying the value for A, which is -1, by the value for B, which is also -1. The result is 1. The intercept is set to one for all rows. The expanded matrix is usually saved so that the design can be analyzed using multiple regression.

To use this design, you would randomly assign these ten points to the ten experimental units.

Factor Section

Name	Number Values	Type	Value1	Value2	Value3
A	3	Quantitative	-1.0000	0.0000	1.0000
B	3	Quantitative	-1.0000	0.0000	1.0000
C	3	Quantitative	-1.0000	0.0000	1.0000

A total of 27 observations will be needed for one replication.

This report summarizes the factors that were included in the design. The last line of this report gives the number of observations required for one complete replication of the experiment. This value is the product of the number of levels for each factor.

Name

The symbol(s) used to represent the factor.

Number Values

The number of values (levels) generated for each factor. For qualitative factors, this value was set in the Qualitative Factors and Levels box of the Design panel. For quantitative factors, this value is one more than the highest exponent used with this term. For example, if the model includes an A*A and nothing of a higher order, this value will be three.

Type

A factor is either quantitative or qualitative.

Value1 – Value 3

These columns list the individual values that are used as the levels of each factor when generating the expanded design matrix based on the model. Notice that the smallest is always -1 and the largest is always 1.

When the expanded design matrix is input directly, these values should be ignored.

Model Terms Section

Variables Needed	Term
1	A
1	B
1	C
1	A*A
1	A*B
1	A*C
1	B*B
1	B*C
1	C*C
9	Model Total

This report shows the terms generated by your model. You should check this report carefully to make sure that the generated model matches what you wanted. The last line of the report gives the total number of degrees of freedom (except for the intercept) required for your model. This number plus one is the minimum size of the D-optimal design for this model.

Variables Needed

The number of degrees of freedom (expanded design variables) required for this term.

Term

The name of each term.

D-Optimal Design

Original Row	Factors		
	A	B	C
1	-1	-1	-1
3	1	-1	-1
5	0	0	-1
7	-1	1	-1
9	1	1	-1
13	-1	0	0
17	0	1	0
20	0	-1	1
25	-1	1	1
27	1	1	1

This report gives the points in the D-optimal design.

Original Row

This is the row number of the point from the list of candidate points. It is only useful in those cases in which you provided the list of candidate points manually.

Factors (A B C)

These are the values of the factors. For example, the first row sets A, B, and C to -1. Remember that these are scaled values. You would transform them back into their original metric using the formula:

$$\text{Original} = (\text{Scaled}(\text{Max} - \text{Min}) + \text{Max} + \text{Min})/2$$

For example, suppose the original metric for factor A is minimum = 10 and maximum = 20. The original values would be calculated as follows:

Scaled	Formula	Original
-1	$(-1(20-10)+20+10)/2$	10
0	$(0(20-10)+20+10)/2$	15
1	$(1(20-10)+20+10)/2$	20

The values 10, 15, and 20 represent the three levels of factor A that are used in the design. They would replace the -1, 0, and 1 displayed in this report.

Determinant Analysis Section

Rank	Determinant of X'X	D-Efficiency	Percent of Maximum
1	1327104	40.95	100.00
2	1327104	40.95	100.00
3	1048576	40.00	79.01
4	1048576	40.00	79.01
5	1048576	40.00	79.01
6	1048576	40.00	79.01
7	1048576	40.00	79.01
8	921600	39.49	69.44
9	921600	39.49	69.44
10	802816	38.95	60.49
11	802816	38.95	60.49
12	802816	38.95	60.49
13	802816	38.95	60.49
14	802816	38.95	60.49
15	802816	38.95	60.49
16	746496	38.66	56.25
17	589824	37.76	44.44
18	589824	37.76	44.44
19	589824	37.76	44.44
20	589824	37.76	44.44

The maximum was achieved on 2 of 30 iterations.

This report shows the largest twenty determinants. The main purpose of this report is to let you decide if enough iterations have been run so that a global maximum has been found. Unless the maximum value was achieved on at least five iterations, you should double the number of iterations and rerun the procedure.

In this example, the top value occurred on only two iterations. In practice we would probably try another 200 iterations to find out if this is the global maximum.

Rank

Only the top twenty are shown on this report. The values are sorted by the determinant.

Determinant of $\mathbf{X}'\mathbf{X}$

This is the value of the determinant of $\mathbf{X}'\mathbf{X}$ which is the statistic that is being maximized. This value is sometimes called the generalized variance of the regression coefficients. Since this value occurs in the denominator of the variance of each regression coefficient, maximizing it has the effect of reducing the variance of the estimated regression coefficients.

D-Efficiency

D-efficiency is the relative number of runs (expressed as a percent) required by a hypothetical orthogonal design to achieve the same determinant value. It provides a way of comparing designs across different sample sizes.

$$DE = 100 \left(\frac{|\mathbf{X}'\mathbf{X}|^{1/p}}{N} \right)$$

where p is the total number of degrees of freedom in the model and N is the number of points in the design.

Percent of Maximum

This is the percentage that the determinant on this row is of the best determinant found.

Individual Degree of Freedom Section

Number	Name	Diagonal of $\mathbf{X}'\mathbf{X}$	Diagonal of $\mathbf{X}'\mathbf{X}$ Inv
1	Intercept	10.0000	0.861111
2	A	7.0000	0.250000
3	B	8.0000	0.166667
4	C	8.0000	0.166667
5	A*A	7.0000	0.722222
6	A*B	6.0000	0.250000
7	A*C	6.0000	0.250000
8	B*B	8.0000	0.861111
9	B*C	7.0000	0.194444
10	C*C	8.0000	0.861111
Determinant		1327104	
D-Efficiency		40.95345	
Trace		4.583333	
A-Efficiency		21.81818	

This report shows the diagonal elements of the $\mathbf{X}'\mathbf{X}$ and its inverse. Since the variance of each term is proportional to diagonal elements from the inverse of $\mathbf{X}'\mathbf{X}$, the last column of this report lets you compare those variances. From this report you can determine if the coefficients will be estimated with the relative precision that is desired.

For example, we can see from this example that the main effects will be estimated with the greatest precision—usually a desirable quality in a design.

Number

An arbitrary sequence number.

Name

The name of the term.

Diagonal of $\mathbf{X}'\mathbf{X}$

The diagonal element of this term in the $\mathbf{X}'\mathbf{X}$ matrix.

Diagonal of $\mathbf{X}'\mathbf{X}$ Inv

The diagonal element of this term in the $\mathbf{X}'\mathbf{X}$ inverse matrix. See the discussion above for an understanding of how this value might be interpreted.

Determinant

This is the value of the determinant of $\mathbf{X}'\mathbf{X}$ which is the statistic that is being maximized. This value is sometimes called the generalized variance of the regression coefficients. Since this value occurs in the denominator of the variance of each regression coefficient, maximizing it has the effect of reducing the variance of the estimated regression coefficients.

D-Efficiency

D-efficiency is the relative number of runs (expressed as a percent) required by a hypothetical orthogonal design to achieve the same determinant value. It provides a way of comparing designs across different sample sizes.

$$DE = 100 \left(\frac{|\mathbf{X}'\mathbf{X}|^{1/p}}{N} \right)$$

where p is the total number of degrees of freedom in the model and N is the number of points in the design.

Trace

This is the value of the trace of $\mathbf{X}'\mathbf{X}$ -inverse which is associated with A-optimality.

A-Efficiency

D-efficiency is the relative number of runs (expressed as a percent) required by a hypothetical orthogonal design to achieve the same trace value. It provides a way of comparing designs across different sample sizes.

$$AE = 100 \left(\frac{p}{\text{trace} \left(N (\mathbf{X}'\mathbf{X})^{-1} \right)} \right)$$

where p is the total number of degrees of freedom in the model and N is the number of points in the design.

Candidate Points Section

Original Row	Factors		
	A	B	C
1	-1	-1	-1
2	0	-1	-1
3	1	-1	-1
4	-1	0	-1
5	0	0	-1
6	1	0	-1
7	-1	1	-1
8	0	1	-1
9	1	1	-1
10	-1	-1	0
11	0	-1	0
12	1	-1	0
13	-1	0	0
14	0	0	0
15	1	0	0
16	-1	1	0
17	0	1	0
18	1	1	0
19	-1	-1	1
20	0	-1	1
21	1	-1	1
22	-1	0	1
23	0	0	1
24	1	0	1
25	-1	1	1
26	0	1	1
27	1	1	1

This report gives a list of candidate points from which the D-optimal design points were selected.

Original Row

This is an arbitrary identification number.

Factors (A B C)

These are the values of the factors. For example, the first row sets A, B, and C to -1. Remember that these are scaled values. You would transform them back into their original metric using the formula:

$$\text{Original} = (\text{Scaled}(\text{Max} - \text{Min}) + \text{Max} + \text{Min})/2$$

For example, suppose the original metric for factor A is minimum = 10 and maximum = 20. The original values would be calculated as follows:

Scaled	Formula	Original
-1	$(-1(20-10)+20+10)/2$	10
0	$(0(20-10)+20+10)/2$	15
1	$(1(20-10)+20+10)/2$	20

The values 10, 15, and 20 represent the three levels of factor A. They would replace the -1, 0, and 1 displayed in this report.

Expanded Design Matrix Section

Row	Variable	Intercept	A	B	C	A*A	A*B	A*C	B*B	B*C	C*C
1		1	-1	-1	-1	1	1	1	1	1	1
2		1	0	-1	-1	0	0	0	1	1	1
3		1	1	-1	-1	1	-1	-1	1	1	1
4		1	-1	0	-1	1	0	1	0	0	1
5		1	0	0	-1	0	0	0	0	0	1
6		1	1	0	-1	1	0	-1	0	0	1
7		1	-1	1	-1	1	-1	1	1	-1	1
8		1	0	1	-1	0	0	0	1	-1	1
9		1	1	1	-1	1	1	-1	1	-1	1
10		1	-1	-1	0	1	1	0	1	0	0
11		1	0	-1	0	0	0	0	1	0	0
12		1	1	-1	0	1	-1	0	1	0	0
13		1	-1	0	0	1	0	0	0	0	0
14		1	0	0	0	0	0	0	0	0	0
15		1	1	0	0	1	0	0	0	0	0
16		1	-1	1	0	1	-1	0	1	0	0
17		1	0	1	0	0	0	0	1	0	0
18		1	1	1	0	1	1	0	1	0	0
19		1	-1	-1	1	1	1	-1	1	-1	1
20		1	0	-1	1	0	0	0	1	-1	1
21		1	1	-1	1	1	-1	1	1	-1	1
22		1	-1	0	1	1	0	-1	0	0	1
23		1	0	0	1	0	0	0	0	0	1
24		1	1	0	1	1	0	1	0	0	1
25		1	-1	1	1	1	-1	-1	1	1	1
26		1	0	1	1	0	0	0	1	1	1
27		1	1	1	1	1	1	1	1	1	1

This report gives a list of candidate points expanded so that each individual term may be seen. The report is useful to show you how the expanded matrix looks. Each variable is generated by multiplying the appropriate factor values. For example, in the first row, A_B is found by multiplying the value for A, which is -1, by the value for B, which is also -1. The result is 1. The intercept is set to one for all rows.

If you want to constrain the design space, you could cut and paste these values back into the spreadsheet and then eliminate points that cannot occur.

Example 2 – Two Factors

This section presents an example of how to generate and analyze a D-optimal design involving two factors. Suppose we want to study the effect of two factor variables, A and B, on a response variable, Y. A and B happen to be quantitative variables and there is reason to believe that a second-order response surface design will work well. A full replication of this design requires nine points. In addition, four more are required to provide an estimate of experimental error. However, we can only afford eight. We will create a D-optimal design with six of the experimental units and use the remaining two as duplicates to provide the estimate of experimental error.

We want to analyze the response surface for values of A between 10 and 20 and values of B between 1 and 3.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **D-Optimal Designs** procedure window by clicking on **Design of Experiments**, then **D-Optimal Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

1 Specify the design and data storage.

- On the D-Optimal Designs window, select the **Design tab**.
- Set **N Per Block** to **6**.
- Set **Optimize the Design for this Model** to **A|A|B|B**.
- Set **Max Term Order** to **2**.
- Check the **Store Data on Spreadsheet** box.
- Set the **First Factor Variable** to **1**.
- Set the **First Expanded Variable** to **4**.
- Check **Rename Factor Variables with Factor Labels**.
- Check **Clear Existing Data**.

2 Specify the reports.

- On the D-Optimal Designs window, select the **Reports tab**.
- Set **Decimal Places** to **0**.

3 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

6-Point, 2 Factor D-Optimal Design

A	B
-1	-1
0	-1
1	-1
0	0
-1	1
1	1

Columns A and B give the design. The Determinant Analysis Section showed that the maximum was achieved on 25 of the 30 iterations. Hence, we assume that the algorithm converged to the global maximum.

Next, we add the two duplicates to the design. When only a few duplicates are available, we like to have them in the middle, so we will duplicate the two rows having zero values. We choose random numbers for the two new response values. The resulting design appears as follows.

6-Point Design with Two Duplicates

A	B
-1	-1
0	-1
0	-1
1	-1
0	0
0	0
-1	1
1	1

Next, we change the factor values back to their original scale. Factor A went from 10 to 20 and factor B went from 1 to 3. We call the two new variables A1 and B1. While we are at it, we also create other variables of the expanded design matrix. The resulting database appears as follows.

6-Point Design in Expanded Form

A	B	A1	B1	A1_B1	A1_A1	B1_B1
-1	-1	10	1	10	100	1
0	-1	15	1	15	225	1
0	-1	15	1	15	225	1
1	-1	20	1	20	400	1
0	0	15	2	30	225	4
0	0	15	2	30	225	4
-1	1	10	3	30	100	9
1	1	20	3	60	400	9

We could continue this exercise by running these data through a multiple regression program, paying particular attention to the Multicollinearity Section and the Eigenvalues of Centered Correlations Section. When we did this, we found that multicollinearity seemed to be a problem in the original scale, but not in the -1 to 1 scale used by the D-optimal algorithm.

Example 3 – Three Factors with Blocking

This section presents an example of how to generate and analyze a D-optimal design involving three factors with blocking.

Suppose we want to study the effect of three quantitative factor variables (A, B, and C) on a response variable. There is reason to believe that a second-order response surface design will work well. A full replication of this design requires twenty-seven experimental units. The manufacturing process that we are studying produces items in batches of four at a time. Because of this and the limited budget available for this study, we decide to use three batches (which we will call ‘Blocks’) of four points each.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **D-Optimal Designs** procedure window by clicking on **Design of Experiments**, then **D-Optimal Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

1 Specify the design and data storage.

- On the D-Optimal Designs window, select the **Design tab**.
- Set **N Per Block** to **4,4,4**.
- Set **Optimize the Design for this Model** to **A|B|C A*A B*B C*C**.
- Set **Max Term Order** to **2**.
- Set **Max Iterations** to **100**.
- Set **Inclusion Points** to **45**. This is approximately $(3)(3)(3)(3)/2$ which is the number of blocks times the product of the number of levels in each factor, all divided by two.
- Set **Removal Points** to **11**. This is one less than the total number of points desired.
- Check the **Store Data on Spreadsheet** box.
- Set the **First Factor Column** to **1**.
- Check **Rename Factor Variables with Factor Labels**.
- Check **Clear Existing Data**.

2 Specify the reports.

- On the D-Optimal Designs window, select the **Reports tab**.
- Set **Decimal Places** to **0**.

3 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

12-Point, 3 Factor D-Optimal Design with Blocking

A	B	C	Blocks
-1	-1	-1	2
1	-1	-1	1
-1	0	-1	3
-1	1	-1	1
1	1	-1	2
0	-1	0	3
-1	0	0	2
-1	-1	1	1
1	-1	1	3
0	0	1	2
-1	1	1	3
1	1	1	1

Columns A, B, C, and Blocks give the design. The Determinant Analysis Section showed that the maximum was achieved on 12 of the 100 iterations. Hence, we assume that the algorithm converged to the global maximum.

Example 4 – Adding Points to an Existing Design

This section presents an example of how to augment additional points to an existing design.

Suppose a standard three factor design has been run. Each factor has two levels. The design was blocked into two blocks of four points each. The design values are contained in the DOPT3.S0 database. This design allows only first-order (linear) terms to be fit.

Suppose that you wish to add more points to the design so that a second-order response surface may be fit. Specifically, suppose you want to add one more block of four points to extend the model from first to second order. What four points should be added?

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **D-Optimal Designs** procedure window by clicking on **Design of Experiments**, then **D-Optimal Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

1 Open DOPT3.S0.

- From the File menu of the PASS Spreadsheet window, select **Open**.
- Select the **Data** directory.
- Select the file **DOPT3.S0**.
- Click the **Ok** button.

2 Specify the design and data storage.

- On the D-Optimal Designs window, select the **Design tab**.
- Set **N Per Block** to **4,4,4**.
- Set **Input Columns (Candidate and Forced)** to **A-C**.
- Set **Forced Points** to **8**.
- Set **Optimize the Design for this Model** to **A|B|C A*A B*B C*C**.
- Set **Max Term Order** to **2**.
- Set **Max Iterations** to **30**.
- Set **Inclusion Points** to **5**.
- Set **Removal Points** to **5**.
- Check the **Store Data on Spreadsheet** box.
- Set the **First Factor Column** to **5**.
- Check **Rename Factor Columns with Factor Labels**.
- Check **Clear Existing Data**.

3 Specify the reports.

- On the D-Optimal Designs window, select the **Reports tab**.
- Set **Decimal Places** to **0**.

4 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Augmented D-Optimal Design with Blocking

A	B	C	Blocks
-1	-1	-1	1
1	1	-1	1
1	-1	1	1
-1	1	1	1
1	-1	-1	2
-1	1	-1	2
-1	-1	1	2
1	1	1	2
-1	0	-1	3
0	1	-1	3
1	-1	0	3
0	0	0	3

Variables A, B, C, and Blocks give the design. The new block is shown as the last four rows of the design.


The Determinant Analysis Section showed that the maximum was achieved on 9 of the 30 iterations. Hence, we assume that the algorithm converged to the global maximum.

Example 5 – Mixture Design

This section presents an example of how to generate a mixture design. Mixture designs are useful in situations in which the factors are constrained to sum to a total. The interest is in the proportions of each factor, not the absolute amounts. For example, the proportions of the components of a chemical solution must sum to one.

Suppose that you wish to design a first-order mixture experiment for a chemical that has three components (which we will label as A, B, and C). In this case, you will not code the factor levels from -1 to 1. Rather, the factor levels will be coded from zero to one. Because of this constraint, the intercept will not be fit in this model.

In this particular case, we will constrain the design space by only entering certain points in the list of candidate points. The candidate points are contained in the database named DOPT_MIXED.S0. The following plots (not generated using *PASS*) show the design space for each pair of factors. Remember that these factors are constrained so that the missing factor is equal to one minus the sum of the other two. Hence, if A is 0.7 and B is 0.2, then C must be 0.1.



The task for the algorithm is to pick the ten best points from the thirteen that are shown here.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **D-Optimal Designs** procedure window by clicking on **Design of Experiments**, then **D-Optimal Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example5** from the Template tab on the procedure window.

1 Open DOPT_MIXED.S0.

- From the File menu of the PASS Spreadsheet window, select **Open**.
- Select the **Data** directory.
- Select the file **DOPT_MIXED.S0**.
- Click the **Ok** button.

2 Specify the design and data storage.

- On the D-Optimal Designs window, select the **Design tab**.
- Set **N Per Block** to **10**.
- Set **Input Columns (Candidate and Forced)** to **1-3**.
- Set **Forced Points** to **0**.
- Set **Optimize the Design for this Model** to **A|B|C**.
- Set **Max Term Order** to **2**.
- Set **Max Iterations** to **30**.
- Set **Inclusion Points** to **5**.
- Set **Removal Points** to **5**.
- **Remove the check** from the **Include Intercept** check box.
- Check the **Store Data on Spreadsheet** box.
- Set the **First Factor Column** to **4**.
- Check **Rename Factor Columns with Factor Labels**.
- Check **Clear Existing Data**.

3 Specify the reports.

- On the D-Optimal Designs window, select the **Reports tab**.
- Set **Decimal Places** to **4**.

4 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Mixture Design


Original Row	Factors		
A	B	C	
1	0.7000	0.1000	0.2000
2	0.2000	0.6000	0.2000
3	0.7000	0.2000	0.1000
4	0.2000	0.2000	0.6000
5	0.3000	0.6000	0.1000
6	0.3000	0.1000	0.6000
8	0.2000	0.4000	0.4000
9	0.5000	0.1000	0.4000
11	0.5000	0.4000	0.1000
13	0.4000	0.3000	0.3000

Columns A, B, and C give the design. The original row from the candidate list is shown as the first column of the report.

The Determinant Analysis Section showed that the maximum was achieved on 30 of the 30 iterations. Hence, we assume that the algorithm converged to the global maximum.

In order to visually analyze the design, we generate the scatter plots (not generated using *PASS*) for each pair of variables in the design.

Plot of Design



It is interesting to compare these plots with those produced earlier to see which points were kept by the algorithm.

Example 6 – Qualitative Factors

This section presents an example of how to design an experiment with qualitative and quantitative factors.

Suppose your experimental situation involves two quantitative variables, A and B, and a qualitative variable C that has five possible levels. You want to fit a second-order response surface to the quantitative variables. Also, you want to fit all two-way interactions among these factors. You have budget for an 18-point design (you will add four duplicates later).

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **D-Optimal Designs** procedure window by clicking on **Design of Experiments**, then **D-Optimal Designs**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example6** from the Template tab on the procedure window.

1 Specify the data storage.

- On the D-Optimal Designs window, select the **Design tab**.
- Set **N Per Block** to **18**.
- Set **Optimize the Design for this Model** to **A|B|C A*A B*B**.
- Set **Max Term Order** to **2**.
- Set **Qualitative Factors and Levels** to **C(5)**.
- Set **Max Iterations** to **30**.
- Set **Inclusion Points** to **20**.
- Set **Removal Points** to **18**.
- **Remove the check** from the **Include Intercept** check box.
- Check the **Store Data on Spreadsheet** box.
- Set the **First Factor Column** to **1**.
- Check **Rename Factor Columns with Factor Labels**.
- Check **Clear Existing Data**.

2 Specify the reports.

- On the D-Optimal Designs window, select the **Reports tab**.
- Set **Decimal Places** to **0**.

3 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Design with Qualitative Factors

Original Row	Factors		
	A	B	C
1	-1	-1	1
3	1	-1	1
5	0	0	1
9	1	1	1
11	0	-1	2
16	-1	1	2
18	1	1	2
19	-1	-1	3
21	1	-1	3
25	-1	1	3
27	1	1	3
28	-1	-1	4
33	1	0	4
34	-1	1	4
38	0	-1	5
40	-1	0	5
42	1	0	5
44	0	1	5

Columns A, B, and C give the design. Notice that column C simply gives the level for factor C—it was not rescaled. Also note that the levels of factor C are numbered arbitrarily. This means that only the pattern is important, not the particular level. For example, in this solution, there are only three level 2's and three level 4's. In the next solution, there might be three level 3's and three level 4's.

The Determinant Analysis Section showed that the maximum was achieved on 5 of the 30 iterations. Hence, we assume that the algorithm converged to the global maximum.

Chapter 889

Design Generator

Introduction

This program generates factorial, repeated measures, and split-plots designs with up to ten factors. The design can be placed in the current spreadsheet.

An introduction to experimental design is presented in Chapter 881 on Two-Level Factorial Designs and will not be repeated here.

Crossed Factors

Two factors are *crossed* if all levels of one factor occur with each level of the second factor. No distinction needs to be made as to whether a factor is random or fixed. Factorial and randomized block designs are examples of designs that contain crossed factors.

Nested Factors

In the repeated measures and split-plot designs, at least one of the factors is nested in another factor. A factor is *nested* when all levels of this factor do not occur with each level of another factor. For example, suppose a study is being made to compare the heart rate of males and females. Five males and five females are selected. One factor in the study would be gender with two levels: male and female. Another factor would be individual with ten levels: P1, P2, ..., and P10. Since five of the ten individuals are in the males group and the other five individuals are in the females group, individuals are nested within gender.

The basic structure of *repeated measures* and *split-plot* designs is identical. The difference between the two is in the way the factor levels are assigned within the individual factor. Consider an exercise study in which heart rate readings are to be made on an individual at five different points in time. If the amounts of exercise is assigned at random before each reading, the design is a split plot. If the amounts of exercise follow the same pattern for each individual, the design is a repeated measures.

Procedure Options

This section describes the options available in this procedure.

Design Tab

This panel specifies the parameters that will be used to create the design values.

Experimental Setup

Factor (1 to 12) Values

The values used to represent the rows are specified here. These values may be letters, digits, words, or numbers. The list is delimited by blanks or commas. The number of levels of a factor corresponds to the number of values that are listed here.

To specify a nested factor, use the word Nested followed by the number of levels within a group. For example, entering 'Nested 4' signifies a design in which four individuals are placed in each group. The number of groups is found by crossing the factors before the nested factor.

An easy way to replicate a design is to specify a nested factor as the last factor with the number of replicates specified as the number of levels.

Data Storage to Spreadsheet

Store Data on Spreadsheet

Check this box to generate the design data on the spreadsheet. The spreadsheet data will be identical to the design data generated on the output window.

Store First Factor In

The first factor is stored in this column. Each additional factor that is specified is stored in the columns immediately to the right of this column. A factor is specified when values are entered into its Factor Values box.

Warning: The program fills these variables with data, so any previous data will be replaced.

Template Tab

The options on this panel allow various sets of options to be loaded (File menu: Load Template) or stored (File menu: Save Template). A template file contains all the settings for this procedure.

Specify the Template File Name

File Name

Designate the name of the template file either to be loaded or stored.

Select a Template to Load or Save

Template Files

A list of previously stored template files for this procedure.

Template Id's

A list of the Template Id's of the corresponding files. This id value is loaded in the box at the bottom of the panel.

Example 1 – Three-by-Four Factorial Design with Three Replicates

This section presents an example of how to degenerate a three-by-four factorial design with three replicates per treatment combination. To run this example, take the following steps. **CAUTION: since the purpose of this routine is to generate data, you should always begin with an empty spreadsheet.**

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Design Generator** procedure window by clicking on **Design of Experiments**, then **Design Generator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Design Generator window, select the **Design** tab.
- Enter **1 2 3** in the **Factor 1 Values (A)** box.
- Enter **1 2 3 4** in the **Factor 2 Values (B)** box.
- Enter **Nested 3** in the **Factor 3 Values (C)** box.
- Check the **Store Data on Spreadsheet** box.
- Enter **1** in the **First Factor Column** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Three-by-Four Design with Three Replicates

Experimental Design

Row	Factor 1	Factor 2	Factor 3
1	1	1	1
2	1	1	2
3	1	1	3
4	1	2	4
5	1	2	5
6	1	2	6
7	1	3	7
8	1	3	8
9	1	3	9
10	1	4	10
11	1	4	11
12	1	4	12
13	2	1	13
14	2	1	14
15	2	1	15
16	2	2	16
17	2	2	17
18	2	2	18
19	2	3	19
20	2	3	20
21	2	3	21
22	2	4	22
23	2	4	23
24	2	4	24
25	3	1	25
26	3	1	26
27	3	1	27
28	3	2	28
29	3	2	29
30	3	2	30
31	3	3	31
32	3	3	32
33	3	3	33
34	3	4	34
35	3	4	35
36	3	4	36

The values were also produced on the spreadsheet.

These values are also generated on the spreadsheet.

Example 2 – Randomized Block Design

This section presents an example of how to degenerate a randomized block design with three blocks and four treatments. To run this example, take the following steps. **CAUTION: since the purpose of this routine is to generate data, you should always begin with an empty spreadsheet.**

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Design Generator** procedure window by clicking on **Design of Experiments**, then **Design Generator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Design Generator window, select the **Design tab**.
- Enter **1 2 3** in the **Factor 1 Values (A)** box.
- Enter **A B C D** in the **Factor 2 Values (B)** box.
- Make sure that the **Factor 3 Values (C)** box is blank.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Randomized Block Design

Experimental Design

Row	Factor 1	Factor 2
1	1	A
2	1	B
3	1	C
4	1	D
5	2	A
6	2	B
7	2	C
8	2	D
9	3	A
10	3	B
11	3	C
12	3	D

It is important to remember that when you use this design, you must randomly assign treatments to the four letters and randomly assign the physical blocks to the three block numbers.

Example 3 – Repeated Measures Design

This section presents an example of how to degenerate a repeated measures design with three groups, two individuals per group, and two treatments which we will label 'Pre' and 'Post.' To run this example, take the following steps. **CAUTION: since the purpose of this routine is to generate data, you should always begin with an empty spreadsheet.**

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Design Generator** procedure window by clicking on **Design of Experiments**, then **Design Generator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

1 Specify the design parameters.

- On the Design Generator window, select the **Design tab**.
- Enter **1 2 3** in the **Factor 1 Values (A)** box.
- Enter **Nested 2** in the **Factor 2 Values (B)** box.
- Enter **Pre Post** in the **Factor 3 Values (C)** box.

2 Run the procedure.

- From the Run menu, select **Run Procedure**. Alternatively, just click the Run button (the left-most button on the button bar at the top).

Repeated Measures Design

Experimental Design

Row	Factor 1	Factor 2	Factor 3
1	1	1	Pre
2	1	1	Post
3	1	2	Pre
4	1	2	Post
5	2	3	Pre
6	2	3	Post
7	2	4	Pre
8	2	4	Post
9	3	5	Pre
10	3	5	Post
11	3	6	Pre
12	3	6	Post

Chapter 900

Chi-Square Effect Size Estimator

Introduction

This procedure calculates the effect size of the Chi-square test for use in power and sample size calculations. Based on your input, the procedure provides effect size estimates for Chi-square goodness-of-fit tests and for Chi-square tests of independence.

The *Chi-square test* is often used to test whether sets of frequencies or proportions follow certain patterns. The two most common cases are in tests of goodness of fit and tests of independence in contingency tables.

The *Chi-square goodness-of-fit* test is used to test whether a set of data follows a particular distribution. For example, you might want to test whether a set of data comes from the normal distribution.

The *Chi-square test for independence* in a contingency table is another common application of this test. Here individuals (people, animals, or things) are classified by two (nominal or ordinal) classification variables into a two-way contingency table. This table contains the counts of the number of individuals in each combination of the row categories and column categories. The Chi-square test determines if there is dependence (association) between the two classification variables.

Effect Size

For each cell of a table containing m cells, there are two proportions considered: one specified by a null hypothesis and the other specified by the alternative hypothesis. Usually, the proportions specified by the alternative hypothesis are those occurring in the data. Define p_{0i} to be the proportion in cell i given by the null hypothesis and p_{1i} to be the proportion in cell i according to the alternative hypothesis. The effect size, W , is calculated using the following formula

$$W = \sqrt{\sum_{i=1}^m \frac{(p_{0i} - p_{1i})^2}{p_{0i}}}.$$

900-2 Chi-Square Effect Size Estimator

The formula for computing the Chi-square value, χ^2 , is

$$\begin{aligned}\chi^2 &= \sum_{i=1}^m \frac{(O_i - E_i)^2}{E_i} \\ &= N \sum_{i=1}^m \frac{(p_{0i} - p_{1i})^2}{p_{0i}},\end{aligned}$$

where N is the total count in all the cells. Hence, the relationship between W and χ^2 is

$$\chi^2 = NW^2$$

or

$$W = \sqrt{\frac{\chi^2}{N}}$$

Contingency Table Tab

This window allows you to enter up to an eight-by-eight contingency table. You can enter percentages or counts. If you enter counts, the Chi-Square and Prob Level values are correct and may be used to test the independence of the row and column variables. If you enter percentages, you should ignore the Chi-Square and Prob Level values.

Note that if you are entering percentages, it does not matter whether you enter table percentages or row (or column) percentages as long as you are consistent.

Example

Suppose you are planning a survey with the primary purpose of testing whether marital status is related to gender. You decide to adopt four marital status categories: never married, married, divorced, widowed. In the population you are studying, previous studies have found the following percentages in each of these categories:

Never Married	27%
Married	39%
Divorced	23%
Widowed	11%

You decide that you want to calculate the effect size when the individual percentages for males and females are

<u>Gender</u>	<u>Male</u>	<u>Female</u>
Never Married	22%	32%
Married	46%	33%
Divorced	22%	24%
Widowed	10%	11%

To complete this example, you would load the Chi-Square Effect Size Estimator procedure from the PASS-Other menu and enter “22 46 22 10” across the top row and “32 33 24 11” across the next row. The value of W turns out to be 0.143626.

Note that even though a Chi-square value (4.13) and probability level (0.248) are displayed, you would ignore them since you have entered percentages, not counts, into the table. If you had entered counts, these results could be used to test the hypothesis of independence.

Multinomial Test Tab

This window allows you to enter a multinomial table with up to fourteen cells. You can enter percentages or counts. If you enter counts, the Chi-Square and Prob Level values are correct and may be used to test the statistical significance of the table. If you enter percentages, you should ignore the Chi-Square and Prob Level values.

Note that if you are using the window to perform a goodness-of-fit test on a set of data, you will need to adjust the degrees of freedom for the number of parameters you are estimating. For example, if you are testing whether the data are distributed normally and you estimate the mean and standard deviation from the data, you will need to reduce the degrees of freedom by two.

Example

Suppose you are going to use the Chi-square goodness-of-fit statistic calculated from a multinomial table to test whether a set of exponential data follow the normal distribution. That is, you want to find a reasonable effect size for comparing exponentially distributed data to the normal distribution.

You decide to divide the data into five groups: 5 or less, 5-10, 10-15, 15-20, 20+

Using tables for the normal and exponential distributions, you find that the probabilities for each group are

900-4 Chi-Square Effect Size Estimator

<u>Category</u>	<u>Normal</u>	<u>Exponential</u>
5 or Less	11%	39%
5 to 10	20%	26%
10 to 15	38%	18%
15 to 20	20%	11%
Above 20	11%	6%

To complete this example, you would set the Chi-Square Effect Size Estimator procedure to the Multinomial Test tab and enter “11 20 38 20 11” down the first column and “39 26 18 11 6” down the second column. The calculated value of W is 0.948271. You would enter this value into the Effect Size option of the Chi-Square Test window in **PASS** to determine the necessary sample size.

Chapter 905

Standard Deviation Estimator

Introduction

Even though it is not of primary interest, an estimate of the *standard deviation (SD)* is needed when calculating the power or sample size of an experiment involving one or more means. Finding such an estimate is difficult not only because the estimate is required before the data are available, but also because the interpretation of the standard deviation is vague and our experience with it may be low. How do you estimate a quantity without data and without a clear understand of what the quantity is? This section will acquaint you with the standard deviation and offer several ways to obtain a rough estimate of it before the experiment begins using the Standard Deviation Estimator.

The Standard Deviation Estimator can also be used to calculate the standard deviation of the means, a quantity used in estimating sample sizes in analysis of variance designs.

Understanding the Standard Deviation

It is difficult to understand the standard deviation solely from the standard deviation formula. There are two general interpretations that can be useful in understanding the standard deviation.

1. The standard deviation may be thought of as the average difference between an observation and the mean, ignoring the sign.
2. The standard deviation may be thought of as the average difference between any two data values, ignoring the sign.

The population standard deviation is calculated using the formula:

$$\sigma = \sqrt{\frac{\sum_{i=1}^N (X_i - \mu)^2}{N}}$$

where N is the number of items in the population, X is the variable being measured, and μ is the mean of X . This formula indicates that the standard deviation is the square root of an average.

905-2 Standard Deviation Estimator

This average is the average of the squared differences between each value and the mean. The differences are squared to remove the sign so that negative values will not cancel out positive values. After summing up these squared differences and dividing by N , the square root is taken to give the result in the original scale. That is, the standard deviation can be thought of as the average difference between the data values and their mean (the terms mean and average are used interchangeably).

Example

Consider the following two sets of numbers

A: 1, 5, 9

B: 4, 5, 6

Both sets have the same mean of 5. However, their standard deviations are quite different. Subtracting the mean and squaring the three items in each set results in

Set A

$$(1-5)(1-5) = 16$$

$$(5-5)(5-5) = 0$$

$$(9-5)(9-5) = 16$$

$$\text{Sum} = 32$$

$$SD_A = \sqrt{\frac{32}{3}} = 3.266$$

Set B

$$(4-5)(4-5) = 1$$

$$(5-5)(5-5) = 0$$

$$(6-5)(6-5) = 1$$

$$\text{Sum} = 2$$

$$SD_B = \sqrt{\frac{2}{3}} = 0.8165$$

The standard deviations show that the data in set A vary more than the data in set B.

Divide by N or N-1?

Note in the example above that we are dividing by N , not $N-1$ as is usually seen in standard deviation calculations. When the standard deviation is computed using all values in the population, N is used as the divisor. However, when the standard deviation is calculated from a sample, $N-1$ is used as the divisor. We stress that the results for a sample by using the lower-case n and naming the *sample standard deviation* S . The value of S is computed from a sample of n values using the formula

$$S = \sqrt{\frac{\sum_{i=1}^n (X_i - \bar{X})^2}{n-1}}$$

The $n-1$ is used instead of n to correct for bias that statisticians have discovered. That is, over the long run, dividing by $n-1$ provides a better estimate of the true standard deviation than does dividing by n . Although we divide by $n-1$ rather than n for sample standard deviations, we recommend that for purposes of interpretation, the divisor is assumed to be n , so that the operation can be thought of as computing an average.

Average Absolute Deviation

If we were to devise a measure of variability with no previous experience, we might first consider the average absolute deviation (AD), sometimes called the mean absolute deviation, or MAD, which is computed by forming the deviations from the mean, taking their absolute values, and computing their average. The absolute value is applied to remove the negative signs, which, in turn, avoids the cancellation of values when the average is taken. The formula for AD is

$$AD = \frac{\sum_{i=1}^N |X_i - M|}{N}$$

This simple average of absolute deviations is much easier to understand, but is very difficult to work with mathematically. Opposingly, the standard deviation is more difficult to interpret directly, but it can be worked with mathematically in statistical problems. The ability to work with the standard deviation mathematically outweighs its deficiency in interpretation. Hence, we generally use the standard deviation rather than the average absolute deviation in practice.

Comparing Average Absolute Deviation and Standard Deviation

Fortunately, the average absolute deviation and the standard deviation are usually close in value. Mathematically, it can be shown that AD is always less than or equal to SD . A small simulation study is summarized below. It shows the relationship between AD and SD for data generated from various distributions.

<u>Distribution</u>	<u>Percent SD > AD</u>	<u>Characteristics</u>
Uniform	15%	Level
Normal	20%	Bell-Shaped
Gamma(5)	30%	Moderately Skewed Right
Gamma(5)^2	45%	Extremely Skewed Right

These distributions were selected for study because they represent a wide range of possibilities. The table shows that, for typical datasets, the standard deviation is from 15 to 30 percent larger

905-4 Standard Deviation Estimator

than the average absolute deviation. And in the case of the normal distribution, the SD is about 20% higher than AD .

Hence, for planning purposes, you can think of the standard deviation as an inflated version of the average absolute deviation.

Example

In our example, we can compute AD for datasets A and B as follows.

Set A

$AD_A = (4+0+4)/3 = 8/3 = 2.667$. Recall that $SD_A = 3.266$.

Set B

$AD_B = (1+0+1)/3 = 2/3 = 0.667$. Recall that $SD_B = 0.8165$.

We see that the values are similar. The degree of difference is likely within the error that we would expect during the planning phase.

Standard Deviation as the Average Difference between Values

The above discussion and formula have pointed out that the standard deviation may be thought of as an average deviation from the mean. In this section, a second interpretation of the standard deviation will be given.

We can manipulate the formula for the sum of squared deviations to show that

$$\sum_{i=1}^N (X_i - M)^2 = \frac{\sum_{i=1}^{N-1} \sum_{j=i+1}^N (X_i - X_j)^2}{N}$$

This formula shows that the squared deviations from the mean are proportional to the squared deviations of each observation from every other observation. Note that the mean is not involved in the expression on the right.

Using the above relationship, the standard deviation may be calculated using the formula

$$SD = \sqrt{\frac{\sum_{i=1}^{N-1} \sum_{j=i+1}^N (X_i - X_j)^2 / N}{N}}$$

Example

Consider again our simple example of sets A and B. Applying this operation to the three possible pairs of the data in set A (1, 5, 9) gives

$$(1-5)(1-5) = 16$$

$$(1-9)(1-9) = 64$$

$$(5-9)(5-9) = 16$$

The sum is 96. Dividing 96 by 3 (the number of pairs) again yields 32. Hence, the standard deviation is computed as $SD = \text{SQRT}(96/3) = 3.266$ (which matches the previous result).

Likewise, for set B (4, 5, 6), the formula results in

$$(4-5)(4-5) = 1$$

$$(4-6)(4-6) = 4$$

$$(5-6)(5-6) = 1$$

so that $SD = \text{SQRT}(6/3) = 0.8165$.

Estimating the Standard Deviation

Our task is to find a rough estimate of the standard deviation. Several possible methods are available in the Standard Deviation Estimator procedure which may be loaded from the PASS-Other menu. *PASS* provides a panel that implements each of these methods for you.

Data Tab – Standard Deviation from Data Values

One method of estimating the standard deviation is to put in a typical set of values and calculate the standard deviation.

This window is also used when you need the standard deviation of a set of hypothesized means in an analysis of variance sample size study.

Pros and Cons of This Method

This method lets you experiment with several different data values. It lets you determine the influence of different data configurations on the standard deviation. In so doing, you can come up with a likely range of SD values.

However, investigators tend to pick trial numbers that are closer to the mean and more uniform than will result in practice. This results in SD's that are underestimated. If you use this method, you should be careful that your range of possible SD values is wide enough to be accurate.

Example 1 – SD for a Set of Values

As an example, suppose that you decide that the following values represent a typical set of data that you would anticipate for one group of individuals:

10, 12, 14, 10, 11, 10, 12, 13, 9, 13, 15, 11

To calculate the appropriate standard deviation, do the following:

1. Load the **Standard Deviation Estimator** window and click on the **Data tab**.
2. The order that the data are entered in does not matter. However, to show the use of the Counts column, we count up the number of times each value occurs. The values and their frequency counts are then entered into the **Values** and **Counts** columns. The data entry goes as follows:

9	1
10	3
11	2
12	2
13	2
14	1
15	1
3. Check the **Use N-1 as divisor** box. (We use the N-1 divisor when estimating sigma from a set of data.)
4. Press the **Calculate N, Mean, SD from Values** button. The standard deviation is 1.825742. We might round this value up to 2.0 for planning purposes.

Example 2 – SD for a Set of Means

In this example, we will show you how to obtain the standard deviation of a set of hypothesized means. Care must be taken that you select the correct divisor— N , not $N-1$.

In this example, a researcher is studying the influence of a drug on heart rate. He estimates that the average heart rate of his group without the drug is 80. His experimental design will apply three different doses. The first dose is expected to lower the heart rate by 10%, the second by 20%, and the third by 30%. Hence, the hypothesized means for the four groups are 80, $80(0.9) = 72$, $80(0.8) = 64$, and $80(0.7) = 56$.

To calculate the appropriate standard deviation, do the following:

1. Load the **Standard Deviation Estimator** window and click on the **Data tab**.
2. Enter the four means into the **Values** column. The Counts column is left blank.

80
72
64
56
3. Make sure the **Use N-1 as divisor** box is not checked since we want the population standard deviation.
4. Press the **Calculate N, Mean, SD from Values** button. The standard deviation of the means is 8.944272.

Standard Error Tab – Standard Deviation from Standard Error

If the value of the standard error of the mean is available from another experiment, it may be used to estimate the standard deviation. The formula estimating the standard deviation from the standard error is

$$SD = SE\sqrt{N}$$

where N is the sample size.

Pros and Cons of This Method

This method is only useful when you have a standard error value available.

Example

To calculate the appropriate standard deviation when a previous study of 23 individuals had a standard error of the mean of 2.7984, do the following:

1. Load the **Standard Deviation Estimator** window and click on the **Standard Error tab**.
2. Enter **23** for **N**.
3. Enter **2.7984** for **Standard Error**.
4. Press the **Calculate Standard Deviation** button. The standard deviation is 1.825742. We might round this value up to 2.0 for planning purposes.

Range Tab – Standard Deviation from Population or Sample Data Range

There are two cases in which the range may be used to estimate the standard deviation. In the first case, the sample size and data range may be available from a previous study. In the second case, a reasonable estimate of the population range may be obtainable. This window allows you to estimate SD in both of these situations.

The basic formula for estimating the standard deviation from the range is

$$SD = \frac{Range}{C}$$

where C is determined by the situation.

Determining C

If the population range can be established, it may be used to estimate sigma by dividing by an appropriate constant. To determine an appropriate value of the constant, statisticians use the fact that most of the data is contained within three standard deviations of the mean—so they set C to six. However, consultants have found that for some reasons (e.g., understated range, or non-normal population) dividing by six tends to understate the standard deviation. So they divide by five or even four. Dividing by a smaller number increases the estimated standard deviation. Our recommendation is to divide by four. To use this method, enter the divisor (4, 5, or 6) in the C (Divisor) box.

905-8 Standard Deviation Estimator

If the data range is available from a previous study, the constant C may be calculated as the median of the distribution of the range for that sample range. This distribution assumes that the data themselves are normally distributed. The median of the distribution of the range is calculated in NCSS using numerical methods, and similar calculations can be made using the Probability Calculator. To use this method, enter the data range and the sample size in the lower region of the Range tab.

Pros and Cons of This Method

The range is a poor substitute for having the standard deviation. This method should be used as a 'last resort'.

Example 1 – Population Range 'Known'

To calculate an estimate of the standard deviation when the population range is known to be 150, do the following:

1. Load the **Standard Deviation Estimator** window and click on the **Range tab**.
2. Enter **150** for **Range (of Population)** in the upper region of the tab.
3. Enter **4** for **C (Divisor)**.
4. Press the **Calculate Standard Deviation** button. The estimate of the standard deviation is 37.5. The value of C used is 4.

Example 2 – Previous Sample Available

To calculate an estimate of the standard deviation when a previous study of 20 animals had a minimum value of 15.3 and a maximum of 18.7, do the following:

1. Load the **Standard Deviation Estimator** window and click on the **Range tab**.
2. Enter **20** for **N (Sample Size)** in the lower region of the tab.
3. Enter **3.4** for **Range (of Sample Data)** in the lower region of the tab. This is 18.7 minus 15.3.
4. Press the **Calculate Standard Deviation** button. The estimate of the standard deviation is 0.92243. The value of C used is 3.685916.

Percentiles Tab – Standard Deviation from Percentiles

If you are willing to assume that the population values are normally distributed (bell-shaped), you can use the values of two percentiles to estimate the standard deviation.

The basic formula for estimating the standard deviation from two percentiles is

$$SD = \left| \frac{X_2 - X_1}{Z\left(\frac{P_2}{100}\right) - Z\left(\frac{P_1}{100}\right)} \right|$$

where X_1 and X_2 are the two percentiles, P_1 and P_2 are the two percentages, and $Z(P)$ is the standard normal deviate that has a tail area of P to the left.

Pros and Cons of This Method

This method works well if you have two accurate percentiles available and the underlying distribution of the data is normal.

Example

To calculate an estimate of the standard deviation when you know that the 25th percentile of the population is 80.25, the 75th percentile of the population is 116.38, and that the population is normally distributed, do the following:

1. Load the **Standard Deviation Estimator** window and click on the **Percentiles tab**.
2. Enter **25** for **Percentage 1**.
3. Enter **80.25** for **Percentile Value 1**.
4. Enter **75** for **Percentage 2**.
5. Enter **116.38** for **Percentile Value 2**.
6. Press the **Calculate Standard Deviation** button. The estimate of the standard deviation is 26.78321.

COV Tab – Standard Deviation from Coefficient of Variation

The coefficient of variation (*COV*) is equal to SD divided by the mean. Hence, if you know the coefficient of variation and the mean, you can estimate SD.

Note that t-tests and the analysis of variance assume that the standard deviations are equal for all groups. If the standard deviations are proportional to their group means, you should use a data transformation (such as the square root or the logarithm) to make the standard deviations equal.

The basic formula for estimating the standard deviation from the *COV* is

$$SD = (COV)(Mean)$$

Pros and Cons of This Method

This method works well if you have estimates of the mean and *COV*.

Example

To calculate an estimate of the standard deviation when you know that the mean is 127 and the *COV* is 0.832, do the following:

1. Load the **Standard Deviation Estimator** window and click on the **COV tab**.
2. Enter **0.832** for **Coefficient of Variation (COV)**.
3. Enter **127** for **Mean**.
4. Press the **Calculate Standard Deviation** button. The estimate of the standard deviation is 105.664.

Confidence Limits Tab – Confidence Limits of the Standard Deviation

Often, you will obtain an estimate of the standard deviation from a previous study or a pilot study. Since this estimate is based on a sample, it is important to understand its precision. This can easily be calculated since the square of the sample standard deviation follows a chi-squared distribution. This confidence interval does assume that the population you are sampling from is normally distributed.

Once a confidence interval has been obtained, it would be wise to enter both values (the confidence limits) into the appropriate place in the sample size calculations to provide a range of possible sample size values (or of statistical power).

Example

Suppose a pilot study of 10 individuals yields a standard deviation of 83.21. Calculate a 95% confidence interval for the population standard deviation using these results.

1. Load the **Standard Deviation Estimator** window and click on the **Confidence Limits** tab.
2. Enter **10** for **N**.
3. Enter **83.21** for **Standard Deviation**.
4. Enter **0.05** for **Alpha**.
5. Press the **Calculate Confidence Limits** button. The confidence limits are 57.23477 and 151.909.

Chapter 910

Odds Ratio and Proportions Estimator

Introduction

The Odds Ratio tab of this procedure calculates any one of three parameters, odds ratio, p_1 , or p_2 , from the other two parameters. Note that p_1 and p_2 are the proportions in groups one and two, respectively. This provides you with a tool to study the relationship between these three parameters. This procedure is most often used when planning the sample size for a test involving two proportions. The procedure may be loaded by selecting *Odds Ratio Estimator* from the *PASS-Other* menu.

As an expanded version of the Odds Ratio tab, the Proportions tab calculates p_1 , p_2 , the difference, ratio, odds ratio, or $\text{Ln}(\text{OR})$ from various combinations of these parameters.

When planning studies involving two proportions, two parameters, p_1 and p_2 , need to be specified. At times, it may be difficult to propose a value for p_2 . In these cases, it might be easier to propose a value for the *odds ratio (OR)*.

The odds of obtaining the response of interest in group 1 are $p_1 / (1 - p_1)$ and the odds of obtaining the response in group 2 are $p_2 / (1 - p_2)$. The ratio of these odds, called the odds ratio, is defined as

$$\begin{aligned} OR &= \frac{p_2 / (1 - p_2)}{p_1 / (1 - p_1)} \\ &= \frac{p_2(1 - p_1)}{p_1(1 - p_2)} \end{aligned}$$

To understand better how to interpret an odds ratio, consider the following example. Suppose the proportion dying from a particular disease during the first five years is 80%. The odds of dying are thus $0.8 / 0.2 = 4.0$. Suppose a treatment reduces the death rate from 80% to 60%. The odds of dying are now $0.6 / 0.4 = 1.5$. The odds ratio is $4.0 / 1.5 = 2.7$. That is, the odds of dying have been reduced by a factor of 2.67.

910-2 Odds Ratio and Proportions Estimator

The odds ratio is reversible. In this example, suppose we talk in terms of surviving instead of dying. The odds of the two groups are now $0.2 / 0.8 = 0.25$ and $0.4 / 0.6 = 0.67$. The odds ratio is $0.67 / 0.25 = 2.67$. The odds of surviving have increased by a factor of 2.67.

In some situations, it may be easier to define a meaning treatment effect in terms of the odds ratio. That is, it might be meaningful to say that a certain treatment increases the odds of survival by 50% ($OR = 1.5$) or by 200% ($OR = 2.0$). If this is the case, a value for p_2 can be calculated from p_1 and OR by solving the above equation for p_2 to find that

$$p_2 = \frac{p_1(OR)}{1 - p_1 + p_1(OR)}$$

Hence, given a value for p_1 and OR , you can calculate an appropriate value for p_2 .

Odds Ratio Tab

This window lets you calculate p_1 , p_2 , or the odds ratio (OR) from the other two parameters.

Example 1 – Solving for P1

Suppose you know that $p_2 = 0.8$ and that $OR = 4$ and you want to find the corresponding value of p_1 .

1. Load the **Odds Ratio and Proportions Estimator** procedure by selecting it from the *PASS-Other* menu.
2. Select the **Odds Ratio tab**.
3. Set **P2** equal to **0.8**.
4. Set **Odds Ratio** equal to **4**.
5. Press the **Calculate P1** button.
6. Read the result in the P1 box. The result is **0.5**.

Example 2 – Solving for P2

Suppose you know that $p_1 = 0.4$ and that $OR = 1.5$ and you want to find the corresponding value of p_2 .

1. Load the **Odds Ratio and Proportions Estimator** procedure by selecting it from the *PASS-Other* menu.
2. Select the **Odds Ratio tab**.
3. Set **P1** equal to **0.4**.
4. Set **Odds Ratio** equal to **1.5**.
5. Press the **Calculate P2** button.
6. Read the result in the P2 box. The result is **0.5**.

Example 3 – Solving for Odds Ratio

Suppose you know that $p_1 = 0.4$ and that $p_2 = 0.8$ and you want to find the corresponding value of the odds ratio.

1. Load the **Odds Ratio and Proportions Estimator** procedure by selecting it from the *PASS-Other* menu.
2. Select the **Odds Ratio tab**.
3. Set **P1** equal to **0.4**.
4. Set **P2** equal to **0.8**.
5. Press the **Calculate Odds Ratio** button.
6. Read the result in the Odds Ratio box. The result is **6**.

Proportions Tab

This window lets you calculate p_1 , p_2 , the difference, ratio, odds ratio, or Ln(OR) from various combinations of these parameters.

Example 1 – Calculating Ln(OR)

Suppose you know that $p_1 = 0.4$ and that $p_2 = 0.8$ and you want to find the corresponding value of Ln(OR).

1. Load the **Odds Ratio and Proportions Estimator** procedure by selecting it from the *PASS-Other* menu.
2. Select the **Proportions tab**.
3. Set **P1** equal to **0.4**.
4. Set **P2** equal to **0.8**.
5. Press the **P1 & P2** button.
6. Read the result in the Ln(O.R.) box. The result is **1.79176**.

910-4 Odds Ratio and Proportions Estimator

Chapter 915

Probability Calculator

Introduction

Most statisticians have a set of probability tables that they refer to in doing their statistical work. This procedure provides you with a set of electronic statistical tables that will let you look up values for various probability distributions.

To run this option, select Probability Calculator from the Other menu of the Analysis menu. A window will appear that will let you indicate which probability distribution you want to use along with various input parameters. Select the Calculate button to find and display the results.

Many of the probability distributions have two selection buttons to the left of them. The first (left) button selects the inverse probability distribution. An inverse probability distribution is in a form so that when you give it a probability, it calculates the associated critical value. The second (right) button selects the regular probability distribution which is formulated so that when you give it a critical value, it calculates the (left tail) probability.

Probability Distributions

Beta Distribution

The beta distribution is usually used because of its relationship to other distributions, such as the t and F distributions. The noncentral beta distribution function is formulated as follows:

$$\Pr(0 \leq x \leq X | A, B, L) = I_X(A, B, L) = \frac{\Gamma(A+B)}{\Gamma(A)\Gamma(B)} \sum_{k=0}^{\infty} \frac{e^{-L} L^k}{k! 2^k} \int_0^X t^{A+k-1} (1-t)^{B-1} dt$$

where

$$0 < A, 0 < B, 0 \leq L, \text{ and } 0 \leq x \leq 1$$

When the noncentrality parameter (NCP), L , is set to zero, the above formula reduces to the *standard* beta distribution, formulated as

$$\Pr(0 \leq x \leq X | A, B) = \frac{\Gamma(A+B)}{\Gamma(A)\Gamma(B)} \int_0^X t^{A-1} (1-t)^{B-1} dt$$

915-2 Probability Calculator

When the inverse distribution is selected, you supply the probability value and the program solves for X . When the regular distribution is selected, you supply X and the program solves for the cumulative (left-tail) probability.

Binomial Distribution

The binomial distribution is used to model the counts of a sequence of independent binary trials in which the probability of a success, P , is constant. The total number of trials (sample size) is N . R represents the number of successes in N trials. The probability of exactly R successes is:

$$\Pr(r = R | N, P) = \binom{N}{R} P^R (1 - P)^{N-R}$$

where

$$\binom{N}{R} = \frac{N!}{R!(N-R)!}$$

The probability of from 0 to R successes is given by:

$$\Pr(0 \leq r \leq R | N, P) = \sum_{r=0}^R \binom{N}{r} P^r (1 - P)^{N-r}$$

When the inverse distribution is selected, you supply the probability value and the program solves for R . When the regular distribution is selected, you supply R and the program solves for the cumulative (left-tail) probability.

Bivariate Normal Distribution

The bivariate normal distribution is given by the formula

$$\Pr(x < h, y < k | r) = \frac{1}{2\pi\sqrt{1-r^2}} \int_{-\infty}^h \int_{-\infty}^k \exp\left\{\frac{-x^2 + 2rxy - y^2}{2(1-r^2)}\right\} dx dy$$

where x and y follow the bivariate normal distribution with correlation coefficient r .

Chi-Square Distribution

The Chi-square distribution arises often in statistics when the normally distributed random variables are squared and added together. DF is the degrees of freedom of the estimated standard error.

The noncentral Chi-square distribution function is used in power calculations. The noncentral Chi-square distribution is calculated using the formula:

$$\Pr(0 \leq x \leq X | df, L) = \sum_{k=0}^{\infty} \frac{L^k e^{-L}}{2^k k!} P(X | df + 2k)$$

where

$$P(X|df) = \frac{1}{2^{df/2} \Gamma\left(\frac{df}{2}\right)} \int_0^X t^{df/2-1} e^{-t/2} dt$$

When the noncentrality parameter (NCP), L , is set to zero, the above formula reduces to the (central) Chi-square distribution.

When the inverse distribution is selected, you supply the probability value and the program solves for X . When the regular distribution is selected, you supply X and the program solves for the cumulative (left-tail) probability.

Correlation Coefficient Distribution

The correlation coefficient distribution is formulated as follows:

$$\Pr(r \leq R|n, \rho) = \int_{-1}^R \frac{2^{n-3}}{\pi(n-3)!} (1-\rho)^{(n-1)/2} (1-r)^{(n-4)/2} \sum_{i=0}^{\infty} \Gamma^2\left(\frac{n+i-1}{2}\right) \frac{(2\rho r)^i}{i!} dr$$

where

$$|r| < 1, |\rho| < 1, \text{ and } |R| < 1$$

When the inverse distribution is selected, you supply the probability value and the program solves for R . When the regular distribution is selected, you supply R and the program solves for the cumulative (left-tail) probability.

F Distribution

The F distribution is used in the analysis of variance and in other places where the distribution of the ratio of two variances is needed. The degrees of freedom of the numerator variance is DF1 and the degrees of freedom of the denominator variance is DF2.

The noncentral-F distribution function is used in power calculations. We calculate the noncentral-F distribution using the following relationship between the F and the beta distribution function.

$$\Pr(0 \leq f \leq F | df_1, df_2, L) = I_x\left(\frac{df_1}{2}, \frac{df_2}{2}, L\right)$$

where

$$X = \frac{F(df_1)}{F(df_1) + df_2}$$

When the noncentrality parameter (NCP), L , is set to zero, the above formula reduces to the *standard* F distribution

When the inverse distribution is selected, you supply the probability value and the program solves for F . When the regular distribution is selected, you supply F and the program solves for the cumulative (left-tail) probability.

Hotelling's T2 Distribution

Hotelling's T -Squared distribution is used in multivariate analysis. We calculate the distribution using the following relationship between the F and the T^2 distribution function.

$$\Pr(0 \leq x \leq \frac{(df - k + 1)}{k(df)} T_{k, df}^2 | k, df) = \Pr(0 \leq x \leq F_{k, df - k + 1} | k, df)$$

where k is the number of variables and df is the degrees of freedom associated with the covariance matrix. When the inverse distribution is selected, you supply the probability value and the program solves for T^2 . When the regular distribution is selected, you supply T^2 and the program solves for the cumulative (left-tail) probability.

Gamma Distribution

The Gamma distribution is formulated as follows:

$$\Pr(0 \leq g \leq G | A, B) = \frac{1}{B^A \Gamma(A)} \int_0^G x^{A-1} e^{-x/B} dx$$

where

$$\Gamma(A) = \int_0^\infty x^{A-1} e^{-x} dx$$

$$0 < A, 0 < B, \text{ and } 0 \leq G$$

When the inverse distribution is selected, you supply the probability value and the program solves for G . When the regular distribution is selected, you supply G and the program solves for the cumulative (left-tail) probability.

Hypergeometric Distribution

The hypergeometric distribution is used to model the following situation. Suppose a sample of size R is selected from a population with N items, M of which have a characteristic of interest. What is the probability that X of the items in the sample have this characteristic.

The probability of exactly X successes is:

$$\Pr(x = X | N, M, R) = \frac{\binom{M}{X} \binom{N-M}{R-X}}{\binom{N}{R}}$$

where

$$\binom{N}{R} = \frac{N!}{R!(N-R)!}$$

$$\text{Maximum}(0, R-N+M) \leq X \leq \text{Minimum}(M, R)$$

Negative Binomial Distribution

The negative binomial distribution is used to model the counts of a sequence of independent binary trials in which the probability of a success, P , is constant. The total number of trials (sample size) is N . R represents the number of successes in N trials. Unlike the binomial distribution, the sample size, N , is the variable of interest.

The question answered by the negative binomial distribution is: how many tosses of a coin (with probability of a head equal to P) is necessary to achieve R heads and X tails.

The probability of exactly R successes is:

$$\Pr(x = X | R, P) = \binom{X + R - 1}{R - 1} P^R (1 - P)^X$$

where

$$\binom{N}{R} = \frac{N!}{R!(N - R)!}$$

Normal Distribution

The normal distribution is formulated as follows:

$$\Pr(x \leq X | \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^x \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right) dx$$

When the mean is 0 and the variance is 1, we have the standard normal distribution. The regular normal distribution uses the variable X . The standard normal distribution uses the variable Z . Any normal distribution may be transformed to the standard normal distribution using the relationship:

$$z = \frac{x - \mu}{\sigma}$$

When the inverse distribution is selected, you supply the probability value and the program solves for R . When the regular distribution is selected, you supply R and the program solves for the cumulative (left-tail) probability.

Poisson Distribution

The Poisson distribution is used to model the following situation. Suppose the average number of accidents at a given intersection is 13.5 per year. What is the probability of having 2 accidents during the next half year?

The probability of exactly X occurrences with a mean occurrence rate of M is:

$$\Pr(x = X | M) = \frac{e^{-M} M^X}{X!}$$

Studentized Range Distribution

The studentized range distribution is used whenever the distribution of the ratio of a range and an independent estimate of its standard error is needed. This distribution is used quite often in multiple comparison tests run after an analysis of variance. DF is the degrees of freedom of the estimated standard error (often the degrees of freedom of the MSE). K is the number of items (means) in the sample. The distribution function is given by:

$$\Pr(0 \leq r \leq R | df, k) = \int_0^R \left(\frac{2^{-df/2+1} df^{df/2} s^{df-1}}{\Gamma\left(\frac{df}{2}\right)} \exp\left(-\frac{dfs^2}{2}\right) P(Rs|n) \right) dx$$

where $P(Rs|n)$ is the probability integral of the range.

When the inverse distribution is selected, you supply the probability value and the program solves for R . When the regular distribution is selected, you supply R and the program solves for the cumulative (left-tail) probability.

Student's t Distribution

The t distribution is used whenever the distribution of the ratio of a statistic and its standard error is needed. DF is the degrees of freedom of the estimated standard error.

The noncentral-t distribution function is used in power calculations. We calculate the noncentral-t distribution using the following relationship between the t and the beta distribution function.

$$\Pr(-\infty \leq t \leq T | df, L) = 1 - \sum_{k=0}^{\infty} e^{-L^2/2} \frac{(L^2/2)^k}{2k!} I_x\left(\frac{df}{2}, \frac{1}{2}, 0\right)$$

where

$$X = \frac{df}{df + T^2}$$

When the noncentrality parameter (NCP), L , is set to zero, the above formula reduces to the (central) Student's t distribution

When the inverse distribution is selected, you supply the probability value and the program solves for T . When the regular distribution is selected, you supply T and the program solves for the cumulative (left-tail) probability.

Weibull Distribution

The Weibull distribution is formulated as follows:

$$\Pr(t \leq T | \lambda, \gamma) = 1 - \exp\left(-(\lambda T)^\gamma\right)$$

When gamma (γ) equal to one, the distribution simplifies to the exponential distribution.

When the inverse distribution is selected, you supply the probability value and the program solves for T . When the regular distribution is selected, you supply T and the program solves for the cumulative (left-tail) probability.

Converting Summary Statistics to Raw Data

Occasionally, you will have summary statistics (mean, count, and standard deviation) but not the raw data used to create these summary statistics. Since you need the original data values, you cannot run descriptive statistics, t-tests, or AOV's on these data. This routine is useful in this case since it generates a set of raw data values with a given mean, count, and standard deviation. Although the generated values are not the same as the original values, they are good enough to use as input into statistical procedures such as analysis of variance or two-sample t-tests.

Let n represent the sample size (count), M represent the sample mean, and s represent the sample standard deviation. (Recall that the standard error is equal to the standard deviation times the square root of n , so if you have been given the standard error, use $s = SQR(n)(s.e.)$.) It is possible to find values, $X1$ and $X2$, so that a variable made up of one $X1$ and $(n-1)$ $X2$'s will have the same values of n , M , and s . The formulas to do this are:

$$X1 = M - (n-1) \frac{s}{\sqrt{n}}$$

and

$$X2 = M + \frac{s}{\sqrt{n}}$$

For example, suppose you have the following two sets of summary statistics and want to run a two-sample t-test.

	<u>Sample One</u>	<u>Sample Two</u>
Sample Size	4	5
Mean	3.2	4.5
Standard Deviation	1.4	1.8

Using the *Convert Mean and S to Data Values* option of the *Probability Calculator* and plugging in these values gives the following results:

	<u>Sample One</u>	<u>Sample Two</u>
X1	1.1	1.280062
N1	1	1
X2	3.9	5.304984
N2	3	4

These data would then be entered into a datasheet as follows:

Row	C1	C2
1	1.1	1.280062
2	3.9	5.304984
3	3.9	5.304984
4	3.9	5.304984
5		5.304984

915-8 Probability Calculator


Now, a two-sample t-test or a one-way analysis of variance could be run on these two variables. Note that since nonparametric tests, tests of assumptions, box plots, and histograms require the original data values, their output would have to be ignored, but the t-test results would be accurate.

Chapter 920

Data Simulator

Introduction

Because of mathematical intractability, it is often necessary to investigate the properties of a statistical procedure using *simulation* (or *Monte Carlo*) techniques. In power analysis, *simulation* refers to the process of generating several thousand random samples that follow a particular distribution, calculating the test statistic from each sample, and tabulating the distribution of these test statistics so that the significance level and power of the procedure may be investigated. This module creates a histogram of a specified distribution as well as a numerical summary of simulated data. By studying the histogram and the numerical summary, you can determine if the distribution has the characteristics you desire. The distribution formula can then be used in procedures that use simulation, such as the new t-test procedures. Below are examples of two distributions that were generated with this procedure.



Technical Details


A random variable's probability distribution specifies its probability over its range of values. Examples of common continuous probability distributions are the normal and uniform distributions. Unfortunately, experimental data often do not follow these common distributions, so other distributions have been proposed. One of the easiest ways to create distributions with desired characteristics is to combine simple distributions. For example, outliers may be added to a distribution by mixing it with data from a distribution with a much larger variance. Thus, to simulate normally distributed data with 5% outliers, we could generate 95% of the sample from a normal distribution with mean 100 and standard deviation 4 and then generate 5% of the sample from a normal distribution with mean 100 and standard deviation 16. Using the standard notation

920-2 Data Simulator

for the normal distribution, the composite distribution of the new random variable Y could be written as

$$Y \sim \delta(0 \leq X < 0.95)N(100,4) + \delta(0.95 \leq X \leq 1.00)N(100,16)$$

where X is a uniform random variable between 0 and 1, $\delta(z)$ is 1 or 0 depending on whether z is true or false, $N(100,4)$ is a normally distributed random variable with mean 100 and standard deviation 4, and $N(100,16)$ is a normally distributed random variable with mean 100 and standard deviation 16. The resulting distribution is shown below. Notice how the tails extend in both directions.



The procedure for generating a random variable, Y , with the mixture distribution described above is

1. Generate a uniform random number, X .
2. If X is less than 0.95, Y is created by generating a random number from the $N(100,4)$ distribution.
3. If X is greater than or equal to 0.95, Y is created by generating a random number from the $N(100,16)$ distribution.

Note that only one uniform random number and one normal random number are generated for any particular random realization from the mixture distribution.

In general, the formula for a mixture random variable, Y , which is to be generated from two or more random variables defined by their distribution function $F_i(Z_i)$ is given by

$$Y \sim \sum_{i=1}^k \delta(a_i \leq X < a_{i+1}) F_i(Z_i), \quad a_1 = 0 < a_2 < \dots < a_{K+1} = 1$$

Note that the a_i 's are chosen so that weighting requirements are met. Also note that only one uniform random number and one other random number actually need to be generated for a particular value. The $F_i(Z_i)$'s may be any of the distributions which are listed below.

Since the test statistics which will be simulated are used to test hypotheses about one or more means, it will be convenient to parameterize the distributions in terms of their means.

Beta Distribution

The beta distribution is given by the density function

$$f(x) = \frac{\Gamma(A+B)}{\Gamma(A)\Gamma(B)} \left(\frac{x-C}{D-C} \right)^{A-1} \left(1 - \frac{x-C}{D-C} \right)^{B-1}, \quad A, B > 0, C \leq x \leq D$$

where A and B are shape parameters, C is the minimum, and D is the maximum. In statistical theory, C and D are usually zero and one, respectively, but the more general formulation used here is more convenient for simulation work. In this program module, a beta random variable is specified as $A(M, A, B, C)$, where M is the mean which is


$$E(X) = M = (D - C) \left[\frac{A}{A + B} \right] + C$$

The parameter D is obtained from M , A , B , and C using the relationship


$$D = \frac{(M - C)(A + B)}{A} + C.$$

The beta density can take a number of shapes depending on the values of A and B :


1. When $A < 1$ and $B < 1$ the density is U-shaped.




2. When $0 < A < 1 \leq B$ the density is J-shaped.




3. When $A=1$ and $B>1$ the density is bounded and decreases monotonically to 0.



4. When $A=1$ and $B=1$ the density is the uniform density.



5. When $A>1$ and $B>1$ the density is unimodal.




Beta random variates are generated using Cheng's rejection algorithm as given on page 438 of Devroye (1986).

Binomial Distribution

The binomial distribution is given by the function

$$\Pr(X = x) = \binom{n}{x} \pi^x (1 - \pi)^{n-x}, \quad x = 0, 1, 2, \dots, n$$

In this program module, the binomial is specified as $\mathbf{B}(M, n)$, where M is the mean which is equal to $n\pi$ and n is the number of trials. The probability of a positive response, π , is not entered directly, but is obtained using the relationship $\pi = M / n$. For this reason, $0 < M < n$.



Binomial random variates are generated using the inverse CDF method. That is, a uniform random variate is generated, and then the CDF of the binomial distribution is scanned to determine which value of X is associated with that probability.


Cauchy Distribution

The Cauchy distribution is given by the density function

$$f(x) = \left[S\pi \left(1 + \left\{ \frac{x - M}{S} \right\}^2 \right) \right]^{-1}, \quad S > 0$$

Although the Cauchy distribution does not possess a mean and standard deviation, M and S are treated as such. Cauchy random numbers are generated using the algorithm given in Johnson, Kotz, and Balakrishnan (1994), page 327.

In this program module, the Cauchy is specified as $\mathbf{C}(M, S)$, where M is a location parameter (median), and S is a scale parameter.



Constant Distribution

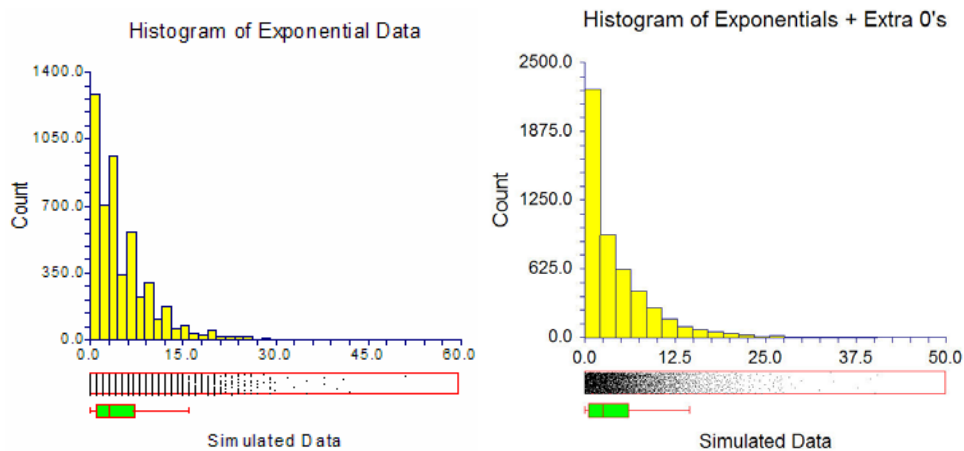
The *constant* distribution occurs when a random variable can only take a single value, X . The constant distribution is specified as $K(X)$, where X is the value.

Data with a Many Zero Values

Sometimes data follow a specific distribution in which there is a large proportion of zeros. This can happen when data are counts or monetary amounts. Suppose you want to generate exponentially distributed data with an extra number of zeros. You could use the following simulation model:

$K(0)[2]; E(5)[9]$

The exponential distribution alone was used to generate the histogram below on the left. The histogram below on the right was simulated by adding extra zeros to the exponential data.




Exponential Distribution

The exponential distribution is given by the density function

$$f(x) = \frac{1}{M} e^{-\frac{x}{M}}, \quad x > 0$$

In this program module, the exponential is specified as $E(M)$, where M is the mean.



Random variates from the exponential distribution are generated using the expression $-M \ln(U)$, where U is a uniform random variate.


F Distribution

Snedecor's F distribution is the distribution of the ratio of two independent chi-square variates. The degrees of freedom of the numerator chi-square variate is A , while that of the denominator chi-square is D . The F distribution is specified as $F(M, A)$, where M is the mean and A is the degrees of freedom of the numerator chi-square. The value of M is related to the denominator chi-square degrees of freedom using the relationship $M=D/(D-2)$.

F variates are generated by first generating a symmetric beta variate, $B(A/2, D/2)$, and transforming it into an F variate using the relationship

$$F_{A,D} = \frac{BD}{A - BD}$$

Below is a histogram for data generated from an F distribution with a mean of 1.1 and $A = 6$.




Gamma Distribution

The three parameter gamma distribution is given by the density function

$$f(x) = \frac{(x)^{A-1}}{B^A \Gamma(A)} e^{-\frac{x}{B}}, \quad x > 0, A > 0, B > 0$$

where A is a shape parameter and B is a scale parameter. In this program module, the gamma is specified as $G(M, A)$, where M is the mean, given by $M=AB$. The parameter B may be obtained using the relationship $B = M/A$.

Gamma variates are generated using the exponential distribution when $A = 1$, Best's XG algorithm given in Devroye (1986), page 410, when $A > 1$, and Vaduva's algorithm given in Devroye (1986), page 415, when $A < 1$.



Multinomial Distribution

The *multinomial* distribution occurs when a random variable has only a few discrete values such as 1, 2, 3, 4, and 5. The multinomial distribution is specified as $M(P_1, P_2, \dots, P_k)$, where P_i is the probability of that the integer i occurs. Note that the values start at one, not zero.

For example, suppose you want to simulate a distribution which has 50% 3's and 1's, 2's, 4's, and 5's all with equal percentages. You would enter $M(1\ 1\ 4\ 1\ 1)$.

As a second example, suppose you wanted to have an equal percentage of 1's, 3's, and 5's, and none of the other percentages. You would enter $M(1\ 0\ 1\ 0\ 0\ 1)$.


Likert-Scale Data

Likert-scale data are common in surveys and questionnaires. To generate data from a five-point Likert-scale distribution, you could use the following simulation model:

$M(6\ 1\ 2\ 1\ 5)$

Note that the weights are relative—they do not have to sum to one. The program will make the appropriate weighting adjustments so that they do sum to one.

The above expression generated the following histogram.




Normal Distribution

The normal distribution is given by the density function

$$f(x) = \phi\left(\frac{x - \mu}{\sigma}\right), \quad -\infty \leq x \leq \infty$$

where $\phi(z)$ is the usual standard normal density. The normal distribution is specified as $N(M, S)$, where M is the mean and S is the standard deviation.

The normal distribution is generated using the Marsaglia and Bray algorithm as given in Devroye (1986), page 390.




Poisson Distribution

The Poisson distribution is given by the function

$$\Pr(X = x) = \frac{e^{-M} M^x}{x!}, \quad x = 0, 1, 2, \dots, M > 0$$

In this program module, the Poisson is specified as $P(M)$, where M is the mean.

Poisson random variates are generated using the inverse CDF method. That is, a uniform random variate is generated and then the CDF of the Poisson distribution is scanned to determine which value of X is associated with that probability.




Student's T Distribution

Student's T distribution is the distribution of the ratio of a unit normal variate and the square root of an independent chi-square variate. The degrees of freedom of the chi-square variate are the degrees of freedom of the T distribution. The T is specified as $T(M, A)$, where M is the mean and A is the degrees of freedom. The central T distribution generated in this program has a mean of zero, so, to obtain a mean of M , M is added to every data value.

T variates are generated by first generating a symmetric beta variate, $B(A/2, A/2)$, with mean equal to 0.5. This beta variate is then transformed into a T variate using the relationship

$$T = \sqrt{A} \frac{X - 0.5}{\sqrt{X(1 - X)}}$$

Here is a histogram for data generated from a T distribution with mean 0 and 5 degrees of freedom.



Tukey's Lambda Distribution

Hoaglin (1985) presents a discussion of a distribution developed by John Tukey for allowing the detailed specification of skewness and kurtosis in a simulation study. This distribution is extended in the work of Karian and Dudewicz (2000). Tukey's idea was to reshape the normal distribution using functions that change the skewness and/or kurtosis. This is accomplished by multiplying a normal random variable by a skewness function and/or a kurtosis function. The general form of the transformation is

$$X = A + B\{G_s(z)H_h(z)z\}$$

where z has the standard normal density. The skewness function Tukey proposed is

$$G_s(z) = \frac{e^{gz} - 1}{gz}$$

The range of g is typically -1 to 1. The value of $G_0(z) \equiv 1$. The kurtosis function Tukey proposed is

$$H_h(z) = e^{hz^2/2}$$

The range of h is also -1 to 1.

Hence, if both g and h are set to zero, the variable X follows the normal distribution with mean A and standard deviation B . As g is increased toward 1, the distribution is increasingly skewed to the right. As g is decreased towards -1, the distribution is increasingly skewed to the left. As h is increased toward 1, the data are stretched out so that more extreme values are probable. As h is decreased toward -1, the data are concentrated around the center—resulting in a beta-type distribution.


The mean of this distribution is given by

$$M = A + B \left(\frac{e^{g^2/2(1-h)} - 1}{g\sqrt{1-h}} \right), \quad 0 \leq h < 1$$

which may be easily solved for A .

Tukey's lambda is specified in the program as $L(M, B, g, h)$ where M is the mean, B is a scale factor (when $g=h=0$, B is the standard deviation), g is the amount of skewness, and h is the amount of kurtosis.

Random variates are generated from this distribution by generating a random normal variate and then applying the skewness and kurtosis modifications. Here are some examples as g is varied from 0 to 0.4 to 0.6. Notice how the amount of skewness is gradually increased. Similar results are achieved when h is varied from 0 to 0.5.




Uniform Distribution

The uniform distribution is given by the density function

$$f(x) = \frac{1}{B-A}, \quad A \leq x \leq B$$

In this program module, the uniform is specified as $U(M, A)$, where M is the mean which is equal to $(A+B)/2$ and A is the minimum of x . The parameter B is obtained using the relationship $B=2M-A$.



Uniform random numbers are generated using Makoto Matsumoto's Mersenne Twister uniform random number generator which has a cycle length greater than $10E+6000$ (that's a one followed by 6000 zeros).

Weibull Distribution

The Weibull distribution is indexed by a shape parameter, B , and a scale parameter, C . The Weibull density function is written as

$$f(x|B, C) = \frac{B}{C} \left(\frac{x}{C} \right)^{(B-1)} e^{-\left(\frac{x}{C} \right)^B}, \quad B > 0, C > 0, x > 0.$$

Shape Parameter – B

The shape parameter controls the overall shape of the density function. Typically, this value ranges between 0.5 and 8.0. One of the reasons for the popularity of the Weibull distribution is that it includes other useful distributions as special cases or close approximations. For example, if


- B = 1 The Weibull distribution is identical to the exponential distribution.
- B = 2 The Weibull distribution is identical to the Rayleigh distribution.
- B = 2.5 The Weibull distribution approximates the lognormal distribution.
- B = 3.6 The Weibull distribution approximates the normal distribution.

Scale Parameter – C

The scale parameter only changes the scale of the density function along the x axis. Some authors use $1/C$ instead of C as the scale parameter. Although this is arbitrary, we prefer dividing by the scale parameter since that is how one usually scales a set of numbers.

The Weibull is specified in the program as $W(M, B)$, where M is the mean which is given by

$$M = C \Gamma \left(1 + \frac{1}{B} \right).$$



Combining Distributions

This section discusses how distributions may be combined to form new distributions. Combining may be done in the form of algebraic manipulation, mixtures, or both.

Creating New Distributions using Expressions

The set of probability distributions discussed above provides a basic set of useful distributions. However, you may want to mimic reality more closely by combining these basic distributions. For example, paired data is often analyzed by forming the differences of the two original variables. If the original data are normally distributed, then the differences are also normally distributed. Suppose, however, that the original data are exponential. The difference of two exponentials is not a common distribution.

Expression Syntax

The basic syntax is

$$C1 D1 operator1 C2 D2 operator2 C3 D3 operator3 \dots$$

where C1, C2, C3, etc. are coefficients (numbers), D1, D2, D3, etc. are probability distributions, and *operator* is one of the four symbols: +, -, *, /. Parentheses are only permitted in the specification of distributions.

Examples of valid expressions include

$$N(4, 5) - N(4, 5)$$

$$2E(3) - 4E(4) + 2E(5)$$

$$N(4, 2)/E(4)-K(5)$$

Notes about the Coefficients: C1, C2, C3

The coefficients may be positive or negative decimal numbers such as 2.3, 5, or -3.2. If no coefficient is specified, the coefficient is assumed to be one.

Notes about the Distributions: D1, D2, D3

The distributions may be any of the distributions listed above such as normal, exponential, or beta. The expressions are evaluated by generating random values from each of the distributions specified and then combining them according to the operators.

Notes about the operators: +, -, *, /

All multiplications and divisions are performed first, followed by any additions and subtractions.

Note that if only addition and subtraction are used in the expression, the mean of the resulting distribution is found by applying the same operations to the individual distribution means. If the expression involves multiplication or division, the mean of the resulting distribution is usually difficult to calculate directly.

Creating New Distributions using Mixtures

Mixture distributions are formed by sampling a fixed percentage of the data from each of several distributions. For example, you may model outliers by obtaining 95% of your data from a normal distribution with a standard deviation of 5 and 5% of your data from a distribution with a standard deviation of 50.

Mixture Syntax

The basic syntax of a mixture is

$$D1[W1]; D2[W2]; \dots; Dk[Wk]$$

where the D 's represent distributions and the W 's represent weights. Note that the weights must be positive numbers. Also note that semi-colons are used to separate the components of the mixture.

Examples of valid mixture distributions include

$N(4, 5)[19]; N(4, 50)[1]$	95% of the distribution is $N(4, 5)$, and the other 5% is $N(4, 50)$.
$W(4, 3)[7]; K(0)[3]$	70% of the distribution is $W(4, 3)$, and the other 30% is made up of zeros.
$N(4, 2)-N(4,3)[2]; E(4)*E(2)[8]$	20% of the distribution is $N(4, 2)-N(4,3)$, and the other 80% is $E(4)*E(2)$.

Notes about the Distributions

The distributions $D1$, $D2$, $D3$, etc. may be any valid distributional expression.

Notes about the Weights

The weights $w1$, $w2$, $w3$, etc. need not sum to one (or to one hundred). The program uses these weights to calculate new, internal weights that do sum to one. For example, if you enter weights of 1, 2, and 1, the internal weights will be 0.25, 0.50, and 0.25.

When a weight is not specified, it is assumed to have the value of '1.' Thus

$$N(4, 5)[19]; N(4,50)[1]$$

is equivalent to

$$N(4, 5)[19]; N(4,50)$$

Special Functions

A set of special functions is available to modify the generator number after all other operations are completed. These special functions are applied in the order they are given next.

Square Root (Absolute Value)

This function is activated by placing a \wedge in the expression. When active, the square root of the absolute value of the number is used.

Logarithm (Absolute Value)

This function is activated by placing a \sim in the expression. When active, the logarithm (base e) of the absolute value of the number is used.

Exponential

This function is activated by placing an $\&$ in the expression. When active, the number is exponentiated to the base e. If the current number x is greater than 70, $\exp(70)$ is used rather than $\exp(x)$.

Absolute Value

This function is activated by placing a $|$ in the expression. When active, the absolute value of the number is used.

Integer

This function is activated by placing a $\#$ in the expression. When active, the number is rounded to the nearest integer.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Data tab. To find out more about using the other tabs such as Template, go to the Procedure Window chapter.

Data Tab

The Data tab contains the parameters used to specify a probability distribution.

Data Simulation

Probability Distribution to be Simulated

Enter the components of the probability distribution to be simulated. One or more components may be entered from among the continuous and discrete distributions listed below the data-entry box.

The W parameter gives the relative weight of that component. For example, if you entered

$P(5)[1];K(0)[2]$, about 33% of the random numbers would follow the $P(5)$ distribution, and 67% would be 0. When only one component is used, the value of W may be omitted. For example, to

generate data from the normal distribution with mean of five and standard deviation of one, you would enter $N(5, 1)$, not $N(5, 1)[1]$.

Each of the possible components were discussed earlier in the chapter.

Data Simulation – Number of Values for Histogram and Summary

Number of Simulated Values

This is the number of values generated from the probability distribution for display in the histogram. We recommend a value of about 5000.

Reports Tab

The following options control the format of the reports.

Select Report

Numerical Summary

This option controls the display of this report.

Select Plot

Histogram

This option controls the display of the histogram.

Report Options

Precision

This allows you to specify the precision of numbers in the report. A single-precision number will show seven-place accuracy, while a double-precision number will show thirteen-place accuracy. Note that the reports are formatted for single precision. If you select double precision, some numbers may run into others. Also note that all calculations are performed in double precision regardless of which option you select here. This is for reporting purposes only.

Report Options – Percentile Options

Percentile Type

This option specifies which of five different methods is used to calculate the percentiles.

RECOMMENDED: **Ave $X_p(n+1)$** since it gives the common value of the median.

In the explanations below, p refers to the fractional value of the percentile (for example, for the 75th percentile $p = .75$), Z_p refers to the value of the percentile, $X[i]$ refers to the i th data value after the values have been sorted, n refers to the total sample size, and g refers to the fractional part of a number (for example, if $np = 23.42$, then $g = .42$). The options are

- **Ave $X_p(n+1)$**

This is the most commonly used option. The 100pth percentile is computed as

$$Z_p = (1-g)X[k1] + gX[k2]$$

where k_1 equals the integer part of $p(n+1)$, $k_2=k_1+1$, g is the fractional part of $p(n+1)$, and $X[k]$ is the k th observation when the data are sorted from lowest to highest.

- **Ave $X_{p(n)}$**

The 100pth percentile is computed as

$$Z_p = (1-g)X[k_1] + gX[k_2]$$

where k_1 equals the integer part of np , $k_2=k_1+1$, g is the fractional part of np , and $X[k]$ is the k th observation when the data are sorted from lowest to highest.

- **Closest to np**

The 100pth percentile is computed as

$$Z_p = X[k_1]$$

where k_1 equals the integer that is closest to np and $X[k]$ is the k th observation when the data are sorted from lowest to highest.

- **EDF**

The 100pth percentile is computed as

$$Z_p = X[k_1]$$

where k_1 equals the integer part of np if np is exactly an integer or the integer part of $np+1$ if np is not exactly an integer. $X[k]$ is the k th observation when the data are sorted from lowest to highest. Note that EDF stands for empirical distribution function.

- **EDF w/Ave**

The 100pth percentile is computed as

$$Z_p = (X[k_1] + X[k_2])/2$$

where k_1 and k_2 are defined as follows: If np is an integer, $k_1=k_2=np$. If np is not exactly an integer, k_1 equals the integer part of np and $k_2 = k_1+1$. $X[k]$ is the k th observation when the data are sorted from lowest to highest. Note that EDF stands for empirical distribution function.

Smallest Percentile

This option lets you assign a different value to the smallest percentile value shown on the percentile report. The default value is 1.0. You can select any value between 0 and 100, including decimal numbers.

Largest Percentile

This option lets you assign a different value to the largest percentile value shown on the percentile report. The default value is 1.0. You can select any value between 0 and 100, including decimal numbers.

Report Options – Report Decimal Places

Means – Values

Specify the number of decimal places used when displaying this item.

GENERAL: Display the entire number without special formatting.

Histogram Tab

This panel sets the options used to define the appearance of the histogram.

Vertical and Horizontal Axis

Label

This is the text of the label. The characters $\{Y\}$ and $\{X\}$ are replaced by appropriate names. Press the button on the right of the field to specify the font of the text.

Minimum

This option specifies the minimum value displayed on the corresponding axis. If left blank, it is calculated from the data.

Maximum

This option specifies the maximum value displayed on the corresponding axis. If left blank, it is calculated from the data.

Tick Label Settings...

Pressing these buttons brings up a window that sets the font, rotation, and number of decimal places displayed in the reference numbers along the vertical and horizontal axes.

Major Ticks and Minor Ticks

These options set the number of major and minor tickmarks displayed on the axis.

Show Grid Lines

This check box indicates whether the grid lines that originate from this axis should be displayed.

Histogram Settings

Plot Style File

Designate a histogram style file. This file sets all options that are not set directly on this panel. Unless you choose otherwise, the default style file (Default) is used. These files are created in the Histograms procedure.

Number of Bars

Specify the number of bars to be displayed. Select 'Automatic' to direct the program to select an appropriate number based on the number of values.

Titles

Plot Title

This is the text of the title. The characters $\{Y\}$, $\{X\}$, and $\{G\}$ are replaced by appropriate names. Press the button on the right of the field to specify the font of the text.

Storage Tab

The Data tab contains the parameters used to specify a probability distribution.

Storage of Simulated Values to Spreadsheet

Store Values in Column

This is the column of the spreadsheet in which the simulated values will be stored. Any data already in this column will be replaced.

Spreadsheet

Press this button to open the spreadsheet window for storage.

Numbers of Values Stored (Maximum of 16000)

This is the number of generated values that are stored in the current spreadsheet.

Example 1 – Generating Normal Data

In this example, 5000 values will be generated from the standard normal (mean zero, variance one) distribution. These values will be displayed in a histogram and summarized numerically.

Setup

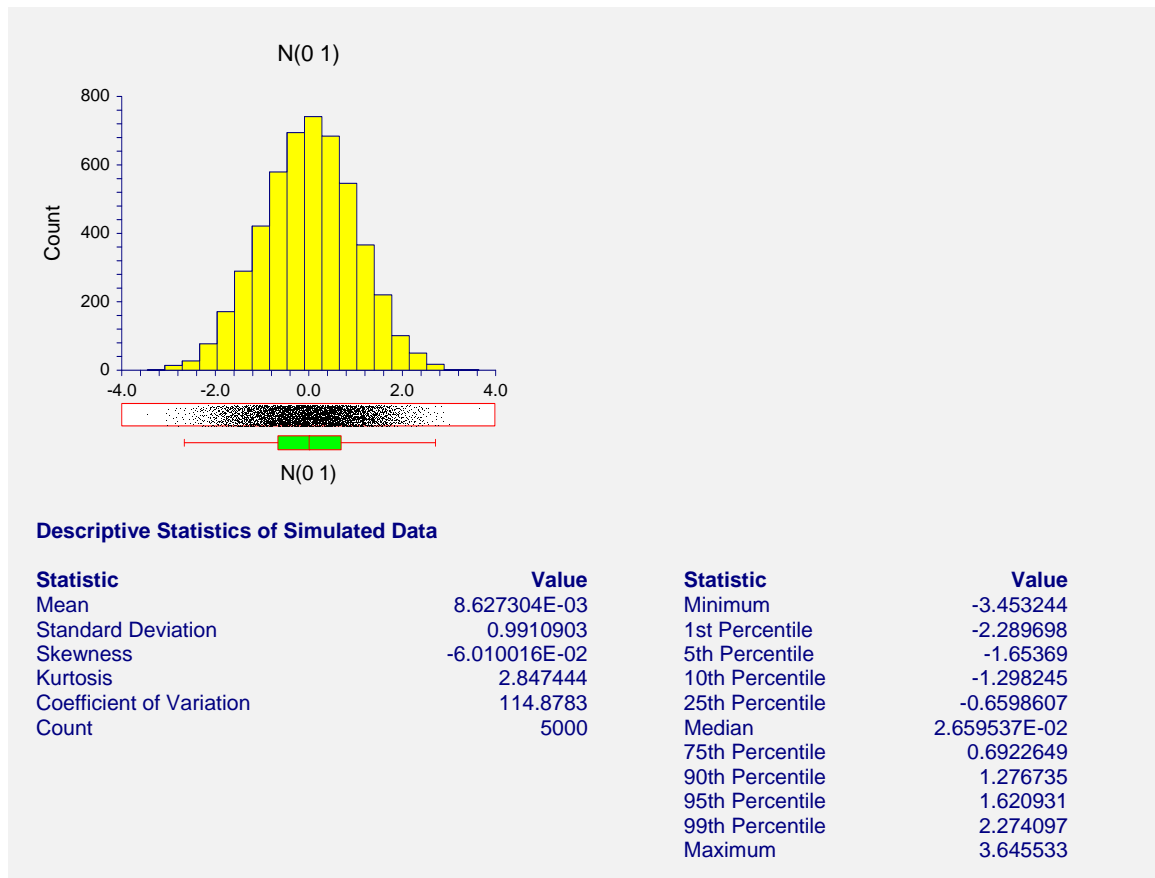
This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Data Simulator** procedure window by clicking on **Helps and Aids**, then **Data Simulator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example1** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Probability Distribution to be Simulated	N(0, 1)
Numbers in of Simulated Values.....	5000
Numbers of Values Stored.....	0

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



This report shows the histogram and a numerical summary of the 5000 simulated normal values. It is interesting to check how well the simulation did. Theoretically, the mean should be zero, the standard deviation one, the skewness zero, and the kurtosis three. Of course, your results will vary from these because these are based on generated random numbers.

Example 2 – Generating Data from a Contaminated Normal

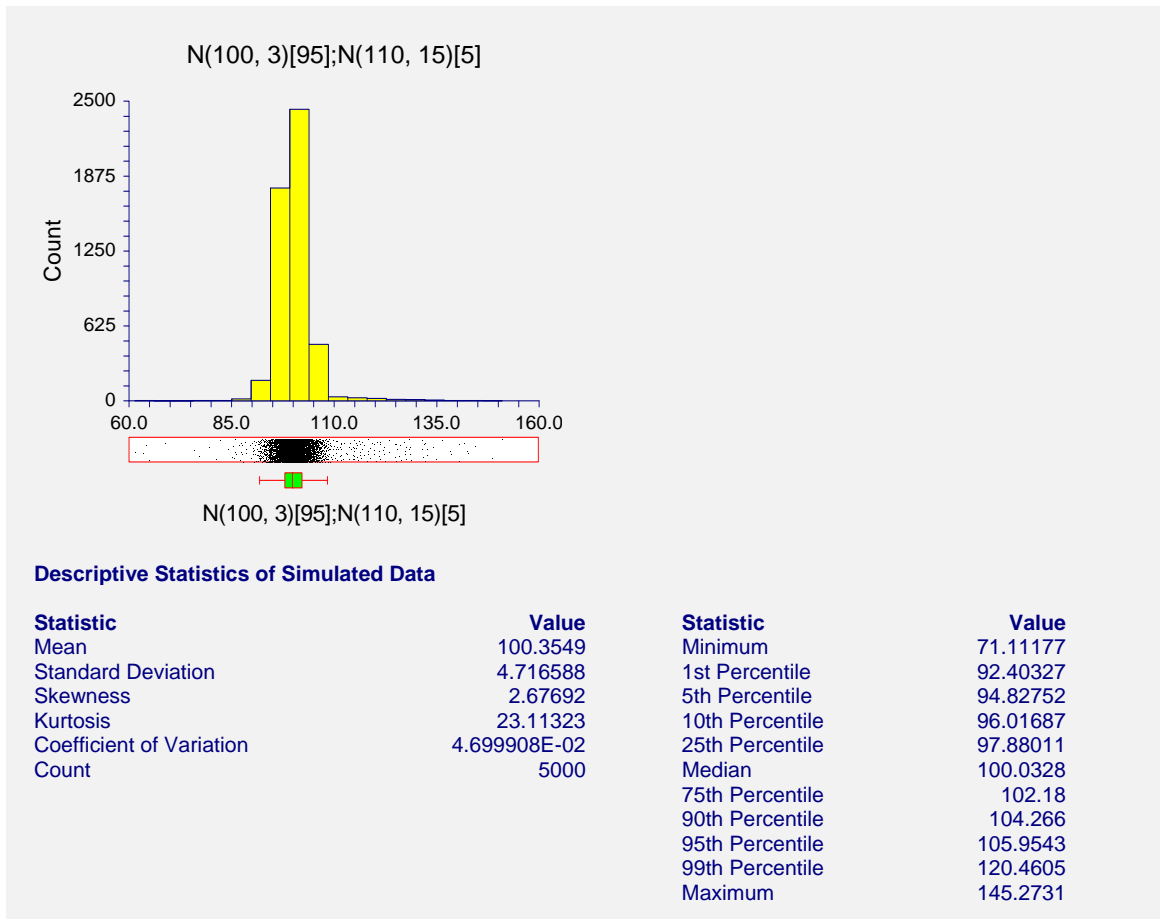
In this example, we will generate data from a contaminated normal. This will be accomplished by generating 95% of the data from a $N(100, 3)$ distribution and 5% from a $N(110, 15)$ distribution.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Data Simulator** procedure window by clicking on **Helps and Aids**, then **Data Simulator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example2** from the Template tab on the procedure window.

Option**Value****Data Tab**Probability Distribution to be Simulated **N(100 3)[95]; N(110 15)[5]**Numbers in of Simulated Values..... **5000**Numbers of Values Stored..... **0****Output**

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots

This report shows the data from the contaminated normal. The mean is close to 100, but the standard deviation, skewness, and kurtosis have non-normal values. Note that there are now some very large outliers.

Example 3 – Likert-Scale Data

In this example, we will generate data following a discrete distribution on a Likert scale. The distribution of the Likert scale will be 30% 1's, 10% 2's, 20% 3's, 10% 4's, and 30% 5's.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Data Simulator** procedure window by clicking on **Helps and Aids**, then **Data Simulator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example3** from the Template tab on the procedure window.

Option

Value

Data Tab

Probability Distribution to be Simulated **M(30 10 20 10 30)**


Numbers in of Simulated Values..... **5000**

Numbers of Values Stored..... **0**

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



Descriptive Statistics of Simulated Data

Statistic	Value	Statistic	Value
Mean	2.9858	Minimum	1
Standard Deviation	1.622506	1st Percentile	1
Skewness	1.169226E-02	5th Percentile	1
Kurtosis	1.436589	10th Percentile	1
Coefficient of Variation	0.5434074	25th Percentile	1
Count	5000	Median	3
		75th Percentile	5
		90th Percentile	5
		95th Percentile	5
		99th Percentile	5
		Maximum	5

This report shows the data from a Likert scale.

Example 4 – Bimodal Data

In this example, we will generate data that have a bimodal distribution. We will accomplish this by combining data from two normal distributions, one with a mean of 10 and the other with a mean of 30. The standard deviation will be set at 4.

Setup


This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Data Simulator** procedure window by clicking on **Helps and Aids**, then **Data Simulator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example4** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Probability Distribution to be Simulated	N(10 4);N(30 4)
Numbers in of Simulated Values.....	5000
Numbers of Values Stored.....	0

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



This report shows the results for the simulated bimodal data.

Example 5 – Gamma Data with Extra Zeros

In this example, we will generate data that have a gamma distribution, except that we will force there to be about 30% zeros. The gamma distribution will have a shape parameter of 5 and a mean of 10.

Setup

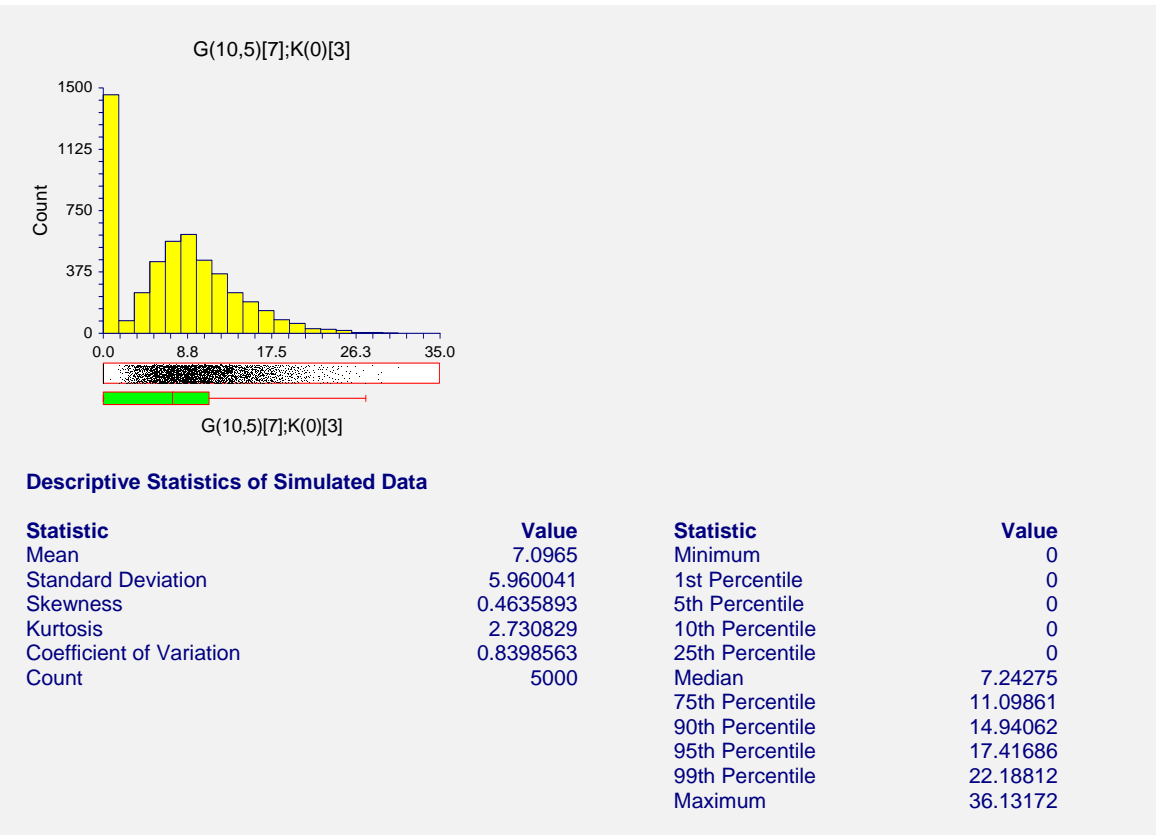
This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Data Simulator** procedure window by clicking on **Helps and Aids**, then **Data Simulator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example5** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Probability Distribution to be Simulated	G(10,5)[7];K(0)[3]
Numbers in of Simulated Values.....	5000
Numbers of Values Stored.....	0

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



This report shows the results for the simulated gamma data with extra zeros.

Example 6 – Mixture of Two Poisson Distributions

In this example, we will generate data that have a mixture of two Poisson distributions. 60% of the data will be from a Poisson distribution with a mean of 10 and 40% from a Poisson distribution with a mean of 20.

Setup


This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Data Simulator** procedure window by clicking on **Helps and Aids**, then **Data Simulator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example6** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Probability Distribution to be Simulated	P(10)[60];P(20)[40]
Numbers in of Simulated Values.....	5000
Numbers of Values Stored.....	0

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



This report shows the results for the simulated mixture-Poisson data.

Example 7 – Difference of Two Identically Distributed Exponentials

In this example, we will demonstrate that the difference of two identically distributed exponential random variables follows a symmetric distribution. This is particularly interesting because the exponential distribution is skewed. In fact, the difference between any two identically distributed random variables follows a symmetric distribution.

Setup


This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Data Simulator** procedure window by clicking on **Helps and Aids**, then **Data Simulator**. You may then follow along here by making the appropriate entries as listed below or load the completed template **Example7** from the Template tab on the procedure window.

<u>Option</u>	<u>Value</u>
Data Tab	
Probability Distribution to be Simulated	E(10)-E(10)
Numbers in of Simulated Values.....	5000
Numbers of Values Stored.....	0

Output

Click the Run button to perform the calculations and generate the following output.

Numeric Results and Plots



This report demonstrates that the distribution of the difference is symmetric.

Chapter 925

The Spreadsheet

Introduction

This chapter discusses the operation of the *Spreadsheet*, used to enter data for a few select **PASS** procedures and to store simulated data. The Spreadsheet is the tool that lets you enter, view, and modify data.

The operation of the **PASS** spreadsheet is similar to the operation of other spreadsheets with which you are familiar. In fact, it has most of the operational features of Microsoft Excel. Since the operation of these spreadsheets is so common, we will not spend a lot of space teaching them to you.

Spreadsheet Menus

You should be familiar with the operation of pull down menus. We will discuss the various options that are on these menus.

File Menu

The File Menu controls the opening and closing of spreadsheets. Note that some of the basic File Menu operations are also provided on the Toolbar.

We will now discuss each of these options.

- **New**

This option closes the current spreadsheet (if any) and creates a new spreadsheet. **PASS** spreadsheets have *.s0* extensions. The save operation will also create required *.s1*, *.s2*, etc. files depending on the number of sheets on the spreadsheet.

- **Open**

The Open option lets you open existing **PASS** spreadsheets. It will cause the Open Dialog box to appear from which you can select a file.

When selecting a spreadsheet, note that you select the file with the *.s0* extension. Attempting to open **PASS** files with other extensions (such as *.s1*, *.s2*, etc.) will produce unpredictable results.

Note that the spreadsheet is copied into memory. Once you open a spreadsheet, you actually have two copies of it—one in memory and one on your disk. No automatic relationship is maintained between the loaded spreadsheet and the disk spreadsheet. Changes made to the copy in memory will not automatically change your disk spreadsheet files unless you save them!

- **Printer Setup**

This option brings up a window that lets you set parameters of your printer(s).

- **Page Setup**

This option lets you specify the format of your datasheet printout. You can specify headers, footers, margins, page order (across or down), and scale (size of the print). It is often used to enable the printing of the row and column labels.

Headers and footers can contain text and special formatting codes. The following table lists the special formatting codes. Header and footer codes can be entered in upper or lower case.

Format Code	Description
&L	Left-aligns the characters that follow.
&C	Centers the characters that follow.
&R	Right-aligns the characters that follow.
&D	Prints the current date.
&T	Prints the current time.
&F	Prints the worksheet name (this is an internal name that may not be useful).
&P	Prints the page number.
&P+ <i>number</i>	Prints the page number plus number.
&P- <i>number</i>	Prints the page number minus number.
&&	Prints an ampersand.
&N	Prints the total number of pages in the document.

The following font codes must appear before other codes and text or they are ignored. The alignment codes (e.g., &L, &C, and &R) restart each section; new font codes can be specified after an alignment code.

Format Code	Description
&B	Use bold font.
&I	Use an italic font.
&U	Underline the header.
&S	Strikeout the header.
&"fontname"	Use the specified font.
&nn	Use the specified font size - must be a two digit number.

- **Print**

This option prints the entire datasheet or a portion that you designate. Note that you must print each datasheet of a spreadsheet separately. If the row and column labels do not show in your printout, select Page Setup from the File menu and check the appropriate selection boxes.

- **Save**
This option saves the current spreadsheet. Remember that a spreadsheet consists of several files (.s0, .s1, .s2, etc.). All of those files will be replaced.
- **Save As**
This option saves the current spreadsheet to a spreadsheet with a different name. For example, if you are working on a spreadsheet called “XYZ” and will be making changes to it, you might want to save a copy of it as “XYZ1” so that any mistakes you might have made will not destroy your original data.
- **Close Spreadsheet**
This option closes the spreadsheet.
- **Exit PASS**
This option quits the NCSS system.
- **Previously Open Files**
A list of previously open spreadsheets is presented. You may select any of them to revert to that spreadsheet directly.

Edit Menu

The Edit Menu controls the editing of spreadsheets. Note that some of the basic Edit Menu operations are provided on the Toolbar.

- **Undo**
Undo allows you to undo the last edit operation made. Note that only the most recent edit operation may be undone.

When you make wholesale changes to your spreadsheet (by cutting or pasting, for example), the Undo system requires a lot of memory to store additional information needed for an undo operation. If you see system resources getting low, make an additional change to a single cell. This will reset the undo system and free up system memory.
- **Cut**
The Cut option copies the currently selected (highlighted) data to the Windows clipboard and clears those cells. This data may be pasted at another location within NCSS or to another Windows program. We will discuss the process of selecting cells later in this chapter.
- **Copy**
The Copy option copies the currently selected (highlighted) data to the Windows clipboard. The selected data is untouched. The copied data may be pasted at another location within NCSS or to another Windows program. We will discuss the process of selecting cells later in this chapter.
- **Paste**
The Paste option copies data from the clipboard to the current datasheet at the currently selected location. The contents of the clipboard may have come from a previous Cut or Copy operation within *PASS* or from another Windows program.

For example, an easy way to analyze the means from an analysis of variance is to copy them from the Output screen and paste them in a datasheet. Furthermore, a quick way to import data from an *Excel* spreadsheet is to copy the data in *Excel* and paste it into a **PASS** datasheet.

The Paste option behaves differently depending on whether part of the datasheet is selected (highlighted). The data on the clipboard acts like a rectangular array of data (it has rows and columns). If there is no selection on the datasheet, the paste operation will place the data on your spreadsheet just as it appeared when it was copied to the clipboard. However, if the spreadsheet has a selected area, the paste operation will do two things. First, it will insert the data inside the selected area only (extra data will be omitted). Second, it will repeat either all or part of the clipboard so that the selected area is filled.

- **Paste Rotated**

This option works like the Paste option (above), except that the data are rotated ninety degrees so that the rows become the columns and the columns become the rows. The following example shows the result of using this option:

Data that was Copied

```
1  2
3  4
5  6
7  8
```

Result of Paste Rotated Operation

```
1  3      5      7
2  4      6      8
```

- **Clear**

The Clear option erases the data that is selected. Note that unlike the Cut option, the Clear option does not put the data on the clipboard.

- **Insert**

The Insert option inserts rows or columns into your datasheet at the current position of the cursor. When you select Insert, a dialog box appears that allows you to indicate whether you want to insert rows or columns (variables). The number of rows (or columns) inserted is determined by the number of rows (or columns) selected.

Hence, the steps to insert columns are as follows:

1. Select the number of columns you want to insert, beginning your selection at the column where you want them added.
2. Select Insert from the Edit Menu.
3. Select Columns from the dialog box.

A datasheet on a spreadsheet contains exactly 256 variables. When you insert new variables, the current variables are shifted to the right. The variables at the right of the datasheet are "pushed off" the datasheet and lost. For this reason, **PASS** first checks to make sure that there are enough empty variables at the edge of the datasheet to accommodate the inserted variables. If you find that you don't have room to insert variables that you need, simply add a new datasheet, cut the last several variables at the right of the datasheet, and paste them to the next datasheet. This will make room for inserting variables.

- **Delete**

The Delete option removes the currently selected rows or columns from your datasheet. When you select Delete, a dialog box appears that lets you indicate whether to delete rows (or columns). The number of rows (or columns) deleted is determined by the number of rows (columns) selected.

- **Fill**

This option brings up the Fill window which fills the current variable (or the currently selected block of cells) with the value specified. The value may be incremented so that special patterns such as 1,2,3,4 may be easily generated.

- **Find**

This option searches through your data for a designated value. Once you have started a find operation, use the Find Next button continues your search. You can search for a single digit or for the complete number.

- **Replace**

The Replace option allows you to quickly replace data throughout your datasheet. You can replace only those cells that match a certain pattern, or you can replace individual letters and digits.

- **Column Name**

This option allows you to change the column names.

- **Column Format**

This option allows you to change the format of spreadsheet columns.

- **Font**

This option allows you to change the font of cells in the spreadsheet.

Window Menu

This menu lets you transfer to one of the other *PASS* windows such as the PASS Home window or one of the currently open procedure windows.

Help Menu

From this menu you can launch the *PASS* Help System to an appropriate topic.

Spreadsheet Toolbar

The toolbar is provided for single-click access to the most commonly used menu options. You will find that each of the options on the toolbar can also be found in the menus. The toolbar has a feature called a "tool tip." This means that when you hold the mouse pointer over a certain square for at least a second, a small help box will appear that explains what this particular toolbar button is for. Most of the buttons on the toolbar follow Windows standards, so you will recognize them right away.

Cell Reference

Two boxes at the bottom of the screen give the *Cell Reference*. The first box gives the variable (column) number of the current location of the cell cursor. The second box gives the row number of the current location of the cell cursor. The cell cursor is the active cell. You can recognize the current cell because it will have an extra dark border.

Cell Edit

The Cell Edit box provides an alternate place to edit data. As you move around the spreadsheet, the contents of the active cell are copied to this Cell Edit box. Occasionally, your data will be longer than can easily be displayed in the cell. Although you could reset the column width, you usually find it easier to edit the data in the Cell Edit box.

Any changes you make will not be entered into the datasheet until you hit the ENTER key or position the mouse on another cell. After making changes to data in the Cell Edit box, you can press the ESC key to withdraw the changes.

Datasheet

This section of the screen shows the data. We will now describe how to use the spreadsheet to modify the data contained in a datasheet. You should know how to select cells, ranges, rows, and columns. Your work will go much faster if you learn how to quickly enter, modify, and delete data. These will all be described in this section.

Navigating the Datasheet

This section describes how to move around the datasheet using the keyboard and the mouse. In addition to moving around the datasheet, we will also describe how to make selections, copy data, and move data.

Active Cell

The datasheet cursor is always located on a single cell, even when a range of cells is selected. The cell on which the datasheet cursor is located is called the *Active Cell*. Any typing that is done will only affect the active cell. The contents of the active cell are displayed in the Cell Edit box. The address of the active cell is displayed in the Cell Reference boxes.

Keyboard Commands

The following commands are used mainly for data entry.

Key	Description
ENTER	Accepts the current entry and moves the active cell down one cell. When a range of cells is selected, accepts the current entry and moves down to the next selected cell. When the bottom of the selection is reached, the active cell moves to the top of the next selected column to the right.
SHIFT-ENTER	Acts like the ENTER key, except that cell-to-cell movement is upward and to the left instead of downward and to the right.

TAB	Accepts the current data entry and moves the cursor one cell to the right. When a range of cells is selected, the tab moves the cursor to the right to the next cell in the selection.
SHIFT-TAB	Acts just like the TAB key, except that cell-to-cell movement is to the left instead of to the right.
F2	Enters edit mode. Pressing F2 a second time brings up a cell-text data entry box.
DEL	Clears the current entry or selection.
ESC	Cancels the current data entry.

The following commands are used mainly for moving about the datasheet.

Key	Description
UP ARROW	Moves the active cell up one row.
DOWN ARROW	Moves the active cell down one row.
LEFT ARROW	Moves the active cell left one column.
RIGHT ARROW	Moves the active cell right one column.
CTRL UP / DOWN / LEFT / RIGHT ARROW	Moves to the next range of cells containing data. If there is no additional data in any of the cells in that direction, the active cell is moved to the edge of the datasheet.
PAGE UP	Moves up one screen.
PAGE DOWN	Moves down one screen.
CTRL PAGE UP	Moves left one screen.
CTRL PAGE DOWN	Moves right one screen.
HOME	Moves to the first column in the current row.
END	Moves to the last column in the current row that contains data.
CTRL HOME	Moves to the upper-left corner of the datasheet (cell 1,1).
CTRL END	Moves to the last row and column that contains data.
SCROLL LOCK	Modifies the action of the above movement keys. This key causes the datasheet window to scroll without changing the current selection. It works with all movement keys except HOME, END, CTRL HOME, and CTRL END.
SHIFT + any movement key	Extends the current selection in the direction indicated.

Mouse Actions

The mouse is used mainly for positioning the active cell and making selections. You can also use the mouse to move a block of cells around the datasheet.

Action	Description
Left Click	Accepts the current entry and moves the active cell to the position of the mouse.
Right Click	Does nothing.
Left Click in a Row or Column Heading	Selects the entire row or column.
Left Double Click	In-cell editing is invoked.
Right Double Click	Does nothing.
Left Click and Drag	Selects a range of cells. If other ranges were selected, they are unselected.
CTRL + Left Click and Drag	Selects a range of cells. If other ranges were selected, they remain selected. Note that edit commands, such as cut and copy, will only work on a single range.
SHIFT + Left Click and Drag	Extends the current selection in the direction indicated.
Dragging a Selection's Copy Handle	Copies the selection to a new location. The copy handle is the small plus sign at the lower-right corner of a selection.
Dragging a Selection's Border	Moves the selection to a new location. Note that when you move data around the datasheet in this fashion, no attempt is made to update the variable names. You will have to do that manually. For this reason, it is best to do this kind of wholesale editing before you attach names and transformations to the variables. Note also that Cell Transformation formulas are updated.

Selecting Cells

Many operations require one or more cells to be selected. There are three types of selections: a single cell, a single rectangular range of cells, and multiple ranges of non-adjacent cells. Cells may be selected either with the mouse or with the keyboard.

Selecting Cells with the Mouse

To select a range of cells with the mouse, click and hold the left mouse button down on the upper-left cell of the range you want to select. Drag the mouse cursor to the lower-right cell, while continuing to hold the left mouse button down. When the desired cells are selected, release the mouse button.

To select multiple ranges with the mouse, press the CTRL key while making each additional selection. Note that multiple selections are only useful for controlling cell cursor movement during data entry. They are not used by any of the edit functions.

To select an entire row or column, click on the row or column heading.

Once a range is selected, you can move the active cell within the selection using the Enter, SHIFT+ENTER, TAB, and SHIFT+TAB keys without destroying the selection.

Selecting Cells with the Keyboard

To select a range of cells with the keyboard, position the active cell at the upper-left corner of your desired selection. While holding down the SHIFT key, use the cursor movement keys (such as the arrow keys) to move to the lower-right corner of the selection.

Editing Datasheets Interactively

Data can be entered into a datasheet in many different ways. In this section, we will explain how you can quickly move and copy ranges of cells by clicking and dragging the copy handle of a selection.

Copying Data Interactively

You can copy a range of cells quickly by using only a few clicks of your mouse. The steps for doing this are enumerated next followed by figures that illustrate the action.

1. Select a range of one or more cells.
2. Select the *copy handle* (the small crosshair that appears at the lower-right corner of a selection) with your mouse cursor by positioning the mouse cursor over the copy handle and pressing the left mouse button.
3. Drag the copy handle through the range of cells that are to receive the copied data.
4. Release the left mouse button.

	C1	C2	C3
1	1	3	
2	2	4	
3			
4			
5			
6			
7			
8			
9			
10			

The copy handle is at the lower right corner of the selection.

925-10 The Spreadsheet

	C1	C2	C3
1	1	3	
2	2	4	
3	1	3	
4	2	4	
5	1	3	
6	2	4	
7	1	3	
8	2	4	
9	1	3	
10	2	4	
11			

The cursor changes to a crosshair as the copy handle is dragged down.

Note that the copied selection is repeated so that the whole copied area has data.

Moving Data Interactively

You can move a range of cells quickly by using only a few clicks of your mouse. The steps for doing this are enumerated next followed by figures that illustrate the action.

1. Select a range of one or more cells.
2. Position the mouse cursor on the border of the selection. When positioned on the border, the pointer changes to an arrow.
3. Drag the selection to the new location.
4. Release the left mouse button.

	C1	C2	C3
1	1	3	
2	2	4	
3			
4			
5			
6			
7			

First, select the cells you wish to move.

Next, position the mouse pointer at the border of the selected area. The mouse pointer will change to an arrow. Once the pointer has changed, do not release the mouse button.

	C1	C2	C3
1	1	3	
2	2	4	
3			
4			
5			
6			
7			

Move the selected cells to their new location.

	C1	C2	C3	
1				
2				
3				
4				
5		1	3	
6		2	4	
7				

Release the left mouse button. The contents of the selected cells will move to the new location.

Caution: the variable names have not changed. You will have to manually adjust variable names and transformations when you move data using this method.

Changing Row Heights and Column Widths

You can interactively resize the height of a row or the width of a column using the mouse. Position the pointer on the right edge of a column heading or the bottom edge of a row heading. The pointer will change shape. Simply drag the pointer to resize the row or column.

If multiple rows are selected when you resize a row, all selected rows are resized as you drag a row border. Multiple columns can be resized in like manner.

You can also set the size of a selected group of columns or rows to equal the size of another row or column. First, select the group of columns or rows you want to resize, including the column or row whose size you want to match. Next, click the right border of the column header or the bottom border of the row whose size you want to match. The columns or rows will all be resized.

Datasheet Tabs

The Datasheet Tabs at the bottom of the spreadsheet let you move quickly from one datasheet to another. These tabs are especially useful for allowing you to move to the Column Names sheet to make changes to variable names and formats.

Chapter 930

Macros

Introduction

This software has an interactive (point and click) user interface which makes it easy to learn and use. At times, however, it is necessary to repeat the same steps over and over. When this occurs, a batch system becomes more desirable. This chapter documents a batch language that lets you create a macro (script or program) and then run that macro. With the click of a single button, you can have the program run a series of procedures.

We begin with a discussion of how to create, modify, and run a macro. Next, we list all of the macro commands and their function.

Macro Command Center

This section describes how to create a macro, edit it, and run it. This is all accomplished from the Macro Command Center window. You can load this window by selecting *Macros* from the *Tools* menu or right-clicking on the *Macro Button* (green triangle) of the toolbar.

We will now describe each of the objects on the Macro Command Center window.

Active Macro

This box displays the name of the currently selected macro file. A preview of this macro is displayed in the Preview window. You can change this name by typing over it or by selecting a different macro from the *Existing Macros* window.

Existing Macros

This box displays a list of all existing macros. Click on a macro in this list to make it the active macro—the macro that is used when one of the control buttons to the right is pressed. Double-clicking a macro causes that macro to open in the window's Notepad program for editing.

The macros are stored in the *MACROS* subdirectory of the *PASS 2008* folder. They always have the file extension “.ncm”.

Preview of Active Macro

This box displays the beginning of the macro that is in the Active Macro box.

Record Macro

Pressing this button causes the commands you issue to be written to the macro file named in the *Active Macro* box. The Macro Command Center window will be replaced by a small *Macro* window that lets you stop recording the macro. When the recording starts, if there is a previous macro file with the same name as that in the Active Macro Name box, the previous macro file is erased.

The *PASS* macro recorder does not record every keystroke and click that you make. Instead, it records major operations. For example, suppose you want to include calculating the sample size for a t-test in your macro. You would load the t-test procedure, change some options, and run it. The macro recorder saves a copy of the t-test settings as a template file and writes a single command line to the macro file that references this template. All of your settings are included in the template file—there is no reference in the macro to the individual settings changes. This makes the macro much smaller and easier to modify.

Functions from the View, File, and Edit menus are not recorded during a macro recording. Those functions are performed using the SendKeys commands or other specific commands. As a general rule, the running of windows for which there is a template tab is recorded during a macro recording.

Edit Macro

Pressing this button causes the Active Macro to be loaded in the Windows NotePad program. This program lets you modify the macro and then save your changes.

Play Macro

Pressing this button causes the Active Macro to be run. Once the macro is finished, the Macro Command Center window will close and you can view the results of running your macro.

Delete Macro

Pressing this button causes the Active Macro to be deleted. The macro file is actually moved to the Recycle Bin from which it can be rescued if you decide it shouldn't have been deleted.

Close

Pressing this button closes the Macro Command Center window.

Select Button Macro

A macro button is displayed in the toolbar of several of the windows (such as the spreadsheet and procedure windows). This button causes the designated macro to be run. The macro that is associated with the macro button is controlled by this section of the Macro Command center. To change the macro that is associated with this button, simply select the desired macro from the *Existing Macros* list and then click the icon. This will associate the macro to the button. This association will remain even if you exit the program.

Syntax of a Macro Command Line

PASS macros are line based. That is, each macro command expression is written on a separate line. The basic structure of a line is that it begins with a *command* followed by one or more options or parameters of the command, called *arguments*. For example, the following macro opens a dataset and runs the Descriptive Statistics procedure on the first five variables in that dataset.

```
LoadProc PSProp2IP  
LoadTemplate PSProp2IP "Template1"  
RunProc PSProp2IP
```

In this example, LoadProc, LoadTemplate, and RunProc are commands, while PSProp2IP and "Template1" are arguments.

Comment Lines

It is often useful to add comment lines to a macro to make it easier to understand later. Comment lines begin with single quotes. When the macro processor encounters a single quote at the beginning of a line, the rest of the line is ignored. Single quotes occurring at a location other than the beginning of a line are treated as text.

Blank lines may also be added to a macro to improve readability. These are also ignored.

Macro Constants and Macro Variables

As stated above, macro commands lines consist of keywords followed by arguments. These arguments are either constants, such as "Template1" or 'PSProp2IP', or macro variables. These will be discussed next.

Macro Constants

Macro constants are fixed values. There are two types of macro constants: text and numeric.

Numeric Constants

Numeric constants are numbers. They may be whole or decimal numbers. They may be positive or negative. They may be enclosed in double quotes, although this is not necessary. When the macro processor expects a number but receives a text value, it sets the numeric value to zero.

Examples of numeric constants are

1, 3.14159, and 0.

Text Constants

Text constants are usually enclosed in double quotes. If a constant is a single word (made of letters and digits with no blanks or special characters), the double quotes are not necessary.

Examples of text constants are

Apple, "Apple Pie", and "D:/Program Files".

Macro Variables

Macro Variables are used to store temporary values for use in macro command lines. Some examples of assigning values to macro variables are

```
A# = 4
B# = 4 + 3
File$ = "C:/Program Files/PASS/Data/ABC.S0"
F$ = "4" & "5"
```

In these examples, A#, B#, File\$, and F\$ are macro variables. The assigned values for each of the variables are 4, 7, "C:\Program Files\PASS\Data\ABC.S0", and 45, respectively.

There are two types of macro variables: text and numeric.

Text Macro Variables

Text macro variables are used to hold text values. The rules for naming them are that the names can contain only letters and numbers (no spaces or special characters) and they must end with a '\$'. The case of the letters is ignored (so 'A\$' is used interchangeably with 'a\$').

930-4 Macros

Examples of text macro variable names are

A\$
Apple\$
FileName\$

Numeric Macro Variables

Numeric variables are used to hold numeric values. The rules for naming them are that the names can contain only letters and numbers (no spaces or special characters), and they must end with a '#'. The case of the letters is ignored (so 'A#' is used interchangeably with 'a#').

Examples of numeric macro variable names are

A#
Apple#
NRows#

Assigning Values to Macro Variables

One type of macro variable expression is that of assigning a value to a variable. The basic syntax for this type of expression is

{variable} = {value}

where the {variable} is text or numeric and {value} is a macro variable or (text or numeric) macro constant. If a text value contains spaces or special characters, it must be enclosed in double quotes.

A text value can be assigned to a numeric macro variable. In this case, the text value is converted to a number. If it cannot be converted to a number (e.g., it is a letter), the numeric macro variable is set to zero.

Following are some examples of valid assignment expressions.

A# = 4
X\$ = John
File\$ = "C:\Program Files\PASS\Data\ABC.S0"
X\$ = File\$
X\$ = A#
A# = F\$
F\$ = 4

Macro Variable Combination Expressions

Macro variables can be combined using simple mathematical expressions. The basic syntax for this type of expression is

{variable} = {value} {operator} {value}.

The available operators are + (add), - (subtract), * (multiply), / (divide), and & (concatenate).

If a text value is involved in a mathematical expression, it is converted to a numeric value before the mathematical expression is evaluated. If it cannot be converted to a number, the text value is set to the numeric value of zero.

Following are some examples of valid assignment expressions.

Expression	Result
A# = 4 + 3	A# = 7
B# = A# * 2	B# = 14
C\$ = "C:/Pgm/"	
D\$ = C\$ & "ABC.S0"	D\$ = "C:/Pgm/ABC.S0"
E# = C\$ * 4	E# = 0
F\$ = "4" + "5"	F\$ = "9"
F\$ = "4" & "5"	F\$ = "45"
A# = 1	
A# = A# + 1	A# = 2

Macro variable assignments are used while the macro is running, but are not saved when the macro has completed.

Displaying Macro Variables

The value of one or more macro variables may be displayed on the output using the PRINT or HEADING commands.

Print

This command outputs the requested values to the printout.

The syntax of this command is

PRINT {p1} {p2} {p3} ...

where

{p1} {p2} ... are assigned macro variables or constants.

Following are some examples of Print commands.

<u>Command</u>	<u>Printed Result</u>
PRINT "Hi World"	Hi World
I# = 1	
J\$ = "C:/PASS/Data/ABC"	
F\$ = J\$ & i#	
F\$ = F\$ & ".s0"	
PRINT "File=" F\$	File=C:/PASS/Data/ABC1.s0
PRINT "1" "2" "3" "4"	1 2 3 4

Heading

This command adds a line to the page heading that is shown at the top of each page.

The syntax of this command is

HEADING {h1}

where

{h1} is an assigned macro variable or constant.

Following are some examples of valid HEADING commands.

<u>Command</u>	<u>Heading</u>
HEADING "Hi World"	Hi World
F\$ = "Heart Study"	
HEADING F\$	Heart Study

Logic and Control Commands

The lines in a macro are processed in succession. These commands allow you to alter the order in which macro lines are processed, allow user-input, or end the program.

List of Logic and Control Commands:

Flag
GOTO
IF
INPUT
END
SendKeys

Flag Statement

A flag is a reference point in the program. The GOTO command sends macro line control to a specific flag. A flag is made up of letters and numbers (no spaces) followed by a colon.

Following are some examples of valid flags. More extensive examples are shown in the description of the IF statement (below)

Examples
Flag1:
A:
Loop1:

GOTO command

This command transfers macro processing to the next statement after a flag.

The syntax of this command is

GOTO {P1}

where

{P1} is a text variable or text constant.

Following are some examples of valid GOTO commands.

Examples
GOTO Flag1
or
F\$ = "Flag1"
GOTO F\$

IF command

This command transfers macro processing to the next statement after a flag if a condition is met. The syntax of this command is

IF {p1} {logic} {p2} GOTO {p3}

where

{p1} is a variable or constant.

{p2} is a variable or constant.

{p3} is a flag.

{logic} is a logic operator. Possible logic operators are =, <, >, <=, >=, and <>.

Following are some examples of valid IF commands.

Examples

IF x1# > 5 GOTO flag1

IF y\$ = "A" GOTO flag2

IF y\$ <> "A" GOTO flag3

INPUT

This command stops macro execution, display a message window, and waits for a value to be input. This value is then stored in the indicated macro variable.

The syntax of this command is

INPUT {variable} {prompt} {title} {default}

where

{variable} is the name of the variable (text or numeric) to receive the value that is input.

{prompt} is the text phrase that is shown on the input window.

{title} is the text phrase that is displayed at the top of the input window.

{default} is the default value for the input.

Following is an example of this command.

Example

INPUT A# "Enter the number of items" "Macro Input Window" 1

END

This command closes the *PASS* system.

The syntax of this command is

END

Example

END

SendKeys

This command sends one or more keystrokes to the program as if you had typed them in from the keyboard. This facility allows you to create macros to accomplish almost anything you can do interactively within the program.

To use this, run the program from the keyboard, noting exactly which keys are pressed. Then, type the appropriate commands into the sendkeys text. Note that spaces are treated as characters, so '{down} {tab}' is different from '{down}{tab}'.

The syntax of this command is

SendKeys {value}

where *{value}* is a text constant or variable.

Remarks

Each key is represented by one or more characters. To specify a single keyboard character, use the character itself. For example, to represent the letter A, use "A" for value. To represent more than one character, append each additional character to the one preceding it. To represent the letters A, B, and C, use "ABC" for string.

The plus sign (+), caret (^), percent sign (%), tilde (~), and parentheses () have special meanings to SendKeys. To specify one of these characters, enclose it within braces ({}). For example, to specify the plus sign, use {+}. Brackets ([]) have no special meaning to SendKeys, but you must enclose them in braces. To specify brace characters, use {{}} and {}.

To specify characters that are not displayed when you press a key, such as ENTER or TAB, and keys that represent actions rather than characters, use the following codes:

<u>Key</u>	<u>Code</u>
Backspace	{bs}
Break	{break}
Caps Lock	{capslock}
Delete	{delete}
Down Arrow	{down}
End	{end}
Enter	{enter}
Esc	{esc}
Home	{home}
Insert	{insert}
Left Arrow	{left}
Num Lock	{numlock}
Page Down	{pgdn}
Page Up	{pgup}
Right Arrow	{right}
Tab	{tab}
Up Arrow	{up}
F1	{F1}
F2	{F2}
F3	{F3}
F4	{F4}
F5	{F5}
F6	{F6}
F7	{F7}

<u>Key</u>	<u>Code</u>
F8	{F8}
F9	{F9}
F10	{F10}
F11	{F11}
F12	{F12}
F13	{F13}
F14	{F14}
F15	{F15}
F16	{F16}

To specify keys combined with any combination of the SHIFT, CTRL, and ALT keys, precede the key code with one or more of the following codes:

<u>Key</u>	<u>Code</u>
Shift	+
Ctrl	^
Alt	%

To specify that any combination of SHIFT, CTRL, and ALT should be held down while several other keys are pressed, enclose the code for those keys in parentheses. For example, to specify to hold down SHIFT while E and C are pressed, use "+(EC)". To specify to hold down SHIFT while E is pressed, followed by C without SHIFT, use "+EC".

To specify repeating keys, use the form {key number}. You must put a space between key and number. For example, {LEFT 4} means press the LEFT ARROW key 4 times; {h 8} means press H 8 times.

Spreadsheet Note: When a macro is run the usual beginning location on the screen is the new page icon (just below File) in the upper left of the screen. A single tab may be entered in the macro to go to the upper left cell position on the spreadsheet.

<u>Examples</u>	<u>Action</u>
SendKeys "ABC {enter}"	(Types ABC and then enter)
SendKeys "{enter right}"	(Enter and then down arrow)
SendKeys "%H{down 2}{enter}"	(Activates the Serial Numbers from the Help menus)
SendKeys "{right}"	(Moves to the right one cell)

Window Position Commands

These commands allow the user to position or hide windows while the macro is running.

List of Window Commands:

WindowLeft
WindowTop

WindowLeft

This command sets the position of the left edge of the spreadsheet, panel, and output windows. This allows you to effectively hide these windows while a macro is running.

930-10 Macros

The syntax of this command is

WINDOWLEFT {value}

where

{value} is the value of the left edge in thousandths of an inch.

Examples

WINDOWLEFT 0

WINDOWLEFT 10000

WINDOWLEFT -10000

Action

(positions the left edge to zero)

(positions the left edge ten inches to the right)

(positions the left edge ten inches to the left)

WindowTop

This command sets the position of the top of the spreadsheet, panel, and output windows. This allows you to effectively hide these windows while a macro is running.

The syntax of this command is

WINDOWTOP {value}

where

{value} is the value of the left edge in thousandths of an inch.

Examples

WINDOWTOP 0

WINDOWTOP -10000

Action

(positions the top to zero)

(positions the top ten inches up)

Spreadsheet Commands

The following commands open, close, and modify an **PASS** spreadsheet.

List of Dataset Commands:

DataOpen

DataNewS0

DataSaveS0

AddASheet

RemoveLastSheet

ResizeRowsCols

GetCell

SetCell

NumRows

GetMaxRows

VarName

VarFormat

NumVars

VarFromList

DataOpen

This command opens the database given by *{file name}*.

The syntax of this command is

DataOpen {filename}

where

{file name} a text variable or text phrase enclosed in double quotes that gives the name of the database to be opened.

Following are some examples of macro snippets that use this command.

```
DataOpen "C:/Program Files/PASS/Sample.s0"
or
F1$="C:/Program Files/PASS/"
F2$="Sample.s0"
F$=F1$ & F2$
DataOpen F$
```

DataNewS0

This command creates a new, untitled spreadsheet-type database. Note that these databases are limited to 16,384 rows of data.

The syntax of this command is

DataNewS0

Following is an example of this command.

```
DataNewS0
```

DataSaveS0

This command saves the current spreadsheet-type database to a file.

The syntax of this command is

DataSaveS0 {file name}

where

{file name} is a text variable or text phrase enclosed between double quotes. Note that the extension of the file name must be "s0". If any data already exists in this database, it will be replaced.

Following are some examples of macro snippets that use this command.

```
DataSaveS0 "C:/Program Files/PASS/Sample.s0"
or
F1$="C:/Program Files/PASS/"
F2$="Sample.s0"
F$=F1$ & F2$
DataSaveS0 F$
```

AddASheet

This command adds another 'sheet' to the current spreadsheet file. Note that each sheet contains 256 variables. You should only add a second sheet if you have used up the 256 variables on the first sheet.

The syntax of this command is

AddASheet

[Example](#)
AddASheet

RemoveLastSheet

This command removes the last 'sheet' of the current spreadsheet file. Note that a sheet cannot be removed if it contains data. Also, the first sheet cannot be removed.

The syntax of this command is

RemoveLastSheet

[Example](#)
RemoveLastSheet

ResizeRowsCols

This command resizes all rows and columns in the spreadsheet or database according to the user-specified resize type. All three types create columns (rows) no narrower (shorter) than the default width (height). If this command is used with a spreadsheet (.S0), the resulting row heights and column widths are saved when the spreadsheet is saved. When this operation is used with a database (.S0Z), the resulting row heights and column widths are not saved when the database is saved. The command must be used each time the database is opened.

The syntax of this command is

ResizeRowsCols {rtype}

where

{rtype} is an integer corresponding to the type of row/column resizing to be done (0 = resize using defaults, 1 = resize using data and titles, 2 = resize using data only).

[Example](#)
ResizeRowsCols 1

GetCell

This command obtains the value of a spreadsheet cell. The syntax of this command is

GetCell {variable} {row} {macro variable}

where

{variable} is the name or number of the variable (column) with the cell to be read.

{row} is the row number of the cell to be read.

{macro variable} is the text or numeric macro variable that holds the value of the spreadsheet cell.

Examples

```
GetCell "HeartRate" 27 H#
GetCell Name 16 Name16$
```

SetCell

This command sets a spreadsheet cell to a specified value. The syntax of this command is

SetCell {v1} {row1} {row2} {value}

where

{v1} is the name or number of the variable to receive the new value.
{row1} is the first row in a range of rows to receive the new value.
{row2} is the last row in a range of rows to receive the new value
{value} is the new value. This value may be text or numeric.

Examples

```
SetCell "HeartRate" 10 10 "100"
SetCell 1 10 20 100
```

NumRows

This command loads the number of rows used in a dataset column into a program variable. The syntax of this command is

NumRows {v1} {n}

where

{v1} is a variable name or number on the current database.
{n} is a numeric macro variable.

Examples

```
NumRows "HeartRate" n1#
NumRows 1 n#
```

GetMaxRows

This command loads the maximum number of rows used by any variable into a program variable.

The syntax of this command is

GetMaxRows {n}

where

{n} is a numeric macro variable.

Examples

```
GetMaxRows n1#
GetMaxRows n#
```

VarName

This command sets the name of the specified variable.

The syntax of this command is

VARNAME *{variable}* *{name}*

where

{variable} is the current name or number of the variable to be renamed.

{name} is the new name of the variable. This name must follow standard PASS variable name restrictions.

Examples

```
VARNAME C1 "HeartRate"  
VARNAME 1 "HeartRate"
```

VarFormat

This command sets the display format of the specified variable.

The syntax of this command is

VARFORMAT *{variable}* *{format}*

where

{variable} is the name or number of the variable to be formatted.

{format} is the format.

Examples

```
VARFORMAT C1 "0.00"  
VARFORMAT 2 "MM/DD/YYYY"
```

NumVars

This command causes the number of variables contained in a list of variables to be loaded into a macro variable.

The syntax of this command is

NUMVARS *{varlist}* *{x}*

where

{varlist} is an expression containing variable names.

{x} is macro variable.

Example

```
NUMVARS "C2:C10, C15" nvars#
```

VarFromList

This command causes number of the *i*th variable in a list to be loaded into a macro variable.

The syntax of this command is

VarFromList *{varlist}* *{i}* *{v}*

where

{varlist} is an expression containing variable names.

{i} is the item number to be selected from the variable list.

{v} is the macro variable that receives the number of the *i*th item in the list.

Example

VARFROMLIST "C1,C2,C3,C10,C20" 4 V1# (V1 becomes 10)

VARFROMLIST "C1,C2,C3,C10,C20" 5 V2# (V2 becomes 20)

Procedure Commands

The following commands open, modify, run and close procedures.

List of Procedure Commands:

LoadProc

RunProc

SaveTemplate

UnloadProc

Option

LoadProc

This command loads the designated procedure window. Once loaded, the options of the procedure may be modified and then the procedure can be executed.

The syntax of this command is

LoadProc {proc} {template}

where

{proc} is a variable or constant that gives the name or number of the procedure to be loaded. Each procedure's name and number is displayed near the bottom of the window under the Template tab.

{template} is an optional text variable or text constant that gives the name of a template file that is loaded with this procedure. If this value is omitted, the default (last) template for this procedure is loaded. Note that the text value does not include the extension or the folder information for the template file.

Following are some examples of valid LOADPROC commands.

Example

LOADPROC 24 "macro 1"

LOADPROC PSProp2IP "macro 1"

LOADPROC PSProp2IP

LOADPROC 24

RunProc

This command executes the indicated procedure. The syntax of this command is

RunProc {proc} {template}

930-16 Macros

where

{proc} is a variable or constant that gives the name or number of the procedure to be run. Each procedure's name and number is displayed near the bottom of the window under the Template tab.

{template} is a required variable or text constant that gives the name of the resulting template file. Note that the text value does not include the extension or the folder information for the template file.

SaveTemplate

This command saves the settings in the last procedure loaded to a template file. Once loaded, the options of the procedure may be modified and then the procedure can be executed.

The syntax of this command is

SaveTemplate {proc} {template} {id}

where

{proc} is a variable or constant that gives the name or number of the procedure whose template is to be saved. Each procedure's name and number is displayed near the bottom of the window under the Template tab.

{template} is a required variable or text constant that gives the name of the resulting template file. Note that the text value does not include the extension or the folder information for the template file.

{id} is an optional text variable or text constant that is stored with the file. This text is displayed with the file name under the Template tab.

Following are some examples of valid SaveTemplate commands.

Example

SaveTemplate PSProp2IP "Template1"

SaveTemplate PSProp2IP "Template1" "This template was created by a macro on January 1"

UnloadProc

This command closes the indicated procedure window. The syntax of this command is

UnloadProc {proc}

where

{proc} is a variable or constant that gives the name or number of the procedure. Each procedure's name and number is displayed near the bottom of the window under the Template tab.

Option

This command lets you set the values of the individual options of a procedure. If the option numbers are not displayed on the procedure windows, go to View or File, Options, View, and check the box next to 'Show Option Numbers'.

For example, you may want to change the value of Alpha or the type of test that is run.

The syntax of this command is

Option {proc} {number} {value1} {value2} {value3} ...

where

{proc} is a variable or constant that gives the name or number of the procedure. Each procedure's name and number is displayed near the bottom of the window under the Template tab.

{number} is the number of the option that is to be set. This number is displayed at the lower, left corner of the procedure window when the mouse is positioned over that option.

If the 'Opt' value is not displayed, activate it by doing the following: from the spreadsheet menus select File or View, Options, and View. Check the option labeled 'Show Option Numbers'.

{value1} is the new value of the option.

{value2}... are the new values of the remaining parameters of the option. Most options only have one value, so a second value is not necessary. However, a few options, such as text properties, bring up a window for option selection. These options have two or more parameters.

Following are some examples of valid OPTION commands.

Example

OPTION 24 2 4

OPTION "PSProp2IP" 3 0.8

Output Commands

The following commands manage the output (word processor) windows.

List of Output Commands:

SaveOutput

ClearOutput

PrintOutput

AddToLog

NewLog

SaveLog

OpenLog

SaveOutput

This command saves the current output to the designated file name.

The syntax of this command is

SaveOutput {filename}

{filename} a text constant or variable that gives the name of the file to receive the output. Note that the extension of the file name should be '.RTF'.

Example

SAVEOUTPUT "C:/Program Files/PASS/Sample.rtf"

ClearOutput

This command clears (erases) the current output.

The syntax of this command is

ClearOutput

Example
CLEAROUTPUT

PrintOutput

This command prints the current output.

The syntax of this command is

PrintOutput

Example
PRINTOUTPUT

AddToLog

This command copies the output in the output window to the log window. Note that nothing is saved by this command.

The syntax of this command is

AddToLog

Example
ADDTOLOG

NewLog

This command clears log window.

The syntax of this command is

NewLog

Example
NEWLOG

SaveLog

This command saves the current contents of the log output window to the designated file name.

The syntax of this command is

SaveLog {filename}

{filename} a text constant or variable that gives the name of the file to receive the log. Note that the extension of the file name should be '.RTF'.

Example
SAVELOG "C:/Program Files/PASS/Reports/Sample.rtf"

OpenLog

This command opens and displays the contents of the specified file.

The syntax of this command is

OpenLog {filename}

{filename} a text constant or variable that gives the name of the file to opened. Note that only RTF files can be opened.

Example

OPENLOG "C:/Program Files/PASS/Sample.rtf"

Alphabetical Macro Command List

AddASheet
 AddToLog
 ClearOutput
 DataNewS0
 DataOpen {filename}
 DataSaveS0 {file name}
 End
 GetCell {variable} {row} {macro variable}
 GetMaxRows {n}
 Goto {P1}
 Heading {h1}
 If {p1} {logic} {p2} goto {p3}
 Input {variable} {prompt} {title} {default}
 LoadProc {proc} {template}
 LoadTemplate {proc} {template}
 NewLog
 NumRows {v1} {n}
 NumVars {varlist} {x}
 OpenLog {filename}
 Option {number} {value1} {value2} {value3} ...
 Print {p1} {p2} {p3} ...
 PrintOutput
 RemoveLastSheet
 ResizeRowsCols
 RunProc {proc}
 SaveLog {filename}
 SaveOutput {filename}
 SaveTemplate {proc} {template} {id}
 SendKeys {value}
 SetCell {v1} {row1} {row2} {value}
 SortBy {v1} {o1} {v2} {o2} {v3} {o3}
 UnloadProc {proc}
 VarFormat {variable} {format}
 VarFromList {varlist} {i} {v}
 VarName {variable} {name}
 WindowLeft {value}
 WindowTop {value}

Examples

The following section provides examples of PASS macros. Our intention is that these examples will help you learn how to write macros to accomplish various repetitive tasks with PASS.

Example 1 – Automatically Run a Procedure

This macro opens the two-sample T-test procedure and runs the procedure with the default values.

```
*** Load the Two-Sample T-Test procedure
*** Note that the default values are used.
LoadProc PSMean2ID

*** Run the analysis
RunProc PSMean2ID
*** End of Macro 1
```

Example 2 – Run Several Procedures

This macro runs three procedures. The two-sample T-test procedure is run with the default values, except that Alpha is set to 0.10. The one-sample T-test procedure is run with default values, except that Alpha is set to 0.15. The linear regression procedure is run with the values set by the template 'My LinReg Template'.

```
*** Load the Two-Sample T-Test procedure
LoadProc PSMean2ID
*** Set Alpha (Option 8) to 0.10
Option PSMean2ID 8 0.10
*** Run the analysis
RunProc PSMean2ID

*** Load the One-Sample T-Test procedure
LoadProc PSMeanCorr
*** Set Alpha (Option 8) to 0.15
Option PSMeanCorr 8 0.15
*** Run the analysis
RunProc PSMeanCorr

*** Load the Linear Regression procedure
LoadProc PSLinReg
*** Load the 'My LinReg Template' template
LoadTemplate PSLinReg "My LinReg Template"
*** Run the analysis
RunProc PSLinReg
```

References

A

- A'Hern, R. P. A.** 2001. "Sample size tables for exact single-stage phase II designs." *Statistics in Medicine*, Volume 20, pages 859-866.
- Al-Sunduqchi, Mahdi S.** 1990. *Determining the Appropriate Sample Size for Inferences Based on the Wilcoxon Statistics*. Ph.D. dissertation under the direction of William C. Guenther, Dept. of Statistics, University of Wyoming, Laramie, Wyoming.
- Armitage, P., and Colton, T.** 1998. *Encyclopedia of Biostatistics*. John Wiley, New York.
- Armitage, P., McPherson, C.K., and Rowe, B.C.** 1969. "Repeated significance tests on accumulating data." *Journal of the Royal Statistical Society, Series A*, 132, pages 235-244.
- Atkinson, A.C.** 1985. *Plots, Transformations, and Regression*. Oxford University Press, Oxford (also in New York). This book goes into the details of regression diagnostics and plotting. It puts together much of the recent work in this area.
- Atkinson, A.C., and Donev, A.N.** 1992. *Optimum Experimental Designs*. Oxford University Press, Oxford. This book discusses D-Optimal designs.

B

- Bain, L.J. and Engelhardt, M.** 1991. *Statistical Analysis of Reliability and Life-Testing Models*. Marcel Dekker. New York. This book contains details for testing data that follow the exponential and Weibull distributions.
- Barnard, G.A.** 1947. "Significance tests for 2 x 2 tables." *Biometrika* 34:123-138.
- Bartholomew, D.J.** 1963. "The Sampling Distribution of an Estimate Arising in Life Testing." *Technometrics*, Volume 5 No. 3, 361-374.
- Beal, S. L.** 1987. "Asymptotic Confidence Intervals for the Difference between Two Binomial Parameters for Use with Small Samples." *Biometrics*, Volume 43, Issue 4, 941-950.
- Benjamini, Y. and Hochberg, Y.** 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing," *Journal of the Royal Statistical Society, Series B (Methodological)*, Vol. 57, No. 1, 289-300.
- Blackwelder, W.C.** 1993. "Sample size and power in prospective analysis of relative risk." *Statistics in Medicine*, Volume 12, 691-698.
- Blackwelder, W.C.** 1998. "Equivalence Trials." In *Encyclopedia of Biostatistics*, John Wiley and Sons. New York. Volume 2, 1367-1372.
- Bonett, Douglas.** 2002. "Sample Size Requirements for Testing and Estimating Coefficient Alpha." *Journal of Educational and Behavioral Statistics*, Vol. 27, pages 335-340.
- Box, G.E.P., Hunter, S. and Hunter, J.S..** 1978. *Statistics for Experimenters*. John Wiley & Sons, New York. This is probably the leading book in the area experimental design in industrial experiments. You definitely should acquire and study this book if you plan anything but a casual acquaintance with experimental design. The book is loaded with examples and explanations.

References-2

- Breslow, N. E. and Day, N. E.** 1980. *Statistical Methods in Cancer Research: Volume 1. The Analysis of Case-Control Studies*. Lyon: International Agency for Research on Cancer.
- Brown, H., and Prescott, R.** 2006. *Applied Mixed Models in Medicine*. 2nd ed. John Wiley & Sons Ltd. Chichester, West Sussex, England.
- Brush, Gary G.** 1988. *Volume 12: How to Choose the Proper Sample Size*, American Society for Quality Control, 310 West Wisconsin Ave, Milwaukee, Wisconsin, 53203. This is a small workbook for quality control workers.

C

- Casagrande, J. T., Pike, M.C., and Smith, P. G.** 1978. "The Power Function of the "Exact" Test for Comparing Two Binomial Distributions," *Applied Statistics*, Volume 27, No. 2, pages 176-180. This article presents the algorithm upon which our Fisher's exact test is based.
- Cochran and Cox.** 1992. *Experimental Designs. Second Edition*. John Wiley & Sons. New York. This is one of the classic books on experimental design, first published in 1957.
- Chen, K.W.; Chow, S.C.; and Li, G.** 1997. "A Note on Sample Size Determination for Bioequivalence Studies with Higher-Order Crossover Designs" *Journal of Pharmacokinetics and Biopharmaceutics*, Volume 25, No. 6, pages 753-765.
- Chen, T. T.** 1997. "Optimal Three-Stage Designs for Phase II Cancer Clinical Trials." *Statistics in Medicine*, Volume 16, pages 2701-2711.
- Chow, S.C. and Liu, J.P.** 1999. *Design and Analysis of Bioavailability and Bioequivalence Studies*. Marcel Dekker. New York.
- Chow, S.C.; Shao, J.; Wang, H.** 2003. *Sample Size Calculations in Clinical Research*. Marcel Dekker. New York.
- Chow, S.-C.; Shao, J.; Wang, H.** 2008. *Sample Size Calculations in Clinical Research, Second Edition*. Chapman & Hall/CRC. Boca Raton, Florida.
- Cohen, Jacob.** 1988. *Statistical Power Analysis for the Behavioral Sciences*, Lawrence Erlbaum Associates, Hillsdale, New Jersey. This is a very nice, clearly written book. There are MANY examples. It is the largest of the sample size books. It does not deal with clinical trials.
- Cohen, Jacob.** 1990. "Things I Have Learned So Far." *American Psychologist*, December, 1990, pages 1304-1312. This is must reading for anyone still skeptical about the need for power analysis.
- Collett, D.** 1991. *Modelling Binary Data*. Chapman & Hall, New York, New York. This book covers such topics as logistic regression, tests of proportions, matched case-control studies, and so on.
- Collett, D.** 1994. *Modelling Survival Data in Medical Research*. Chapman & Hall, New York, New York. This book covers such survival analysis topics as Cox regression and log rank tests.
- Conlon, M. and Thomas, R.** 1993. "The Power Function for Fisher's Exact Test." *Applied Statistics*, Volume 42, No. 1, pages 258-260. This article was used to validate the power calculations of Fisher's Exact Test in PASS. Unfortunately, we could not use the algorithm to improve the speed because the algorithm requires equal sample sizes.
- Cook, R. D. and Weisburg, S.** 1999. *Applied Regression Including Computing and Graphics*. John Wiley and Sons, Inc.
- Cox, D. R.** 1972. "Regression Models and life tables." *Journal of the Royal Statistical Society, Series B*, Volume 34, Pages 187-220. This article presents the proportional hazards regression model.

D

- D'Agostino, R.B., Chase, W., Belanger, A.** 1988. "The Appropriateness of Some Common Procedures for Testing the Equality of Two Independent Binomial Populations.", *The American Statistician*, August 1988, Volume 42 Number 3, pages 198-202.
- Davies, Owen L.** 1971. *The Design and Analysis of Industrial Experiments*. Hafner Publishing Company, New York. This was one of the first books on experimental design and analysis. It has many examples and is highly recommended.
- DeLong, E.R., DeLong, D.M., and Clarke-Pearson, D.L.** 1988. "Comparing the Areas Under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach." *Biometrics*, 44, pages 837-845.
- DeMets, D.L. and Lan, K.K.G.** 1984. "An overview of sequential methods and their applications in clinical trials." *Communications in Statistics, Theory and Methods*, 13, pages 2315-2338.
- DeMets, D.L. and Lan, K.K.G.** 1994. "Interim analysis: The alpha spending function approach." *Statistics in Medicine*, 13, pages 1341-1352.
- Demidenko, E.** 2004. *Mixed Models – Theory and Applications*. John Wiley & Sons. Hoboken, New Jersey.
- Desu, M. M. and Raghavarao, D.** 1990. *Sample Size Methodology*. Academic Press. New York. (Presents many useful results for determining sample sizes.)
- Devroye, Luc.** 1986. *Non-Uniform Random Variate Generation*. Springer-Verlag. New York. This book is currently available online at <http://jeff.cs.mcgill.ca/~luc/rnbookindex.html>.
- Diggle, P.J., Liang, K.Y., and Zeger, S.L.** 1994. *Analysis of Longitudinal Data*. Oxford University Press. New York, New York.
- Dixon, W. J. and Tukey, J. W.** 1968. "Approximate behavior of the distribution of Winsorized t," *Technometrics*, Volume 10, pages 83-98.
- Donnelly, Thomas G.** 1980. "ACM Algorithm 462: Bivariate Normal Distribution," *Collected Algorithms from ACM*, Volume II, New York, New York.
- Donner, Allan.** 1984. "Approaches to Sample Size Estimation in the Design of Clinical Trials--A Review," *Statistics in Medicine*, Volume 3, pages 199-214. This is a well done review of the clinical trial literature. Although it is becoming out of date, it is still a good place to start.
- Donner, A. and Klar, N.** 1996. "Statistical Considerations in the Design and Analysis of Community Intervention Trials." *The Journal of Clinical Epidemiology*, Vol. 49, No. 4, 1996, pages 435-439.
- Donner, A. and Klar, N.** 2000. *Design and Analysis of Cluster Randomization Trials in Health Research*. Arnold. London.
- Draghici, S.** 2003. *Data Analysis Tools for DNA Microarrays*. Chapman & Hall/CRC. London. This is an excellent overview of most areas of Microarray analysis.
- Dudoit, S., Yang, Y.H., Callow, M.J., and Speed, T.P.** 2002. "Statistical Methods for Identifying Differentially Expressed Genes in Replicated cDNA Experiments," *Statistica Sinica*, Volume 12, pages 111-139.
- Dunn, O. J.** 1961. "Multiple comparisons among means," *Journal of the American Statistical Association*, Volume 56, pages 52-64.
- Dunn, O. J.** 1964. "Multiple comparisons using rank sums," *Technometrics*, Volume 6, pages 241-252.
- Dunnett, C. W.** 1955. "A Multiple comparison procedure for Comparing Several Treatments with a Control," *Journal of the American Statistical Association*, Volume 50, pages 1096-1121.
- Dupont, William.** 1988. "Power Calculations for Matched Case-Control Studies," *Biometrics*, Volume 44, pages 1157-1168.

References-4

Dupont, William and Plummer, Walton D. 1990. "Power and Sample Size Calculations--A Review and Computer Program," *Controlled Clinical Trials*, Volume 11, pages 116-128. Documents a nice public-domain program on sample size and power analysis.

E

Edgington, E. 1987. *Randomization Tests*. Marcel Dekker. New York. A comprehensive discussion of randomization tests with many examples.

Edwards, L.K. 1993. *Applied Analysis of Variance in the Behavior Sciences*. Marcel Dekker. New York. Chapter 8 of this book is used to validate the repeated measures module of PASS.

Efron, B. 1971. "Forcing a Sequential Experiment to be Balanced." *Biometrika*. Volume 58, pages 403-417.

Efron, B. and Tibshirani, R. J. 1993. *An Introduction to the Bootstrap*. Chapman & Hall. New York.

Elandt-Johnson, R.C. and Johnson, N.L. 1980. *Survival Models and Data Analysis*. John Wiley. NY, NY. This book devotes several chapters to population and clinical life-table analysis.

Epstein, Benjamin. 1960. "Statistical Life Test Acceptance Procedures." *Technometrics*. Volume 2.4, pages 435-446.

F

Farrington, C. P. and Manning, G. 1990. "Test Statistics and Sample Size Formulae for Comparative Binomial Trials with Null Hypothesis of Non-Zero Risk Difference or Non-Unity Relative Risk." *Statistics in Medicine*, Vol. 9, pages 1447-1454. This article contains the formulas used for the Equivalence of Proportions module in PASS.

Feldt, L.S., Woodruff, D.J., and Salih, F.A. 1987. "Statistical inference for coefficient alpha." *Applied Psychological Measurement*, Vol. 11, pages 93-103.

Feldt, L.S. and Ankenmann, R.D. 1999. "Determining Sample Size for a Test of the Equality of Alpha Coefficients When the Number of Part-Tests is Small." *Psychological Methods*, Vol. 4(4), pages 366-377.

Fisher, R. A. 1921. "On the probable error of a coefficient of correlation deduced from a small sample." *Metron*, i (4), 1-32.

Flack, V. F., Afifi, A. A., Lachenbruch, P. A., and Schouten, H. J. A. 1988. "Sample Size Determinations for the Two Rater Kappa Statistic." *Psychometrika*, Volume 53, No. 3, pages 321-325.

Fleiss, Joseph L. 1981. *Statistical Methods for Rates and Proportions*. John Wiley & Sons. New York. This book provides a very good introduction to the subject.

Fleiss, Joseph L. 1986. *The Design and Analysis of Clinical Experiments*. John Wiley & Sons. New York. This book provides a very good introduction to clinical trials. It may be a bit out of date now, but it is still very useful.

Fleiss, J. L., Levin, B., Paik, M.C. 2003. *Statistical Methods for Rates and Proportions. Third Edition*. John Wiley & Sons. New York. This book provides a very good introduction to the subject.

Fleming, T. R. 1982. "One-sample multiple testing procedure for Phase II clinical trials." *Biometrics*, Volume 38, pages 143-151.

Freedman, L.S. 1982. "Tables of the Number of Patients Required in Clinical Trials using the Logrank Test." *Statistics in Medicine*, 1:121-129.

G

Gans. 1984. "The Search for Significance: Different Tests on the Same Data." *The Journal of Statistical Computation and Simulation*, 1984, pages 1-21.

Gart, John J. and Nam, Jun-mo. 1988. "Approximate Interval Estimation of the Ratio in Binomial Parameters: A Review and Corrections for Skewness." *Biometrics*, Volume 44, Issue 2, 323-338.

Gart, John J. and Nam, Jun-mo. 1990. "Approximate Interval Estimation of the Difference in Binomial Parameters: Correction for Skewness and Extension to Multiple Tables." *Biometrics*, Volume 46, Issue 3, 637-643.

Gehlback, Stephen. 1988. *Interpreting the Medical Literature: Practical Epidemiology for Clinicians*. Second Edition. McGraw-Hill. New York. Telephone: (800)722-4726. The preface of this book states that its purpose is to provide the reader with a useful approach to interpreting the quantitative results given in medical literature. We reference it specifically because of its discussion of ROC curves.

Gentle, James E. 1998. *Random Number Generation and Monte Carlo Methods*. Springer. New York.

Gibbons, J. 1976. *Nonparametric Methods for Quantitative Analysis*. Holt, Rinehart and Winston. New York.

Gibbons, J. 1985. *Nonparametric Methods for Quantitative Analysis (2nd Edition)*. American Sciences Press. New York.

Goldstein, Richard. 1989. "Power and Sample Size via MS/PC-DOS Computers," *The American Statistician*, Volume 43, Number 4, pages 253-260. A comparative review of power analysis software that was available at that time.

Greenwood, J. A. and Sandomire, M. M. 1950. "Sample Size Required for Estimating the Standard Deviation as a Per Cent of its True Value", *Journal of the American Statistical Association*, Vol. 45, No. 250, pp. 257-260.

Griffiths, P. and Hill, I.D. 1985. *Applied Statistics Algorithms*, The Royal Statistical Society, London, England. See page 243 for ACM algorithm 291.

Gross and Clark 1975. *Survival Distributions: Reliability Applications in Biomedical Sciences*. John Wiley, New York.

Guenther, William C. 1977. "Desk Calculation of Probabilities for the Distribution of the Sample Correlation Coefficient," *The American Statistician*, Volume 31, Number 1, pages 45-48.

Guenther, William C. 1977. *Sampling Inspection in Statistical Quality Control*. Griffin's Statistical Monographs, Number 37. London.

H

Hahn, G. J. and Meeker, W.Q. 1991. *Statistical Intervals*. John Wiley & Sons. New York.

Hanley, J. A. and McNeil, B. J. 1982. "The Meaning and Use of the Area under a Receiver Operating Characteristic (ROC) Curve." *Radiology*, 143, 29-36. April, 1982.

References-6

- Hanley, J. A. and McNeil, B. J.** 1983. "A Method of Comparing the Areas under Receiver Operating Characteristic Curves Derived from the Same Cases." *Radiology*, 148, 839-843. September, 1983.
- Harris, M., Horvitz, D. J., and Mood, A. M.** 1948. "On the Determination of Sample Sizes in Designing Experiments", *Journal of the American Statistical Association*, Volume 43, No. 243, pp. 391-402.
- Hernandez-Bermejo, B. and Sorribas, A.** 2001. "Analytical Quantile Solution for the S-distribution, Random Number Generation and Statistical Data Modeling." *Biometrical Journal* 43, 1007-1025.
- Hoaglin, Mosteller, and Tukey.** 1985. *Exploring Data Tables, Trends, and Shapes*. John Wiley. New York.
- Hochberg, Y. and Tamhane, A. C.** 1987. *Multiple Comparison Procedures*. John Wiley & Sons. New York.
- Howe, W.G.** 1969. "Two-Sided Tolerance Limits for Normal Populations—Some Improvements." *Journal of the American Statistical Association*, 64, 610-620.
- Hosmer, D. and Lemeshow, S.** 1989. *Applied Logistic Regression*. John Wiley & Sons. New York. This book gives an advanced, in depth look at logistic regression.
- Hosmer, D. and Lemeshow, S.** 1999. *Applied Survival Analysis*. John Wiley & Sons. New York.
- Hotelling, H.** 1933. "Analysis of a complex of statistical variables into principal components." *Journal of Educational Psychology* 24, 417-441, 498-520.
- Hsieh, F.Y.** 1989. "Sample Size Tables for Logistic Regression," *Statistics in Medicine*, Volume 8, pages 795-802. This is the article that was the basis for the sample size calculations in logistic regression in PASS 6.0. It has been superceded by the 1998 article.
- Hsieh, F.Y., Block, D.A., and Larsen, M.D.** 1998. "A Simple Method of Sample Size Calculation for Linear and Logistic Regression," *Statistics in Medicine*, Volume 17, pages 1623-1634. The sample size calculation for logistic regression in PASS are based on this article.
- Hsieh, F.Y. and Lavori, P.W.** 2000. "Sample-Size Calculations for the Cox Proportional Hazards Regression Model with Nonbinary Covariates," *Controlled Clinical Trials*, Volume 21, pages 552-560. The sample size calculation for Cox regression in PASS are based on this article.
- Hsu, Jason.** 1996. *Multiple Comparisons: Theory and Methods*. Chapman & Hall. London. This book gives a beginning to intermediate discussion of multiple comparisons, stressing the interpretation of the various MC tests. It provides details of three popular MC situations: all pairs, versus the best, and versus a control. The power calculations used in the MC module of PASS came from this book.

J

- Johnson, N.L., Kotz, S., and Kemp, A.W.** 1992. *Univariate Discrete Distributions, Second Edition*. John Wiley & Sons. New York.
- Johnson, N.L., Kotz, S., and Balakrishnan, N.** 1994. *Continuous Univariate Distributions Volume 1, Second Edition*. John Wiley & Sons. New York.
- Johnson, N.L., Kotz, S., and Balakrishnan, N.** 1995. *Continuous Univariate Distributions Volume 2, Second Edition*. John Wiley & Sons. New York.
- Julious, Steven A.** 2004. "Tutorial in Biostatistics. Sample sizes for clinical trials with Normal data." *Statistics in Medicine*, 23:1921-1986.
- Jung, Sin-Ho.** 2005. "Sample size for FDR-control in microarray data analysis" *Bioinformatics*, 21(14):3097-3104.

Jung, Sin-Ho; Kang, Sun J.; McCall, Linda M.; Blumenstein, Brent. 2005. "Sample Sizes Computation for Two-Sample Noninferiority Log-Rank Test", *J. of Biopharmaceutical Statistics*, Volume 15, pages 969-979.

Juran, J.M. 1979. *Quality Control Handbook*. McGraw-Hill. New York.

K

Kalbfleisch, J.D. and Prentice, R.L. 1980. *The Statistical Analysis of Failure Time Data*. John Wiley, New York.

Karian, Z.A and Dudewicz, E.J. 2000. *Fitting Statistical Distributions*. CRC Press, New York.

Katz, D., Baptista, J., Azen, S. P., and Pike, M. C. 1978. "Obtaining Confidence Intervals for the Risk Ratio in Cohort Studies," *Biometrics*, 34, pages 469-474.

Kendall, M. and Stuart, A. 1987. *Kendall's Advanced Theory of Statistics. Volume 1: Distribution Theory*. Oxford University Press. New York. This is a fine math-stat book for graduate students in statistics. We reference it because it includes formulas that are used in the program.

Kenward, M. G. and Roger, J. H. 1997. "Small Sample Inference for Fixed Effects from Restricted Maximum Likelihood," *Biometrics*, 53, pages 983-997.

Kirk, Roger E. 1982. *Experimental Design: Procedures for the Behavioral Sciences*. Brooks/Cole. Pacific Grove, California. This is a respected reference on experimental design and analysis of variance.

Klein, J.P. and Moeschberger, M.L.. 1997. *Survival Analysis*. Springer-Verlag. New York. This book provides a comprehensive look at the subject complete with formulas, examples, and lots of useful comments. It includes all the more recent developments in this field. I recommend it.

Koch, G.G.; Atkinson, S.S.; Stokes, M.E. 1986. *Encyclopedia of Statistical Sciences*. Volume 7. John Wiley. New York. Edited by Samuel Kotz and Norman Johnson. The article on Poisson Regression provides a very good summary of the subject.

Kraemer, H. C. and Thiemann, S. 1987. *How Many Subjects*, Sage Publications, 2111 West Hillcrest Drive, Newbury Park, CA. 91320. This is an excellent introduction to power analysis.

Kupper, L. L. and Hafner, K. B. 1989. 'How Appropriate are Popular Sample Size Formulas?', *The American Statistician*, Volume 43, No. 2, pp. 101-105.

L

Lachin, John M. 2000. *Biostatistical Methods*. John Wiley & Sons. New York. This is a graduate-level methods book that deals with statistical methods that are of interest to biostatisticians such as odds ratios, relative risks, regression analysis, case-control studies, and so on.

Lachin, John M. and Foulkes, Mary A. 1986. "Evaluation of Sample Size and Power for Analyses of Survival with Allowance for Nonuniform Patient Entry, Losses to Follow-up, Noncompliance, and Stratification," *Biometrics*, Volume 42, September, pages 507-516.

Lakatos, Edward. 1988. "Sample Sizes Based on the Log-Rank Statistic in Complex Clinical Trials", *Biometrics*, Volume 44, March, pages 229-241.

Lakatos, Edward. 2002. "Designing Complex Group Sequential Survival Trials", *Biometrika*, Volume 70, pages 1969-1989.

Lan, K.K.G. and DeMets, D.L. 1983. "Discrete sequential boundaries for clinical trials." *Biometrika*, 70, pages 659-663.

References-8

- Lan, K.K.G. and Zucker, D.M.** 1993. "Sequential monitoring of clinical trials: the role of information and Brownian motion." *Statistics in Medicine*, 12, pages 753-765.
- Lance, G.N. and Williams, W.T.** 1967. "A general theory of classificatory sorting strategies. I. Hierarchical systems." *Comput. J.* 9, pages 373-380.
- Lance, G.N. and Williams, W.T.** 1967. "Mixed-data classificatory programs I. Agglomerative systems." *Aust. Comput. J.* 1, pages 15-20.
- Lawless, J.F.** 1982. *Statistical Models and Methods for Lifetime Data*. John Wiley, New York.
- Lawson, John.** 1987. *Basic Industrial Experimental Design Strategies*. Center for Statistical Research at Brigham Young University. Provo, Utah. 84602. This is a manuscript used by Dr. Lawson in courses and workshops that he provides to industrial engineers. It is the basis for many of our experimental design procedures.
- Lee, E.T.** 1980. *Statistical Methods for Survival Data Analysis*. Lifetime Learning Publications. Belmont, California.
- Lee, E.T.** 1992. *Statistical Methods for Survival Data Analysis*. Second Edition. John Wiley & Sons. New York. This book provides a very readable introduction to survival analysis techniques.
- Lee, M.-L. T.** 2004. *Analysis of Microarray Gene Expression Data*. Kluwer Academic Publishers. Norwell, Massachusetts.
- Lenth, Russell V.** 1987. "Algorithm AS 226: Computing Noncentral Beta Probabilities," *Applied Statistics*, Volume 36, pages 241-244.
- Lenth, Russell V.** 1989. "Algorithm AS 243: Cumulative Distribution Function of the Non-central t Distribution," *Applied Statistics*, Volume 38, pages 185-189.
- Lewis, J.A.** 1999. "Statistical principles for clinical trials (ICH E9) an introductory note on an international guideline." *Statistics in Medicine*, 18, pages 1903-1942.
- Lipsey, Mark W.** 1990. *Design Sensitivity Statistical Power for Experimental Research*, Sage Publications, 2111 West Hillcrest Drive, Newbury Park, CA. 91320. This is an excellent introduction to power analysis.
- Littell, R. C. et al.** 2006. *SAS for Mixed Models – Second Edition*. SAS Institute Inc., Cary, North Carolina.
- Liu, J., Hsueh, H., Hsieh, E., and Chen, J.J.** 2002. "Tests for equivalence or non-inferiority for paired binary data," *Statistics in Medicine*, Volume 21, pages 231-245.
- Liu, H. and Wu, T.** 2005. "Sample Size Calculation and Power Analysis of Time-Averaged Difference," *Journal of Modern Applied Statistical Methods*, Vol. 4, No. 2, pages 434-445.
- Lu, Y. and Bean, J.A.** 1995. "On the sample size for one-sided equivalence of sensitivities based upon McNemar's test," *Statistics in Medicine*, Volume 14, pages 1831-1839.
- Locke, C.S.** 1984. "An exact confidence interval for untransformed data for the ratio of two formulation means," *J. Pharmacokinet. Biopharm.*, Volume 12, pages 649-655.

M

- Machin, D., Campbell, M., Fayers, P., and Pinol, A.** 1997. *Sample Size Tables for Clinical Studies*, 2nd Edition. Blackwell Science. Malden, Mass. A very good & easy to read book on determining appropriate sample sizes in many situations.
- Marubini, E. and Valsecchi, M.G.** 1996. *Analysing Survival Data from Clinical Trials and Observational Studies*. John Wiley: New York, New York.
- Matsumoto, M. and Nishimura, T.** 1998. "Mersenne twister: A 623-dimensionally equidistributed uniform pseudorandom number generator" *ACM Trans. On Modeling and Computer Simulations*.

- McClish, D.K.** 1989. "Analyzing a Portion of the ROC Curve." *Medical Decision Making*, 9: 190-195
- Metz, C.E.** 1978. "Basic principles of ROC analysis." *Seminars in Nuclear Medicine*, Volume 8, No. 4, pages 283-298.
- Miettinen, O.S. and Nurminen, M.** 1985. "Comparative analysis of two rates." *Statistics in Medicine* 4: 213-226.
- Montgomery, Douglas.** 1984. *Design and Analysis of Experiments*. John Wiley & Sons, New York. A textbook covering a broad range of experimental design methods. The book is not limited to industrial investigations, but gives a much more general overview of experimental design methodology.
- Moore, D. S. and McCabe, G. P.** 1999. *Introduction to the Practice of Statistics*. W. H. Freeman and Company. New York.
- Moura, Eduardo C.** 1991. *How To Determine Sample Size And Estimate Failure Rate in Life Testing*. ASQC Quality Press. Milwaukee, Wisconsin.
- Mukerjee, H., Robertson, T., and Wright, F.T.** 1987. "Comparison of Several Treatments With a Control Using Multiple Contrasts." *Journal of the American Statistical Association*, Volume 82, No. 399, pages 902-910.
- Muller, K. E., and Barton, C. N.** 1989. "Approximate Power for Repeated-Measures ANOVA Lacking Sphericity." *Journal of the American Statistical Association*, Volume 84, No. 406, pages 549-555.
- Muller, K. E., LaVange, L.E., Ramey, S.L., and Ramey, C.T.** 1992. "Power Calculations for General Linear Multivariate Models Including Repeated Measures Applications." *Journal of the American Statistical Association*, Volume 87, No. 420, pages 1209-1226.
- Muller, K. E. and Stewart, P.W.** 2006. *Linear Model Theory: Univariate, Multivariate, and Mixed Models*. John Wiley & Sons Inc. Hoboken, New Jersey.
- Myers, R.H.** 1990. *Classical and Modern Regression with Applications*. PWS-Kent Publishing Company. Boston, Massachusetts. This is one of the bibles on the topic of regression analysis.

N

- Nam, Jun-mo.** 1987. "A Simple Approximation for Calculating Sample Sizes for Detecting Linear Trend in Proportions," *Biometrics*, Volume 43, pages 701-705.
- Nam, Jun-mo.** 1992. "Sample Size Determination for Case-Control Studies and the Comparison of Stratified and Unstratified Analyses," *Biometrics*, Volume 48, pages 389-395.
- Nam, Jun-mo.** 1997. "Establishing equivalence of two treatments and sample size requirements in matched-pairs design," *Biometrics*, Volume 53, pages 1422-1430.
- Nam, J-m. and Blackwelder, W.C.** 2002. "Analysis of the ratio of marginal probabilities in a matched-pair setting," *Statistics in Medicine*, Volume 21, pages 689-699.
- Neter, J., Kutner, M., Nachtsheim, C., and Wasserman, W.** 1996. *Applied Linear Statistical Models*. Richard D. Irwin, Inc. Chicago, Illinois. This mammoth book covers regression analysis and analysis of variance thoroughly and in great detail. We recommend it.
- Neter, J., Wasserman, W., and Kutner, M.** 1983. *Applied Linear Regression Models*. Richard D. Irwin, Inc. Chicago, Illinois. This book provides you with a complete introduction to the methods of regression analysis. We suggest it to non-statisticians as a great reference tool.
- Newcombe, Robert G.** 1998a. "Two-Sided Confidence Intervals for the Single Proportion: Comparison of Seven Methods." *Statistics in Medicine*, Volume 17, 857-872.

References-10

Newcombe, Robert G. 1998b. "Interval Estimation for the Difference Between Independent Proportions: Comparison of Eleven Methods." *Statistics in Medicine*, Volume 17, 873-890.

Newcombe, Robert G. 1998c. "Improved Confidence Intervals for the Difference Between Binomial Proportions Based on Paired Data." *Statistics in Medicine*, Volume 17, 2635-2650.

O

O'Brien, P.C. and Fleming, T.R. 1979. "A multiple testing procedure for clinical trials." *Biometrics*, 35, pages 549-556.

O'Brien, R.G. and Kaiser, M.K. 1985. "MANOVA Method for Analyzing Repeated Measures Designs: An Extensive Primer." *Psychological Bulletin*, 97, pages 316-333.

Obuchowski, N. 1998. "Sample Size Calculations in Studies of Test Accuracy." *Statistical Methods in Medical Research*, 7, pages 371-392.

Obuchowski, N. and McClish, D. 1997. "Sample Size Determination for Diagnostic Accuracy Studies Involving Binormal ROC Curve Indices." *Statistics in Medicine*, 16, pages 1529-1542.

Odeh, R.E. and Fox, M. 1991. *Sample Size Choice*. Marcel Dekker, Inc. New York, NY.

O'Neill and Wetherill. 1971 "The Present State of Multiple Comparison Methods," *The Journal of the Royal Statistical Society, Series B*, vol.33, 218-250).

Ostle, B. and Malone, L. C. 1988. *Statistics in Research. Fourth Edition*. Iowa State Press. Ames, Iowa. A comprehension book on statistical methods.

Owen, Donald B. 1956. "Tables for Computing Bivariate Normal Probabilities," *Annals of Mathematical Statistics*, Volume 27, pages 1075-1090.

Owen, Donald B. 1965. "A Special Case of a Bivariate Non-Central t-Distribution," *Biometrika*, Volume 52, pages 437-446.

P

Parmar, M.K.B. and Machin, D. 1995. *Survival Analysis*. John Wiley and Sons. New York.

Pan, Z. and Kupper, L. 1999. "Sample Size Determination for Multiple Comparison Studies Treating Confidence Interval Width as Random." *Statistics in Medicine* 18, 1475-1488.

Pearson, E.S. and Hartley, H.O. 1976. *Biometrika Tables For Statistics, Volume 1*. Biometrika Trust. London.

Phillips, Kem F. 1990. "Power of the Two One-Sided Tests Procedure in Bioequivalence," *Journal of Pharmacokinetics and Biopharmaceutics*, Volume 18, No. 2, pages 137-144.

Piantadosi, S. 2005. *Clinical Trials – A Methodological Perspective*. John Wiley & Sons. New Jersey.

Pocock, S.J. 1977. "Group sequential methods in the design and analysis of clinical trials." *Biometrika*, 64, pages 191-199.

Pocock, S.J. 1983. *Clinical Trials – A Practical Approach*. John Wiley & Sons. New York.

Price, K., Storn R., and Lampinen, J. 2005. *Differential Evolution – A Practical Approach to Global Optimization*. Springer. Berlin, Germany.

Prihoda, Tom. 1983. "Convenient Power Analysis For Complex Analysis of Variance Models." *Poster Session of the American Statistical Association Joint Statistical Meetings*, August 15-18, 1983, Toronto, Canada. Tom is currently at the University of Texas Health Science Center. This article includes FORTRAN code for performing power analysis.

R

- Ramsey, Philip H.** 1978 "Power Differences Between Pairwise Multiple Comparisons," *JASA*, vol. 73, no. 363, pages 479-485.
- Rao, C.R. , Mitra, S.K., & Matthai, A.** 1966. *Formulae and Tables for Statistical Work*. Statistical Publishing Society, Indian Statistical Institute, Calcutta, India.
- Reboussin, D.M., DeMets, D.L., Kim, K, and Lan, K.K.G.** 1992. "Programs for computing group sequential boundaries using the Lan-DeMets Method." Technical Report 60, Department of Biostatistics, University of Wisconsin-Madison.
- Rencher, Alvin C.** 1998. *Multivariate Statistical Inference and Applications*. John Wiley. New York, New York. This book provides a comprehensive mixture of theoretical and applied results in multivariate analysis. My evaluation may be biased since Al Rencher took me fishing when I was his student.
- Robins, Greenland, and Breslow.** 1986. "A General Estimator for the Variance of the Mantel-Haenszel Odds Ratio," *American Journal of Epidemiology*, vol.42, pages 719-723.
- Robins, Breslow, and Greenland.** 1986. "Estimators of the Mantel-Haenszel variance consistent in both sparse data and large-strata limiting models," *Biometrics*, vol. 42, pages 311-323.
- Rosenberger, W.F., and Lachin, J.M.** 2002. *Randomization in Clinical Trials – Theory and Practice*. John Wiley & Sons. New York.

S

- Sachs, Lothar.** 1984. *Applied Statistics: A Handbook of Techniques*. Springer-Verlag. New York, New York.
- Sahai, Hardeo & Khurshid, Anwer.** 1995. *Statistics in Epidemiology*. CRC Press. Boca Raton, Florida.
- Schilling, Edward.** 1982. *Acceptance Sampling in Quality Control*. Marcel-Dekker. New York.
- Schlesselman, Jim.** 1981. *Case-Control Studies*. Oxford University Press. New York. This presents a complete overview of case-control studies. It was our primary source for the Mantel-Haenszel test.
- Schoenfeld, David A.** 1983. "Sample-Size Formula for the Proportional-Hazards Regression Model" *Biometrics*, Volume 39, pages 499-503.
- Schoenfeld, David A. and Richter, Jane R.** 1982. "Nomograms for Calculating the Number of Patients Needed for a Clinical Trial with Survival as an Endpoint," *Biometrics*, March 1982, Volume 38, pages 163-170.
- Schork, M. and Williams, G.** 1980. "Number of Observations Required for the Comparison of Two Correlated Proportions." *Communications in Statistics-Simula. Computa.*, B9(4), 349-357.
- Schuirmann, Donald.** 1981. "On hypothesis testing to determine if the mean of a normal distribution is contained in a known interval," *Biometrics*, Volume 37, pages 617.
- Schuirmann, Donald.** 1987. "A Comparison of the Two One-Sided Tests Procedure and the Power Approach for Assessing the Equivalence of Average Bioavailability," *Journal of Pharmacokinetics and Biopharmaceutics*, Volume 15, Number 6, pages 657-680.
- Senn, Stephen.** 1993. *Cross-over Trials in Clinical Research*. John Wiley & Sons. New York.
- Senn, Stephen.** 2002. *Cross-over Trials in Clinical Research*. Second Edition. John Wiley & Sons. New York.

References-12

- Shuster, Jonathan J.** 1990. *CRC Handbook of Sample Size Guidelines for Clinical Trials*. CRC Press, Boca Raton, Florida. This is an expensive book (\$300) of tables for running log-rank tests. It is well documented, but at this price it better be.
- Signorini, David.** 1991. "Sample size for Poisson regression," *Biometrika*, Volume 78, 2, pages 446-450.
- Simon, Richard.** "Optimal Two-Stage Designs for Phase II Clinical Trials," *Controlled Clinical Trials*, 1989, Volume 10, pages 1-10.
- Smith, R.L.** 1984. "Sequential Treatment Allocation using Biased Coin Designs." *Journal of the Royal Statistical Society B*. Volume 46, pages 519-543.
- Stekel, D.** 2003. *Microarray Bioinformatics*. Cambridge University Press. Cambridge, United Kingdom.
- Statxact 5.** 2001. *Statistical Software for exact nonparametric inference, user manual*. Cytel Software Corporation. Cambridge, Massachusetts.
- Swets, John A.** 1996. *Signal Detection Theory and ROC Analysis in Psychology and Diagnostics - Collected Papers*. Lawrence Erlbaum Associates. Mahway, New Jersey.

T

- Tango, Toshiro.** 1998. "Equivalence Test and Confidence Interval for the Difference in Proportions for the Paired-Sample Design." *Statistics in Medicine*, Volume 17, 891-908.
- Therneau, T.M. and Grambsch, P.M.** 2000. *Modeling Survival Data*. Springer: New York, New York. At the time of the writing of the Cox regression procedure, this book provides a thorough, up-to-date discussion of this procedure as well as many extensions to it. Recommended, especially to those with at least a masters in statistics.
- Thode, Henry C.** 2002. *Testing for Normality*. Marcel Dekker, Inc. New York.
- Thompson, Simon G.** 1998. *Encyclopedia of Biostatistics, Volume 4*. John Wiley & Sons. New York. Article on Meta-Analysis on pages 2570-2579.
- Tubert-Bitter, P., Manfredi, R., Lellouch, J., Begaud, B.** 2000. "Sample size calculations for risk equivalence testing in pharmacoepidemiology." *Journal of Clinical Epidemiology* 53, 1268-1274.
- Tukey, J.W. and McLaughlin, D.H.** 1963. "Less Vulnerable confidence and significance procedures for location based on a single sample: Trimming/Winsorization." *Sankhya, Series A* 25, 331-352.

U

- Upton, G.J.G.** 1982. "A Comparison of Alternative Tests for the 2 x 2 Comparative Trial.", *Journal of the Royal Statistical Society, Series A*, Volume 145, pages 86-105.

W

- Walter, S.D., Eliasziw, M., and Donner, A.** 1998. "Sample Size and Optimal Designs For Reliability Studies." *Statistics in Medicine*, 17, 101-110.
- Wei, L.J., and Lachin, J.M.** 1988. "Properties of the Urn Randomization in Clinical Trials." *Controlled Clinical Trials*. Volume 9, pages 345-364.

- Welch, B.L.** 1938. "The significance of the difference between two means when the population variances are unequal." *Biometrika*, 29, 350-362.
- Westlake, W.J.** 1981. "Bioequivalence testing—a need to rethink," *Biometrics*, Volume 37, pages 591-593.
- Whittemore, Alice.** 1981. "Sample Size for Logistic Regression with Small Response Probability," *Journal of the American Statistical Association*, Volume 76, pages 27-32.
- Wilson, E.B..** 1927. "Probable Inference, the Law of Succession, and Statistical Inference," *Journal of the American Statistical Association*, Volume 22, pages 209-212. This article discusses the 'score' method that has become popular when dealing with proportions.
- Winer, B.J.** 1991. *Statistical Principles in Experimental Design (Third Edition)*. McGraw-Hill. New York, NY. A very complete analysis of variance book.
- Wolfinger, R., Tobias, R. and Sall, J.** 1994. "Computing Gaussian likelihoods and their derivatives for general linear mixed models," *SIAM Journal of Scientific Computing*, 15, no.6, pages 1294-1310.
- Woolson, R.F., Bean, J.A., and Rojas, P.B.** 1986. "Sample Size for Case-Control Studies Using Cochran's Statistic," *Biometrics*, Volume 42, pages 927-932.

Y

- Yateman, Nigel A. and Skene, Allan M.** 1992. "Sample Sizes for Proportional Hazards Survival Studies with Arbitrary Patient Entry and Loss to Follow-Up Distributions." *Statistics in Medicine*, 11:1103-1113.
- Yuen, K.K. and Dixon, W. J.** 1973. "The approximate behavior and performance of the two-sample trimmed t," *Biometrika*, Volume 60, pages 369-374.
- Yuen, K.K.** 1974. "The two-sample trimmed t for unequal population variances," *Biometrika*, Volume 61, pages 165-170.

Z

- Zar, Jerrold H.** 1984. *Biostatistical Analysis (Second Edition)*. Prentice-Hall. Englewood Cliffs, New Jersey. This introductory book presents a nice blend of theory, methods, and examples for a long list of topics of interest in biostatistical work.
- Zhou, X., Obuchowski, N., McClish, D.** 2002. *Statistical Methods in Diagnostic Medicine*. John Wiley & Sons, Inc. New York, New York. This is a great book on the designing and analyzing diagnostic tests. It is especially useful for its presentation of ROC curves.

Index

Index entries are of the form “chapter-page”. A list of chapters is given in the Table of Contents.

3

3D parameters, 4-18

A

Abbreviations, 4-15
 Abort, 4-5
 Absolute deviation
 standard deviation estimator, 905-3
 Accrual time
 logrank, 705-6
 A-efficiency
 D-optimal designs, 888-13
 Agreement between two raters - kappa, 811-1
 Alias
 two-level designs, 881-2
 All-contrasts power
 simulation, 590-4
 All-pair power
 multiple comparisons - simulation, 580-4
 simulation, 585-4
 Alpha, 7-3
 adjusting, 9-5
 Alpha - simulation
 multiple comparisons, 580-1, 585-2
 multiple contrasts, 590-1
 Alpha spending function
 survival - group sequential, 710-1
 two means - group sequential, 475-1
 two proportions - group sequential, 220-1
 Alternative hypothesis, 7-1, 9-1
 Analysis of variance
 fixed effects ANOVA, 560-1
 one-way, 550-1
 one-way - simulation, 555-1
 randomized block, 565-1
 repeated measures, 570-1
 Anderson-Darling test, 670-2
 ANOVA
 factorial, 560-1
 fixed effects, 560-1
 multiple comparisons, 575-2
 multiple comparisons - simulation, 580-2
 multiple contrasts - simulation, 590-2
 one-way, 550-1
 one-way - simulation, 555-1
 randomized block, 565-1
 repeated measures, 570-1
 three-way, 560-1
 two-way, 560-1
 ANOVA - fixed effects
 examples, 560-13
 validation, 560-19
 ANOVA - one-way
 assumptions, 550-2
 examples, 550-10
 validation, 550-19
 ANOVA - one-way - simulation
 examples, 555-10
 validation, 555-16
 ANOVA - randomized block
 examples, 565-9
 validation, 565-11
 ANOVA - repeated measures
 examples, 570-29
 validation, 570-49
 Any-contrast power
 simulation, 590-4
 Any-pair power
 multiple comparisons - simulation, 580-4
 simulation, 585-4
 Assigning subjects to groups, 880-1
 AUC
 ROC curve (one), 260-3
 ROC curves (two), 265-1, 265-6
 Autocorrelation
 repeated measures ANOVA, 570-25
 Average absolute deviation
 standard deviation estimator, 905-3
 Axes/Legend/Grid tab, 4-11
 Axis
 format, 4-14
 maximum, 4-11
 minimum, 4-11
 parameters, 4-11
 range, 4-11
 Axis color, 4-12

B

Balaam's design, 530-1, 535-1, 540-1, 545-1
 Balance, 880-1

Index-2

Balanced incomplete block designs, 883-1
 examples, 883-4
Bar chart options, 4-16
Beta
 calculating, 7-4
Beta distribution
 probability calculator, 915-1
 simulation, 920-3
Between standard deviation
 repeated measures ANOVA, 570-14
BIB designs, 883-1
Bimodal data
 simulation, 920-23
Binomial distribution
 probability calculator, 915-2
 simulation, 920-5
Binomial model
 one proportion, 100-2
Binomial probabilities, 7-2
Binomial test
 one proportion, 100-4
Binormal model
 ROC curve (one), 260-2
 ROC curves (two), 265-2
Bioequivalence
 two proportions, 215-1
Bivariate normal distribution
 probability calculator, 915-2
Blackwelder
 risk ratio, 205-26
Blocking
 two-level designs, 881-2
Bonferroni adjustment
 t-test - one group, 610-5
 t-test - two groups, 615-4
Bonferroni test
 multiple contrasts - simulation, 590-3
Bootstrap test
 simulation, 410-3
Bootstrap test - paired means
 simulation, 490-4
Bootstrap test - paired means - equivalence
 simulation, 495-5
Box-Behnken designs, 885-1
Brownian motion
 survival - group sequential, 710-5
 two means - group sequential, 475-4
 two proportions - group sequential, 220-4
Buttons
 output window, 5-7, 6-2
 PASS home window, 3-3
 procedure window, 4-6

C

Cage
 edge color, 4-18
 flip, 4-18
 thin walls, 4-18
 wall color, 4-18
Candidate points
 D-optimal designs, 888-14
Carryover effects
 repeated measures ANOVA, 570-3
Case-control
 matched, 155-1
Cauchy distribution
 simulation, 920-5
Central-composite designs, 885-1
Changing fonts, 5-6
Chen
 three-stage phase II trials, 130-1
Chi-square
 estimator, 900-1
 one variance, 650-1
Chi-square distribution
 probability calculator, 915-2
Chi-Square estimator
 examples, 900-2
Chi-square test, 250-1
 examples, 250-6
 two proportions, 200-5
 validation, 250-9
Circularity
 repeated measures ANOVA, 570-13
Clinical trial
 three-stage, 130-1
Clinical trials
 one proportion, 120-1
Clopper-Pearson confidence interval, 115-2
Cluster randomized - equivalence
 two proportions, 240-1
Cluster randomized - non-inferiority
 two proportions, 235-1
Cluster randomized - superiority
 two proportions, 235-1
Cluster randomized design
 two means, 480-1
 two proportions, 230-1
Cochran-Armitage test, 255-1
 examples, 255-12
 validation, 255-19
Cochran-Mantel-Haenszel, 225-1
Coefficient alpha
 one, 815-1
 two, 820-1
Coefficient alpha (one)
 examples, 815-4
 validation, 815-7
Coefficient alphas (two)
 examples, 820-6
 validation, 820-10
Coefficient of variation
 cross-over, 505-2
 cross-over - equivalence, 525-3
 cross-over - higher-order - equivalence, 545-5
 cross-over - higher-order - non-inferiority, 535-5
 cross-over - non-inferiority, 515-3
 mean ratio, 445-2
 mean ratio - equivalence, 470-3

- mean ratio - non-inferiority, 455-2
 - standard deviation estimator, 905-9
- Cohort study
 - post-marketing surveillance, 135-2
- Color
 - axis, 4-12
 - grid lines, 4-12
 - legend, 4-14
 - symbols, 4-19
- Color max, 4-19
- Color min, 4-18
- Color palette, 4-18
- Column widths, 925-11
- Comparisons
 - multiple comparisons, 575-1
 - one-way ANOVA, 550-1
- Comparisonwise error rate
 - simulation, 580-1, 585-2, 590-1
- Complete randomization, 880-3
- Compound symmetry
 - repeated measures ANOVA, 570-13
- Confidence interval
 - one proportion, 115-1
 - paired means, 496-1
- Confidence intervals
 - one correlation, 801-1
 - one mean, 420-1
 - one mean - tolerance probability, 421-1
 - one standard deviation, 640-1
 - one standard deviation - relative error, 642-1
 - one standard deviation - tolerance probability, 641-1
 - one variance, 651-1
 - one variance - relative error, 653-1
 - one variance - tolerance probability, 652-1
 - paired means - tolerance probability, 497-1
 - slope - simple linear regression, 856-1
 - two means, 471-1
 - two means - tolerance probability, 472-1
 - two proportions - difference, 216-1
 - two proportions - odds ratio, 216-1
 - two proportions - ratio, 216-1
 - variance ratio, 656-1
 - variance ratio - relative error, 657-1
- Confounding
 - two-level designs, 881-2
- Constant distribution
 - simulation, 920-6
- Consumer's risk
 - exponential mean (one), 405-2
- Contaminated normal simulation, 920-20
- Contingency table, 250-1
 - chi-square estimator, 900-1
- Continuity correction
 - one proportion, 100-5
 - two proportions - group sequential, 220-7
 - two proportions - offset, 205-7
- Contrast matrix
 - repeated measures ANOVA, 570-10
- Contrasts
 - MANOVA, 605-4
 - mixed models, 201-6, 431-7
 - one-way ANOVA, 550-1
- Contrasts - multiple
 - simulation, 590-1
- Controlled variables
 - multiple regression, 865-4
- Correlated proportions
 - matched case-control, 155-1
 - McNemar test, 150-1
- Correlation
 - intraclass, 810-1
- Correlation (one)
 - confidence interval, 801-1
 - examples, 800-4
 - validation, 800-7
- Correlation (one) - confidence interval
 - examples, 801-4
 - validation, 801-7
- Correlation coefficient, 800-1
- Correlation coefficient distribution
 - probability calculator, 915-3
- Correlation test
 - one, 800-1
 - two, 805-1
- Correlations (two)
 - examples, 805-5
 - validation, 805-8
- Covariance matrix
 - repeated measures ANOVA, 570-13
- Covariance patterns
 - AR(1), 201-11, 431-4
 - banded, 201-11, 431-4
 - compound symmetry, 201-11, 431-4
 - simple, 201-12
 - simple, 431-5
- Cox regression, 850-1
 - examples, 850-4
 - validation, 850-6
- Creating data
 - simulation, 920-1
- Creating randomization lists, 880-1
- Cronbach's alpha
 - one, 815-1
 - two, 820-1
- Crossed factors
 - design generator, 889-1
- Cross-over
 - higher-order - equivalence, 540-1
 - higher-order - non-inferiority, 530-1
 - higher-order - superiority, 530-1
 - ratio - equivalence, 525-1
 - ratio - higher-order - equivalence, 545-1
 - ratio - higher-order - non-inferiority, 535-1
 - ratio - higher-order - superiority, 535-1
 - ratio - non-inferiority, 515-1
 - ratio - superiority, 515-1
 - repeated measures ANOVA, 570-44
 - two means, 500-1
 - two means - equivalence, 520-1
 - two means - non-inferiority, 510-1
 - two means - ratio, 505-1

Index-4

- two means - superiority, 510-1
- Cross-over - higher-order - equivalence
 - examples, 540-7
 - validation, 540-10
- Cross-over - higher-order - non-inferiority
 - examples, 530-9
 - validation, 530-12
- Cross-over - ratio - equivalence
 - examples, 525-7
 - validation, 525-9
- Cross-over - ratio - higher-order - equivalence
 - examples, 545-9
 - validation, 545-12
- Cross-over - ratio - higher-order - non-inferiority
 - examples, 535-10
 - validation, 535-13
- Cross-over - ratio - non-inferiority
 - examples, 515-6
 - validation, 515-8
- Cross-over (two means)
 - examples, 500-7
 - validation, 500-10
- Cross-over (two means) - equivalence
 - examples, 520-7
 - validation, 520-10
- Cross-over (two means) - non-inferiority
 - examples, 510-7
 - validation, 510-10
- Cross-over (two means) - ratio
 - examples, 505-6
 - validation, 505-8
- Crossover analysis
 - mixed models, 571-1

D

- D'Agostino kurtosis, 670-2
- D'Agostino omnibus test, 670-3
- D'Agostino skewness test, 670-4
- Data
 - simulation of, 920-1
- Data entry, 925-1, 925-6
- Data simulation
 - examples, 920-19
- Data simulator, 920-1
- Data tab, 4-7
- Default template, 4-2
- D-efficiency
 - D-optimal designs, 888-12
- Delete template button, 4-4, 4-22
- Depth, 4-18
- Design generator, 889-1
 - examples, 889-3
- Design of experiments
 - randomization lists, 880-1
- Designs
 - Box-Behnken, 885-1
 - central-composite, 885-1
 - design generator, 889-1

- factorial, 881-2
- fractional factorial, 882-1
- Plackett-Burman, 886-1
- response surface, 885-1
- screening, 886-1
- Taguchi, 887-1
- two-level factorial, 881-1, 889-1
- Determinant
 - D-optimal designs, 888-13
- Determinant analysis
 - D-optimal designs, 888-11
- Diagnostic testing
 - ROC curve (one), 260-1
- Difference
 - proportions, 8-2
- Difference data
 - simulation, 920-26
- Discordant pairs
 - McNemar test, 150-2
- Distributions
 - combining, 920-13
 - mixing, 920-13
 - simulation, 920-1
- Documentation
 - printing, 1-9
- Donner and Klar
 - cluster randomized design, 230-2
- D-optimal designs, 888-1
 - examples, 888-8
- Drift
 - survival - group sequential, 710-5
 - two means - group sequential, 475-4
 - two proportions - group sequential, 220-4
- Dunn's test
 - multiple contrasts - simulation, 590-3
- Dunnett's test
 - multiple comparisons, 575-3
- Dunnett's test
 - simulation, 585-1, 585-3
- Dunn's test
 - simulation, 590-1
- Duplicates
 - D-optimal designs, 888-4

E

- Edge color, 4-18
- Edit menu
 - output window, 5-4
 - spreadsheet, 925-3
- Effect size, 7-5
 - ANOVA, 560-2, 560-7
 - chi-square estimator, 900-1
 - chi-square test, 250-2
 - multiple regression, 865-2
 - one-way ANOVA, 550-3
 - randomized block ANOVA, 565-2
- Efron's biased coin randomization, 880-3
- Entering matrices, 925-1

- Entering procedure options, 4-1
- Equivalence
 - correlated proportions, 165-1
 - means (two), 460-1
 - means (two) - ratio, 470-1
 - means (two) - simulation, 465-1
 - one proportion, 110-1
 - two proportions, 215-1
- Equivalence - two correlated proportions
 - examples, 165-11
 - validation, 165-14
- Equivalence hypothesis, 7-9
- Equivalence limits - paired means
 - simulation, 495-7
- Equivalence margin
 - two means, 450-8
 - two means - non-inferiority, 455-5
- Equivalence test data
 - simulating, 495-3
- Error rates - simulation
 - multiple comparison, 580-1, 585-2
 - multiple contrasts, 590-1
- Errors, 7-2
- Exact binomial test
 - one proportion, 100-4
- Examples
 - balanced incomplete block designs, 883-4
 - Chi-square estimator, 900-2
 - Chi-square test, 250-6
 - Cochran-Armitage test, 255-12
 - correlation (one) - confidence interval, 801-4
 - Cox regression, 850-4
 - cross-over - higher-order - equivalence, 540-7
 - cross-over - higher-order - non-inferiority, 530-9
 - cross-over - ratio - equivalence, 525-7
 - cross-over - ratio - higher-order - equivalence, 545-9
 - cross-over - ratio - higher-order - non-inferiority, 535-10
 - cross-over - ratio - non-inferiority, 515-6
 - cross-over (two means), 500-7
 - cross-over (two means) - equivalence, 520-7
 - cross-over (two means) - non-inferiority, 510-7
 - cross-over (two means) - ratio, 505-6
 - data simulation, 920-19
 - design generator, 889-3
 - D-optimal designs, 888-8
 - equivalence - two correlated proportions, 165-11
 - exponential mean (one), 405-8
 - exponential means (two), 435-5
 - fixed effects ANOVA, 560-13
 - fractional factorial designs, 882-4
 - Hotelling's T₂, 600-8
 - intraclass correlation, 810-4
 - kappa, 811-6
 - Latin square designs, 884-5
 - linear regression, 855-5
 - logistic regression, 860-9
 - logrank tests - non-inferiority, 706-7
 - logrank tests (Lakatos), 715-16
 - Mann-Whitney test, 430-16
 - MANOVA, 605-12
 - many proportions - trend, 255-12
 - matched case-control - proportions, 155-6
 - McNemar test - two correlated proportions, 150-7
 - mean ratio - non-inferiority, 455-6
 - microarray one-sample or paired t-test, 610-12
 - microarray two-sample t-test, 615-12
 - mixed models, 571-35
 - multiple comparisons, 575-13
 - multiple comparisons - simulation, 580-13
 - multiple comparisons - vs control - simulation, 585-12
 - multiple contrasts - simulation, 590-14
 - multiple regression, 865-6
 - non-inferiority - two correlated proportions, 160-10
 - normality - simulation, 670-8
 - odds ratio estimator, 910-2
 - one coefficient alpha, 815-4
 - one correlation, 800-4
 - one mean - confidence interval, 420-5
 - one mean - tolerance - confidence interval, 421-6
 - one proportion, 100-13
 - one proportion - confidence interval, 115-6
 - one proportion - equivalence, 110-11
 - one proportion - non-inferiority, 105-10
 - one standard deviation - confidence interval, 640-5
 - one standard deviation - relative error - confidence interval, 642-4
 - one standard deviation - tolerance - confidence interval, 641-6
 - one variance - confidence interval, 651-4
 - one variance - relative error - confidence interval, 653-4
 - one variance - tolerance - confidence interval, 652-6
 - one-way ANOVA, 550-10
 - one-way ANOVA - simulation, 555-10
 - paired means - confidence interval, 496-5
 - paired means - tolerance - confidence interval, 497-6
 - Poisson regression, 870-7
 - post-marketing surveillance, 135-6
 - proportions estimator, 910-2
 - randomization lists, 880-9
 - randomized block ANOVA, 565-9
 - ratio of two means, 445-6
 - ratio of two means - equivalence, 470-7
 - regression - confidence interval, 856-7
 - repeated measures - two means, 431-13
 - repeated measures - two proportions, 201-16
 - repeated measures ANOVA, 570-29
 - response surface designs, 885-3
 - ROC curve (one), 260-8
 - ROC curves (two), 265-8
 - screening designs, 886-3
 - single-stage phase II trials, 120-3
 - standard deviation estimator, 905-6
 - survival - group sequential, 710-11
 - survival - logrank, 700-6
 - survival - logrank - advanced, 705-8

Index-6

- Taguchi designs, 887-4
- three-stage phase II trials, 130-8
- t-test (one mean), 400-8
- t-test (one mean) - non-inferiority, 415-9
- t-test (one mean) - simulation, 410-10
- t-test (paired means) - equivalence - simulation, 495-11
- t-test (paired means) - simulation, 490-11
- t-test (paired), 400-13
- t-test (two means), 430-10
- t-test (two means) - simulation, 440-12
- two coefficient alphas, 820-6
- two correlated proportions - equivalence, 165-11
- two correlated proportions - non-inferiority, 160-10
- two correlations, 805-5
- two means - cluster randomized, 480-4
- two means - confidence interval, 471-5
- two means - equivalence, 460-5
- two means - equivalence - simulation, 465-12
- two means - group sequential, 475-10
- two means - non-inferiority, 450-10
- two means - ratio, 445-6
- two means - ratio - equivalence, 470-7
- two means - tolerance - confidence interval, 472-6
- two proportions - cluster - equivalence, 240-8
- two proportions - cluster - non-inferiority, 235-7
- two proportions - cluster randomized, 230-9
- two proportions - confidence interval, 216-26
- two proportions - equivalence, 215-17
- two proportions - group sequential, 220-9
- two proportions - inequality, 200-14
- two proportions - non-inferiority, 210-17
- two proportions - offset, 205-18
- two proportions - stratified design, 225-8
- two-level designs, 881-6
- two-stage phase II trials, 125-8
- variance (one), 650-4
- variance ratio - confidence interval, 656-5
- variance ratio - relative error - confidence interval, 657-4
- variances (two), 655-4
- Exiting PASS, 4-4
- Experiment (run)
 - two-level designs, 881-2
- Experimental design, 881-1
 - two-level designs, 881-2
- Experimental error
 - two-level designs, 881-2
- Experimentwise error rate
 - simulation, 580-1, 585-2, 590-1
- Exponential
 - logrank, 705-3
 - log-rank, 715-2
- Exponential distribution
 - simulation, 920-6
- Exponential mean (one), 405-1
 - examples, 405-8
 - validation, 405-10
- Exponential means (two), 435-1
 - examples, 435-5

- validation, 435-7
- Exponential test
 - simulation, 410-4
- Exposure
 - Poisson regression, 870-1
- Exposure probability
 - matched case-control, 155-4

F

- F distribution
 - probability calculator, 915-3
 - simulation, 920-7
- Factor scaling
 - D-optimal designs, 888-2
- Factorial ANOVA, 560-1
- Factorial designs
 - two-level designs, 881-1, 881-2
- False discovery rate adjustment
 - t-test
 - two groups, 615-4
 - t-test - one group, 610-6
- Familywise error rate
 - simulation, 580-1, 585-2, 590-1
- Farrington-Manning confidence interval
 - two proportions, 216-5
- Farrington-Manning test
 - two proportions - equivalence, 215-9
 - two proportions - non-inferiority, 210-10
 - two proportions - offset, 205-9
- File menu
 - output window, 5-2
 - procedure window, 4-3
 - spreadsheet, 925-1
- Finite population correction
 - t test, 400-8
- Finite population size
 - one mean - confidence interval, 420-2
- Fisher's exact test
 - two proportions, 200-4
- Fisher-z transformation
 - two correlations, 805-1
- Fixed effects ANOVA
 - examples, 560-13
 - validation, 560-19
- Fixed effects models
 - mixed models, 571-1
- Fixed factor
 - ANOVA, 560-5
- Fleiss Confidence intervals
 - two proportions, 216-12, 216-16
- Fleming algorithm
 - single-stage phase II trials, 120-1
- Follow-up
 - logrank, 705-4
- Fonts
 - changing, 5-6
- Forced points
 - D-optimal designs, 888-5

Format
 tick labels, 4-14
 Format menu
 output window, 5-6
 Format toolbar, 5-5
 FPR
 ROC curve (one), 260-7
 ROC curves (two), 265-1
 Fractional factorial designs, 882-1
 examples, 882-4
 Freedman
 logrank, 700-1
 F-test
 Geisser-Greenhouse, 570-4
 simulation, 555-1
 two variances, 655-1

G

Games-Howell - simulation
 multiple comparison, 580-1
 Games-Howell test
 simulation, 580-3
 Gamma data
 simulation, 920-24
 Gamma distribution
 probability calculator, 915-4
 simulation, 920-7
 Gart-Nam confidence interval
 two proportions, 216-7
 Gart-Nam test
 two proportions - equivalence, 215-10
 two proportions - non-inferiority, 210-11
 two proportions - offset, 205-10
 Geisser-Greenhouse
 repeated measures ANOVA, 570-1
 Geisser-Greenhouse F-test
 repeated measures ANOVA, 570-4
 General linear multivariate model
 MANOVA, 605-2
 repeated measures ANOVA, 570-3
 Generating data, 920-1
 Goodness of fit
 chi-square estimator, 900-1
 Chi-square test, 250-1
 Graeco-Latin square designs, 884-1
 Grid color, 4-12
 Grid line style, 4-12
 Grid lines, 4-12
 Group sequential test
 log-rank, 710-1
 survival, 710-1
 two means, 475-1
 two proportions, 220-1

H

Hazard rate
 Cox regression, 850-1
 Hazard rate parameterization
 logrank, 715-2
 Hazard rates
 logrank, 700-2
 logrank - advanced, 705-2
 time dependent, 715-22
 Hazard ratio
 group sequential, 710-2
 Help menu
 output window, 5-7
 PASS home window, 3-3
 procedure window, 4-6
 spreadsheet, 925-5
 Help system, 1-5
 contents window, 1-7
 index window, 1-6
 navigating, 1-6
 printing documentation, 1-9
 search window, 1-8
 Heterogeneous variances
 mixed models, 571-46
 Home window, 3-1
 Horizontal viewing angle, 4-17
 Hotelling's T2, 600-1
 examples, 600-8
 validation, 600-11
 Hotelling's T2 distribution
 probability calculator, 915-4
 Hotelling-Lawley trace, 605-1
 MANOVA, 605-3
 repeated measures ANOVA, 570-1, 570-7
 Hypergeometric distribution
 probability calculator, 915-4
 Hypergeometric model
 one proportion, 100-2
 Hypothesis
 equivalence, 7-9
 inequality, 7-7
 introduction, 7-1
 means, 9-1
 non-inferiority, 7-7
 superiority, 7-8
 types, 7-6
 Hypothesis testing
 introduction, 7-1

I

Icons
 output window, 5-7, 6-2
 PASS home window, 3-3
 procedure window, 4-6
 Incidence rate
 post-marketing surveillance, 135-2

Index-8

- Inclusion points
 - D-optimal designs, 888-6
- Independence test, 250-1
- Inequality
 - one proportion, 100-1
 - two means ratio, 445-1
 - two proportions, 200-1
 - two proportions - offset, 205-1
- Inequality hypothesis, 7-7
- Installation, 1-1
 - folders, 1-1
- Interaction
 - two-level designs, 881-3
- Intercept
 - linear regression, 855-1
- Interim analysis
 - survival, 710-1
 - three-stage trial, 130-1
 - two means, 475-1
 - two proportions, 220-1
- Intraclass correlation, 810-1
 - examples, 810-4
 - validation, 810-6
- Intraclass correlation
 - cluster randomization - two means, 480-3
 - cluster randomized - non-inferiority, 235-4
 - cluster randomized design, 230-2
- Introduction to power analysis, 7-1
- Isometric, 4-17
- Iterations
 - maximum, 4-20
- Iterations tab, 4-20

K

- Kappa
 - examples, 811-6
 - validation, 811-11
- Kappa test for rater agreement, 811-1
- Kenward and Roger method
 - mixed models, 571-11
- Kolmogorov-Smirnov test, 670-2
- Kruskal-Wallis
 - multiple comparisons - simulation, 580-1
 - simulation, 555-1
- Kruskal-Wallis test
 - multiple comparisons - simulation, 580-3, 585-1
 - simulation, 585-3

L

- Labels of plots, 4-13
- Lachin and Foulkes
 - logrank test, 705-1
- Lakatos
 - logrank, 715-1
- Lan-DeMets

- survival - group sequential, 710-1
- two means - group sequential, 475-1
- two proportions - group sequential, 220-1
- Latin square
 - ANOVA, 560-17
- Latin square designs, 884-1
 - examples, 884-5
- Legend
 - color, 4-14
 - parameter, 4-12
 - percent of vertical space, 4-12
 - position, 4-12
- Likelihood ratio test
 - two proportions, 200-7
- Likert-scale
 - simulation, 410-21, 920-8
 - simulation, 920-22
- Lilliefors' critical values
 - normality tests, 670-2
- Line chart options, 4-16
- Linear model
 - ANOVA, 560-2
 - repeated measures ANOVA, 570-3
- Linear regression, 855-1
 - confidence interval, 856-1
 - examples, 855-5
 - validation, 855-7
- Load template button, 4-4, 4-21
- Loading a procedure, 2-2
- Log file, 5-2
- Log transformation
 - cross-over, 505-2
 - cross-over - equivalence, 525-2
 - cross-over - higher-order - equivalence, 545-4
 - cross-over - higher-order - non-inferiority, 535-4
 - cross-over - non-inferiority, 515-2
 - mean ratio, 445-2
 - mean ratio - equivalence, 470-2
 - mean ratio - non-inferiority, 455-2
- Logistic regression, 860-1
 - examples, 860-9
 - validation, 860-14
- Logit
 - logistic regression, 860-2
- Logrank
 - hazard rate parameterization, 715-2
 - median survival time parameterization, 715-2
 - mortality parameterization, 715-2
 - proportion surviving parameterization, 715-2
- Log-rank
 - group sequential test, 710-1
- Logrank procedure comparison, 715-3
- Logrank test, 700-1, 715-1
 - Lachin and Foulkes, 705-1
 - non-inferiority, 706-1
- Logrank tests - non-inferiority
 - examples, 706-7
 - validation, 706-11
- Logrank tests (Lakatos)
 - examples, 715-16
 - validation, 715-21

Longitudinal data models
 mixed models, 571-2
 Longitudinal models, 571-1

M

Macros, 930-1
 command list, 930-19
 commands, 930-6
 examples, 930-20
 syntax, 930-2
 Mann-Whitney test, 430-1
 assumptions, 430-4
 examples, 430-16
 non-inferiority, 450-5
 simulation, 440-5
 Mann-Whitney test - equivalence
 simulation, 465-5
 Mann-Whitney test - simulation
 equivalence, 465-1
 MANOVA, 605-1
 assumptions, 605-1
 examples, 605-12
 validation, 605-17
 Mantel Haenszel test
 two proportions, 200-7
 Mantel-Haenszel, 225-1
 Many proportions - trend
 examples, 255-12
 validation, 255-19
 Map window, 6-1
 Margin of equivalence
 two means, 450-8
 two means - non-inferiority, 455-5
 Martinez-Iglewicz test, 670-3
 Matched case-control, 155-1
 post-marketing surveillance, 135-2
 Matched case-control - proportions
 examples, 155-6
 validation, 155-10
 Matrices, 925-1
 Maximum iterations, 4-20
 Maximum likelihood
 mixed models, 571-8
 Maximum on axis, 4-11
 McNemar test, 150-1
 McNemar test - two correlated proportions
 examples, 150-7
 validation, 150-9
 Mean (one)
 confidence interval, 420-1
 exponential, 405-1
 simulation, 410-1
 Mean (one) - confidence interval
 examples, 420-5
 validation, 420-7
 Mean (one) - tolerance - confidence interval
 examples, 421-6
 validation, 421-8
 Mean (one) - tolerance probability
 confidence interval, 421-1
 Mean absolute deviation
 standard deviation estimator, 905-3
 Mean ratio
 equivalence, 470-1
 inequality, 445-1
 Means
 introduction, 9-1
 one-way - simulation, 555-1
 Means - ratio - equivalence
 cross-over, 525-1
 Means - ratio - equivalence - higher-order
 cross-over, 545-1
 Means - ratio - non-inferiority
 cross-over, 515-1
 Means - ratio - non-inferiority - higher-order
 cross-over, 535-1
 Means - ratio - superiority
 cross-over, 515-1
 Means - ratio - superiority - higher-order
 cross-over, 535-1
 Means (paired)
 confidence interval, 496-1
 simulation, 490-1
 Means (paired) - confidence interval
 examples, 496-5
 validation, 496-6
 Means (paired) - equivalence
 simulation, 495-1
 Means (paired) - tolerance - confidence interval
 examples, 497-6
 validation, 497-8
 Means (paired) - tolerance probability
 confidence interval, 497-1
 Means (two)
 cluster randomized design, 480-1
 confidence interval, 471-1
 cross-over, 500-1
 equivalence, 460-1
 exponential, 435-1
 group sequential test, 475-1
 interim analysis, 475-1
 non-inferiority, 450-1
 ratio, 445-1
 simulation, 440-1
 T-test, 430-1
 t-test - equivalence, 460-1
 T-test - non-inferiority, 450-1
 Means (two) - cluster randomized design
 examples, 480-4
 validation, 480-6
 Means (two) - confidence interval
 examples, 471-5
 validation, 471-7
 Means (two) - equivalence
 cross-over, 520-1
 examples, 460-5
 validation, 460-7
 Means (two) - equivalence - higher-order
 cross-over, 540-1

Index-10

- Means (two) - equivalence - simulation
 - examples, 465-12
 - validation, 465-19
- Means (two) - group sequential
 - examples, 475-10
 - validation, 475-17
- Means (two) - non-inferiority
 - cross-over, 510-1
 - examples, 450-10
 - validation, 450-13
- Means (two) - non-inferiority - higher-order
 - cross-over, 530-1
- Means (two) - ratio
 - cross-over, 505-1
 - equivalence, 470-1
 - non-inferiority, 455-1
 - superiority, 455-1
- Means (two) - ratio - non-inferiority
 - examples, 455-6
 - validation, 455-8
- Means (two) - simulation
 - equivalence, 465-1
 - t-test - equivalence, 465-1
- Means (two) - superiority
 - cross-over, 510-1
- Means (two) - superiority - higher-order
 - cross-over, 530-1
- Means (two) - tolerance - confidence interval
 - examples, 472-6
 - validation, 472-8
- Means (two) - tolerance probability
 - confidence interval, 472-1
- Means matrix
 - MANOVA, 605-4
 - repeated measures ANOVA, 570-7
- Measurement error
 - randomized block ANOVA, 565-2
- Median survival time parameterization
 - logrank, 715-2
- Menus
 - output window, 5-2
 - PASS home window, 3-2
 - procedure window, 4-3
 - spreadsheet, 925-1
- Microarray data
 - one sample t-test, 610-1
 - paired T-test, 610-1
 - two sample t-test, 615-1
- Microarray one-sample or paired t-test
 - examples, 610-12
 - validation, 610-19
- Microarray two-sample t-test
 - examples, 615-12
 - validation, 615-19
- Miettinen-Nurminen confidence interval
 - two proportions, 216-6
- Miettinen-Nurminen test
 - two proportions - equivalence, 215-7
 - two proportions - non-inferiority, 210-8
 - two proportions - offset, 205-7
- Minimum detectable difference
 - mixed models, 571-14
 - multiple comparisons, 575-11
 - one-way ANOVA, 550-17
 - t-test, 400-12
 - two-sample t test, 430-14
- Minimum on axis, 4-11
- Mixed model
 - defined, 571-2
- Mixed models, 571-1
 - differential evolution, 571-12
 - examples, 571-35
 - F test, 571-11
 - Fisher scoring, 571-12
 - fixed effects, 571-1
 - G matrix, 571-5
 - heterogeneous variances, 571-46
 - Kenward and Roger method, 571-11
 - L matrix, 571-9
 - likelihood formulas, 571-8
 - maximum likelihood, 571-8
 - minimum detectable difference, 571-14
 - MIVQUE, 571-12
 - Newton-Raphson, 571-12
 - pairwise contrasts, 201-6, 431-7
 - R matrix, 571-5
 - restricted maximum likelihood, 571-9
 - simulation steps, 571-17
 - types, 571-1
 - validation, 571-41
- Mixture data
 - simulation, 920-25
- Mixture design
 - D-optimal designs, 888-20
- Monte Carlo, 9-3
- Monte Carlo simulation, 920-1
- Mortality parameterization
 - logrank, 715-2
- Moving data, 925-10
- MTBF
 - exponential mean (one), 405-1
- Multinomial
 - chi-square estimator, 900-3
- Multinomial distribution
 - simulation, 920-8
- Multiple comparisons, 575-1
 - Dunnett's test - simulation, 585-1
 - examples, 575-13
 - Games-Howell - simulation, 580-1
 - pair-wise - simulation, 580-1
 - validation, 575-19
 - vs control - simulation, 585-1
 - with a control, 575-2
 - with best, 575-5
- Multiple comparisons - simulation
 - examples, 580-13
 - power, 580-4
 - validation, 580-22
- Multiple comparisons - vs control - simulation
 - examples, 585-12
 - validation, 585-21
- Multiple contrasts

- simulation, 590-1
- Multiple contrasts - simulation
 - examples, 590-14
 - validation, 590-24
- Multiple regression, 865-1
 - examples, 865-6
 - validation, 865-10
- Multiple testing adjustment
 - t-test - two groups, 615-2
 - t-test -one group, 610-4
- Multivariate analysis of variance, 605-1

N

- Navigating the help system, 1-6
- Negative binomial distribution
 - probability calculator, 915-5
- Nested factors
 - design generator, 889-1
- New template, 4-3
- Noncentrality
 - one-way ANOVA, 550-3
- Noncentrality parameter
 - one-way ANOVA, 550-4
- Non-inferiority
 - correlated proportions, 160-1
 - logrank, 705-13, 706-1
 - means (two), 450-1
 - means (two) - ratio, 455-1
 - one proportion, 105-1
 - paired t-test, 415-1
 - survival, 706-1
 - two proportions, 210-1
- Non-inferiority - two correlated proportions
 - examples, 160-10
 - validation, 160-13
- Non-inferiority hypothesis, 7-7
- Non-Inferiority test (two means)
 - simulation, 440-17
- Nonparametric
 - t-test, 400-7
- Nonparametric tests
 - Wilcoxon test, 400-1
- Normal distribution
 - probability calculator, 915-5
 - simulation, 920-9, 920-19
- Normality - simulation
 - examples, 670-8
 - validation, 670-11
- Normality tests, 670-1
 - Anderson-Darling test, 670-2
 - D'Agostino kurtosis, 670-2
 - D'Agostino omnibus, 670-3
 - D'Agostino skewness, 670-4
 - Kolmogorov-Smirnov, 670-2
 - Lilliefors' critical values, 670-2
 - Martinez-Iglewicz, 670-3
 - range, 670-4
 - Shapiro-Wilk, 670-4
- Nuisance parameter, 9-2
- Nuisance parameters, 7-6
- Null hypothesis, 7-1, 9-1

O

- O'Brien-Fleming
 - survival - group sequential, 710-1
 - two means - group sequential, 475-1
 - two proportions - group sequential, 220-1
- Odds ratio
 - logistic regression, 860-2
 - matched case-control, 155-1
 - McNemar test, 150-3
 - one proportion, 100-6
 - proportions, 8-3
 - two proportions, 200-3
- Odds ratio estimator, 910-1
 - examples, 910-2
- One proportion
 - equivalence, 110-1
 - inequality, 100-1
 - non-inferiority, 105-1
 - superiority, 105-1
- One-sample t-test
 - microarray data, 610-1
- One-way ANOVA, 550-1
 - examples, 550-10
 - simulation, 555-1
 - validation, 550-19
- One-way ANOVA - simulation
 - examples, 555-10
 - validation, 555-16
- Open template, 4-3
- Options, 4-4
- Orthogonal arrays, 887-1
- Orthogonal sets of Latin squares, 884-2
- Outliers
 - multiple comparisons - simulation, 580-23
 - one-way ANOVA - simulation, 555-18
- Outliers (two means)
 - simulation, 440-19
- Outline
 - PASS home window, 3-4
- Output, 2-4
- Output window, 5-1
 - edit menu, 5-4
 - file menu, 5-2
 - format menu, 5-6
 - help menu, 5-7
 - toolbar, 5-7, 6-2
 - view menu, 5-5
 - window menu, 5-6

P

- P value, 7-4
- Page setup
 - spreadsheet, 925-2
- Paired design
 - microarray data, 610-1
- Paired designs
 - non-inferiority, 415-1
- Paired distributions
 - simulating, 490-2
- Paired means
 - simulation, 490-1
- Paired means - equivalence
 - simulation, 495-1
- Paired proportions
 - equivalence, 165-1
 - non-inferiority, 160-1
- Paired t-test, 400-1
 - assumptions, 400-3
 - microarray data, 610-1
 - non-inferiority, 415-1
 - superiority, 415-1
- Pairwise comparisons
 - multiple comparisons, 575-7
- Pair-wise comparisons
 - simulation, 580-1
- Pairwise contrasts
 - mixed models, 201-6
- Pairwise contrasts
 - mixed models, 431-7
- Panel, 4-1
- Parameters
 - 3D, 4-18
 - abbreviations, 4-15
 - axis, 4-11
 - entering, 4-7
 - legend, 4-12
- PASS help system, 1-5
- PASS home window, 3-1
 - help menu, 3-3
 - outline, 3-4
 - procedure menus, 3-2
 - toolbar, 3-3
 - tools menu, 3-2
 - view menu, 3-2
 - window menu, 3-3
- Password, 5-7
- Patient entry
 - logrank, 705-3
- Perspective, 4-17, 4-18
- Phase I trials
 - definition, 120-1
- Phase II clinical trials
 - single-stage one proportion, 120-1
 - three-stage one proportion, 130-1
 - two-stage one proportion, 125-1
- Phase II trials
 - definition, 120-1
- Pillai-Bartlett trace, 605-1
- MANOVA, 605-3
 - repeated measures ANOVA, 570-1, 570-6
- Plackett-Burman designs, 886-1
- Planned comparisons
 - one-way ANOVA, 550-1
- Plot labels, 4-13
- Plot text tab, 4-13
- Plot titles, 4-13
- Plot type, 4-15
- Plot type tab, 4-15
- Pocock
 - survival - group sequential, 710-1
 - two means - group sequential, 475-1
 - two proportions - group sequential, 220-1
- Poisson distribution
 - probability calculator, 915-5
 - simulation, 920-9
- Poisson incidence
 - post-marketing surveillance, 135-1
- Poisson regression, 870-1
 - examples, 870-7
 - validation, 870-9
- Population size
 - t-test, 400-8
- Post-marketing surveillance, 135-1
 - examples, 135-6
 - validation, 135-8
- Power
 - calculating, 7-4
 - introduction, 7-1
 - means, 9-1
 - multiple comparisons - simulation, 580-4
- Prevalence
 - correlated proportions, 160-2, 165-2
- Printer setup
 - spreadsheet, 925-2
- Printing documentation, 1-9
- Printing output, 5-4
- Probability calculator, 915-1
 - Beta distribution, 915-1
 - Binomial distribution, 915-2
 - Bivariate normal distribution, 915-2
 - Chi-square distribution, 915-2
 - Correlation coefficient distribution, 915-3
 - F distribution, 915-3
 - Gamma distribution, 915-4
 - Hotelling's T2 distribution, 915-4
 - Hypergeometric distribution, 915-4
 - Negative binomial distribution, 915-5
 - Normal distribution, 915-5
 - Poisson distribution, 915-5
 - Student's t distribution, 915-6
 - Studentized range distribution, 915-6
 - Weibull distribution, 915-6
- Procedure menus
 - PASS home window, 3-2
 - procedure window, 4-5
- Procedure options, 4-1
- Procedure window, 4-1
 - file menu, 4-3
 - help menu, 4-6

- procedure menus, 4-5
 - run menu, 4-5
 - tabs, 4-7
 - toolbar, 4-6
 - tools menu, 4-5
 - window menu, 4-5
 - Producer's risk
 - exponential mean (one), 405-2
 - Projection method, 4-17
 - Proportion (one)
 - binomial model, 100-2
 - confidence interval, 115-1
 - continuity correction, 100-5
 - equivalence, 110-1
 - exact binomial test, 100-4
 - examples, 100-13
 - hypergeometric model, 100-2
 - inequality, 100-1
 - non-inferiority, 105-1
 - odds ratio, 100-6
 - saw-tooth power function, 100-16
 - single-stage phase II trials, 120-1
 - superiority, 105-1
 - three-stage phase II trials, 130-1
 - two-stage phase II trials, 125-1
 - validation, 100-20
 - Z test, 100-4
 - Proportion (one) - confidence interval
 - examples, 115-6
 - validation, 115-9
 - Proportion (one) - equivalence
 - examples, 110-11
 - validation, 110-17
 - Z test, 110-6
 - Proportion (one) - non-inferiority
 - examples, 105-10
 - validation, 105-16
 - Proportion surviving parameterization
 - logrank, 715-2
 - Proportion trend test, 255-1
 - Proportional hazards regression, 850-1
 - Proportions
 - comparing, 8-1
 - difference, 8-2
 - interpretation, 8-3
 - introduction, 8-1
 - logistic regression, 860-1
 - odds ratio, 8-3
 - odds ratio estimator, 910-1
 - paired (equivalence), 165-1
 - paired (non-inferiority), 160-1
 - ratio, 8-2
 - Proportions (many)
 - trend, 255-1
 - Proportions (many) - trend
 - examples, 255-12
 - validation, 255-19
 - Proportions (multiple)
 - Chi-square test, 250-1
 - Proportions (two)
 - Chi-square test, 200-5
 - cluster randomized - equivalence, 240-1
 - cluster randomized - non-inferiority, 235-1
 - cluster randomized - superiority, 235-1
 - cluster randomized design, 230-1
 - confidence intervals, 216-1
 - equivalence, 215-1
 - Fisher's exact, 200-4
 - group sequential test, 220-1
 - independent, 200-1
 - independent - equivalence, 215-1
 - independent - non-inferiority & superiority, 210-1
 - independent - offset, 205-1
 - inequality, 200-1
 - inequality - offset, 205-1
 - interim analysis, 220-1
 - matched case control, 155-1
 - McNemar test, 150-1
 - non-inferiority & superiority, 210-1
 - stratified, 225-1
 - Proportions (two) - cluster - equivalence
 - examples, 240-8
 - validation, 240-12
 - Proportions (two) - cluster - non-inferiority
 - examples, 235-7
 - validation, 235-11
 - Proportions (two) - cluster randomized design
 - examples, 230-9
 - validation, 230-17
 - Proportions (two) - confidence interval
 - examples, 216-26
 - validation, 216-28
 - Proportions (two) - equivalence
 - examples, 215-17
 - validation, 215-23
 - Proportions (two) - group sequential
 - examples, 220-9
 - validation, 220-17
 - Proportions (two) - inequality
 - examples, 200-14
 - validation, 200-19
 - Proportions (two) - non-inferiority
 - examples, 210-17
 - validation, 210-24
 - Proportions (two) - offset
 - examples, 205-18
 - validation, 205-24
 - Proportions (two) - stratified design
 - examples, 225-8
 - validation, 225-11
 - Proportions (two) correlated
 - equivalence, 165-1
 - non-inferiority, 160-1
 - Proportions estimator, 910-1
 - examples, 910-2
-
- ## Q
- Qualitative factors
 - D-optimal designs, 888-5

Index-14

Quick launch window, 6-1
Quick-access buttons
 map window, 6-2
 output window, 5-9
 PASS home window, 3-3
 procedure window, 4-6
Quitting PASS, 4-4

R

Random assignment, 880-1
Random effects models, 571-1
 mixed models, 571-1
Random factor
 ANOVA, 560-5
Random numbers, 920-1
 multiple comparisons - simulation, 580-5
 simulation, 585-5, 590-5
Random sorting, 880-3
Random sorting using max % deviation, 880-5
Randomization
 complete, 880-3
 Efron's biased coin, 880-3
 Latin square designs, 884-2
 random sorting, 880-3
 random sorting using max % deviation, 880-5
 Smith, 880-4
 Wei's urn, 880-4
Randomization lists, 880-1
 examples, 880-9
Randomized block ANOVA, 565-1
 examples, 565-9
 validation, 565-11
Randomized block design
 design generator, 889-5
Range on axis, 4-11
Range test, 670-4
Rater agreement - kappa, 811-1
Rating data
 ROC curve (one), 260-4
 ROC curves (two), 265-3
Ratio
 proportions, 8-2
 two means, 445-1
 two means - equivalence, 470-1
Ratio (two means)
 non-inferiority, 455-1
Ratio of two means
 examples, 445-6
 validation, 445-8
Ratio of two means - equivalence
 examples, 470-7
 validation, 470-9
Ratio -two means
 cross-over, 505-1
Regression
 confidence interval, 856-1
 Cox, 850-1
 linear, 855-1

 logistic, 860-1
 multiple, 865-1
 Poisson, 870-1
Regression - confidence interval
 examples, 856-7
 validation, 856-10
Rejection region, 7-3
Repeated measures
 mixed models, 571-1
 two means, 431-1
 two proportions, 201-1
Repeated measures - two means
 examples, 431-13
 validation, 431-19
Repeated measures - two proportions
 examples, 201-16
 validation, 201-24
Repeated measures ANOVA, 570-1
 examples, 570-29
 validation, 570-49
Repeated measures design
 design generator, 889-6
Replication
 two-level designs, 881-3
Reports tab, 4-8
Resetting a template, 4-3
Response surface designs, 885-1
 examples, 885-3
Restricted maximum likelihood
 mixed models, 571-9
Risk ratio
 Blackwelder, 205-26
ROC curve (one), 260-1
 examples, 260-8
 validation, 260-12
ROC curves (two), 265-1
 examples, 265-8
 validation, 265-13
Row heights, 925-11
R-squared
 logistic regression, 860-8
 multiple regression, 865-1
RTF, 5-3
RTF files, 5-1
Ruler, 5-5
Run menu
 procedure window, 4-5
Running a procedure, 2-3
Running PASS, 2-1

S

Sample size
 introduction, 7-1
Save template, 4-4
Save template button, 4-4, 4-22
Saw-tooth power function
 one proportion, 100-16
Scaling factors

- D-optimal designs, 888-2
- Score test
 - two proportions - equivalence, 215-7
 - two proportions - non-inferiority, 210-8
 - two proportions - offset, 205-7
- Screening designs, 886-1
 - examples, 886-3
- Sensitivity
 - correlated proportions, 160-2, 165-2
 - ROC curve (one), 260-1
- Serial numbers, 1-4, 5-7
- Setting options, 4-4
- Shapiro-Wilk test, 670-4
- Show tickmarks, 4-13
- Sign test
 - simulation, 410-3
- Sign test - paired means
 - simulation, 490-4
- Sign test - paired means - equivalence
 - simulation, 495-5
- Significance level, 7-3
 - adjusting, 9-5
- Significance level - simulation
 - multiple comparisons, 580-1, 585-2
 - multiple contrasts, 590-1
- Simon
 - two-stage phase II trials, 125-1
- Simple linear regression
 - confidence interval, 856-1
- Simulation, 9-3, 920-1
 - Beta distribution, 920-3
 - Binomial distribution, 920-5
 - Cauchy distribution, 920-5
 - Constant distribution, 920-6
 - contaminated normal, 920-20
 - Exponential distribution, 920-6
 - F distribution, 920-7
 - Gamma distribution, 920-7
 - Likert-scale, 920-8, 920-22
 - Multinomial distribution, 920-8
 - multiple comparisons, 580-1, 580-5, 585-1, 585-4
 - multiple contrasts, 590-1, 590-4
 - Normal distribution, 920-9, 920-19
 - normality tests, 670-1
 - one mean, 410-1
 - one-way ANOVA, 555-1
 - paired means, 490-1
 - paired means - equivalence, 495-1
 - Poisson distribution, 920-9
 - random number generation, 580-5, 585-5, 590-5
 - size, 9-3
 - skewed distribution, 920-10
 - Student's T distribution, 920-10
 - syntax, 920-13
 - T distribution, 920-10
 - Tukey's lambda distribution, 920-10
 - two means, 440-1
 - two means - equivalence, 465-1
 - Uniform distribution, 920-11
 - Weibull distribution, 920-12
- Simulation steps
 - mixed models, 571-17
- Single-stage design
 - one proportion, 120-1
- Single-stage phase II trials
 - examples, 120-3
 - validation, 120-3
- Skewed data
 - one-way ANOVA - simulation, 555-20
 - simulation, 410-12
- Skewed data (two means)
 - simulation, 440-20
- Skewed distribution
 - simulation, 920-10
- Skewness test, 670-4
- Slope
 - linear regression, 855-1
- Slope - simple linear regression
 - confidence interval, 856-1
- Smith's randomization, 880-4
- Specificity
 - correlated proportions, 160-2, 165-2
 - ROC curve (one), 260-1
- Spending functions
 - survival - group sequential, 710-3
 - two means - group sequential, 475-2
 - two proportions - group sequential, 220-2
- Split plot analysis
 - mixed models, 571-1
- Spreadsheet, 925-1
 - data entry, 925-6
 - edit menu, 925-3
 - file menu, 925-1
 - help menu, 925-5
 - menus, 925-1
 - navigating, 925-6
 - window menu, 925-5
- Standard deviation, 9-2
 - estimator, 905-1
 - interpretation, 905-1
- Standard deviation (one)
 - confidence interval, 640-1
- Standard deviation (one) - confidence interval
 - examples, 640-5
 - validation, 640-7
- Standard deviation (one) - relative error
 - confidence interval, 642-1
- Standard deviation (one) - relative error - confidence interval
 - examples, 642-4
 - validation, 642-6
- Standard deviation (one) - tolerance - confidence interval
 - examples, 641-6
 - validation, 641-8
- Standard deviation (one) - tolerance probability
 - confidence interval, 641-1
- Standard deviation estimator
 - examples, 905-6
- Standard deviation test
 - one, 650-1
 - two, 655-1

- Starting PASS, 1-3, 2-1
- Stratified designs
 - two proportions, 225-1
- Student's t distribution
 - probability calculator, 915-6
- Studentized range distribution
 - probability calculator, 915-6
- Student's T distribution
 - simulation, 920-10
- Style
 - grid lines, 4-12
- Superiority
 - means (two), 450-1
 - means (two) - ratio, 455-1
 - one proportion, 105-1
 - paired t-test, 415-1
 - two proportions, 210-1
- Superiority hypothesis, 7-8
- Surface chart options, 4-17
- Survival
 - logrank, 700-1, 715-1
 - logrank - Lachin and Foulkes, 705-1
 - non-inferiority, 706-1
- Survival - group sequential
 - examples, 710-11
 - validation, 710-18
- Survival - logrank
 - examples, 700-6
 - validation, 700-9
- Survival - logrank - Lachin and Foulkes
 - examples, 705-8
 - validation, 705-12
- Syntax
 - macros, 930-2
- System requirements, 1-1

T

- T distribution
 - simulation, 920-10
- T test
 - two proportions - equivalence, 215-7
 - two proportions - non-inferiority, 210-8
- T2
 - Hotelling's, 600-1
- Tabs
 - axes/legend/grid, 4-11
 - data, 4-7
 - iterations, 4-20
 - plot text, 4-13
 - plot type, 4-15
 - reports, 4-8
 - template, 4-21
- Tabs on the procedure window, 4-7
- Taguchi designs, 887-1
 - examples, 887-4
- Technical support, 1-11
- Template tab, 4-21
- Templates, 4-1
 - automatic, 4-2
 - creating a new, 4-3
 - default, 4-2
 - definition, 4-3
 - deleting, 4-4, 4-22
 - file extension, 4-4, 4-22
 - file name, 4-21
 - loading, 4-3, 4-4, 4-21
 - opening, 4-3
 - saving, 4-4, 4-22
 - storage location, 4-4, 4-22
 - template id, 4-21
- Test statistics, 7-6
- Thin walls, 4-18
- Three-stage design
 - one proportion, 130-1
- Three-stage phase II trials
 - examples, 130-8
 - validation, 130-8
- Tickmarks, 4-12
 - show, 4-13
- Time averaged difference
 - binary data, 201-1
 - normal data, 431-1
 - power for, 201-2, 431-2
 - two means, 431-1
 - two proportions, 201-1
- Titles of plots, 4-13
- Toolbar
 - output window, 5-7, 6-2
 - PASS home window, 3-3
 - procedure window, 4-6
- Toolbars
 - customizing, 3-4
 - customizing using drag-and-drop, 6-3
 - format, 5-5
- Tools menu
 - PASS home window, 3-2
 - procedure window, 4-5
- Trend in proportions, 255-1
- Trimmed t-test (two means)
 - simulation, 440-4
- Trimmed t-test (two means) - equivalence
 - simulation, 465-4
- T-test
 - cross-over, 500-3
 - microarray data, 610-1, 615-1
 - one group - multiple testing adjustment, 610-4
 - one mean, 400-1
 - one mean - simulation, 410-1
 - paired, 400-1
 - paired - equivalence - simulation, 495-1
 - paired - simulation, 490-1
 - two groups - multiple testing adjustment, 615-2
 - two means, 430-1
 - two means - equivalence, 460-1
 - two means - non-inferiority, 450-1
 - two means - simulation, 440-1
 - two proportions, 200-8
- T-test - equivalence
 - cross-over, 520-3

- T-test - simulation
 - two means - equivalence, 465-1
- T-test (one mean)
 - assumptions, 400-3
 - examples, 400-8
 - non-inferiority, 415-1
 - superiority, 415-1
 - validation, 400-17
- T-test (one mean) - non-inferiority
 - examples, 415-9
 - validation, 415-12
- T-test (one mean) - simulation
 - examples, 410-10
 - validation, 410-16
- T-test (paired means) - equivalence - simulation
 - examples, 495-11
 - validation, 495-17
- T-test (paired means) - simulation
 - examples, 490-11
 - validation, 490-16
- T-test (paired)
 - examples, 400-13
- T-test (two means)
 - examples, 430-10
 - validation, 430-18
- T-test (two means) - simulation
 - examples, 440-12
 - validation, 440-16
- T-tests (two means)
 - assumptions, 430-3
- Tukey-Kramer
 - simulation, 580-1
- Tukey-Kramer test
 - multiple comparisons, 575-7
 - simulation, 580-3
- Tukey's lambda distribution
 - simulation, 920-10
- Two correlated proportions - equivalence
 - examples, 165-11
 - validation, 165-14
- Two correlated proportions - non-inferiority
 - examples, 160-10
 - validation, 160-13
- Two means
 - cluster randomized design, 480-1
 - cross-over, 500-1
- Two means - cluster randomized design
 - examples, 480-4
 - validation, 480-6
- Two means - equivalence
 - cross-over, 520-1
 - examples, 460-5
 - validation, 460-7
- Two means - equivalence - higher-order
 - cross-over, 540-1
- Two means - equivalence - simulation
 - examples, 465-12
 - validation, 465-19
- Two means - group sequential
 - examples, 475-10
 - validation, 475-17
- Two means - non-inferiority
 - cross-over, 510-1
 - examples, 450-10
 - validation, 450-13
- Two means - non-inferiority - higher-order
 - cross-over, 530-1
- Two means - ratio
 - cross-over, 505-1
 - examples, 445-6
 - validation, 445-8
- Two means - ratio - equivalence
 - cross-over, 525-1
 - examples, 470-7
 - validation, 470-9
- Two means - ratio - equivalence - higher-order
 - cross-over, 545-1
- Two means - ratio - non-inferiority
 - cross-over, 515-1
- Two means - ratio - non-inferiority - higher-order
 - cross-over, 535-1
- Two means - ratio - superiority
 - cross-over, 515-1
- Two means - ratio - superiority - higher-order
 - cross-over, 535-1
- Two means - superiority
 - cross-over, 510-1
- Two means - superiority - higher-order
 - cross-over, 530-1
- Two means (ratio) - non-inferiority
 - examples, 455-6
 - validation, 455-8
- Two proportions
 - cluster randomized - equivalence, 240-1
 - cluster randomized - non-inferiority, 235-1
 - cluster randomized - superiority, 235-1
 - cluster randomized design, 230-1
 - equivalence, 215-1
 - inequality, 200-1
 - inequality - offset, 205-1
 - matched case-control, 155-1
 - McNemar test, 150-1
 - non-inferiority & superiority, 210-1
- Two proportions - cluster - equivalence
 - examples, 240-8
 - validation, 240-12
- Two proportions - cluster - non-inferiority
 - examples, 235-7
 - validation, 235-11
- Two proportions - cluster randomized design
 - examples, 230-9
 - validation, 230-17
- Two proportions - equivalence
 - examples, 215-17
 - validation, 215-23
- Two proportions - group sequential
 - examples, 220-9
 - validation, 220-17
- Two proportions - inequality
 - examples, 200-14
 - validation, 200-19
- Two proportions - non-inferiority

- examples, 210-17
- validation, 210-24
- Two proportions - offset
 - examples, 205-18
 - validation, 205-24
- Two proportions - stratified design
 - examples, 225-8
 - validation, 225-11
- Two-channel arrays, 610-1
- Two-level designs, 881-1
 - examples, 881-6
- Two-level factorial designs, 881-1
- Two-sample t-test, 430-1
 - equivalence, 460-1
 - microarray data, 615-1
 - non-inferiority, 450-1
 - superiority, 450-1
- Two-sample t-test - simulation
 - equivalence, 465-1
- Two-stage design
 - one proportion, 125-1
- Two-stage phase II trials
 - examples, 125-8
 - validation, 125-10
- Type-I error, 7-2
- Type-II error, 7-2

U

- Uniform distribution
 - simulation, 920-11

V

- Validation
 - Chi-square test, 250-9
 - Cochran-Armitage test, 255-19
 - correlation (one) - confidence interval, 801-7
 - Cox regression, 850-6
 - cross-over - higher-order - equivalence, 540-10
 - cross-over - higher-order - non-inferiority, 530-12
 - cross-over - ratio - equivalence, 525-9
 - cross-over - ratio - higher-order - equivalence, 545-12
 - cross-over - ratio - higher-order - non-inferiority, 535-13
 - cross-over - ratio - non-inferiority, 515-8
 - cross-over (two means), 500-10
 - cross-over (two means) - equivalence, 520-10
 - cross-over (two means) - non-inferiority, 510-10
 - cross-over (two means) - ratio, 505-8
 - equivalence - two correlated proportions, 165-14
 - exponential mean (one), 405-10
 - exponential means (two), 435-7
 - fixed effects ANOVA, 560-19
 - Hotelling's T₂, 600-11
 - intraclass correlation, 810-6

- kappa, 811-11
- linear regression, 855-7
- logistic regression, 860-14
- logrank tests - non-inferiority, 706-11
- logrank tests (Lakatos), 715-21
- MANOVA, 605-17
- many proportions - trend, 255-19
- matched case-control - proportions, 155-10
- McNemar test - two correlated proportions, 150-9
- mean ratio - non-inferiority, 455-8
- microarray one-sample or paired t-test, 610-19
- microarray two-sample t-test, 615-19
- mixed models, 571-41
- multiple comparisons, 575-19
- multiple comparisons - simulation, 580-22
- multiple comparisons - vs control - simulation, 585-21
- multiple contrasts - simulation, 590-24
- multiple regression, 865-10
- non-inferiority - two correlated proportions, 160-13
- normality - simulation, 670-11
- one coefficient alpha, 815-7
- one correlation, 800-7
- one mean - confidence interval, 420-7
- one mean - tolerance - confidence interval, 421-8
- one proportion, 100-20
- one proportion - confidence interval, 115-9
- one proportion - equivalence, 110-17
- one proportion - non-inferiority, 105-16
- one standard deviation - confidence interval, 640-7
- one standard deviation - relative error - confidence interval, 642-6
- one standard deviation - tolerance - confidence interval, 641-8
- one variance - confidence interval, 651-6
- one variance - relative error - confidence interval, 653-6
- one variance - tolerance - confidence interval, 652-8
- one-way ANOVA, 550-19
- one-way ANOVA - simulation, 555-16
- paired means - confidence interval, 496-6
- paired means - tolerance - confidence interval, 497-8
- Poisson regression, 870-9
- post-marketing surveillance, 135-8
- randomized block ANOVA, 565-11
- ratio of two means, 445-8
- ratio of two means - equivalence, 470-9
- regression - confidence interval, 856-10
- repeated measures - two means, 431-19
- repeated measures - two proportions, 201-24
- repeated measures ANOVA, 570-49
- ROC curve (one), 260-12
- ROC curves (two), 265-13
- single-stage phase II trials, 120-3
- survival - group sequential, 710-18
- survival - logrank, 700-9
- survival - logrank - Lachin and Foulkes, 705-12
- three-stage phase II trials, 130-8

- t-test (one mean), 400-17
 - t-test (one mean) - non-inferiority, 415-12
 - t-test (one mean) - simulation, 410-16
 - t-test (paired means) - equivalence - simulation, 495-17
 - t-test (paired means) - simulation, 490-16
 - t-test (two means), 430-18
 - t-test (two means) - simulation, 440-16
 - two coefficient alphas, 820-10
 - two correlated proportions - equivalence, 165-14
 - two correlated proportions - non-inferiority, 160-13
 - two correlations, 805-8
 - two means - cluster randomized, 480-6
 - two means - confidence interval, 471-7
 - two means - equivalence, 460-7
 - two means - equivalence - simulation, 465-19
 - two means - group sequential, 475-17
 - two means - non-inferiority, 450-13
 - two means - ratio, 445-8
 - two means - ratio - equivalence, 470-9
 - two means - tolerance - confidence interval, 472-8
 - two proportions - cluster - equivalence, 240-12
 - two proportions - cluster - non-inferiority, 235-11
 - two proportions - cluster randomized, 230-17
 - two proportions - confidence interval, 216-28
 - two proportions - equivalence, 215-23
 - two proportions - group sequential, 220-17
 - two proportions - inequality, 200-19
 - two proportions - non-inferiority, 210-24
 - two proportions - offset, 205-24
 - two proportions - stratified design, 225-11
 - two-stage phase II trials, 125-10
 - variance (one), 650-7
 - variance ratio - confidence interval, 656-8
 - variance ratio - relative error - confidence interval, 657-6
 - variances (two), 655-8
 - Variance (one)
 - confidence interval, 651-1
 - examples, 650-4
 - validation, 650-7
 - Variance (one) - confidence interval
 - examples, 651-4
 - validation, 651-6
 - Variance (one) - relative error
 - confidence interval, 653-1
 - Variance (one) - relative error - confidence interval
 - examples, 653-4
 - validation, 653-6
 - Variance (one) - tolerance - confidence interval
 - examples, 652-6
 - validation, 652-8
 - Variance (one) - tolerance probability
 - confidence interval, 652-1
 - Variance ratio
 - confidence interval, 656-1
 - Variance ratio - confidence interval
 - examples, 656-5
 - validation, 656-8
 - Variance ratio - relative error
 - confidence interval, 657-1
 - Variance ratio - relative error - confidence interval
 - examples, 657-4
 - validation, 657-6
 - Variance test
 - one, 650-1
 - two, 655-1
 - Variances (two)
 - examples, 655-4
 - validation, 655-8
 - Vertical viewing angle, 4-18
 - View menu
 - output window, 5-5
 - PASS home window, 3-2
 - Viewing angle
 - horizontal, 4-17
 - vertical, 4-18
 - Viewing output, 2-4
-
- ## W
- Wall color, 4-18
 - Walter's confidence intervals
 - two proportions, 216-11
 - Wei's urn randomization, 880-4
 - Weibull distribution
 - probability calculator, 915-6
 - simulation, 920-12
 - Welch test
 - multiple contrasts - simulation, 590-4
 - Welch's test - simulation
 - equivalence, 465-1
 - Welch's t-test
 - non-inferiority, 450-5
 - simulation, 440-3
 - Wilcoxon test, 400-1, 400-7, 415-7, 450-9
 - assumptions, 400-3
 - non-inferiority, 415-1
 - simulation, 410-3
 - superiority, 415-1
 - Wilcoxon test - paired means
 - simulation, 490-4
 - Wilcoxon test - paired means - equivalence
 - simulation, 495-4
 - Wilks' Lambda, 605-1
 - MANOVA, 605-2
 - repeated measures ANOVA, 570-1, 570-6
 - Wilson score limits
 - one proportion, 115-2
 - Wilson's score confidence interval
 - two proportions, 216-7
 - Window menu
 - output window, 5-6
 - PASS home window, 3-3
 - procedure window, 4-5
 - spreadsheet, 925-5
 - Winsorized test (two means) - equivalence
 - simulation, 465-4
 - Within standard deviation

Index-20

repeated measures ANOVA, 570-15
Within-subjects design
repeated measures ANOVA, 570-3
Word processor, 5-1

Z

Z test
one proportion, 100-4
one proportion - equivalence, 110-6
two proportions - equivalence, 215-5
two proportions - non-inferiority, 210-6
two proportions - offset, 205-6