

Chapter 336

Multi-Arm Non-Inferiority Tests for the Ratio of Treatment and Control Means (Log-Normal Data)

Introduction

This module computes power and sample size for multiple non-inferiority tests of treatment means versus a control mean when the data are assumed to follow the log-normal distribution and the statistical hypotheses are expressed in terms of mean ratios.

The details of this t-test are given in Julious (2004). The multiple comparison aspect of this procedure is based on the results in Machin, Campbell, Tan, and Tan (2018).

In this parallel-group design, there are k treatment groups and one control group. A mean is measured in each group. A total of k hypothesis tests are anticipated, each comparing a treatment group with the common control group using a t-test based on the difference of two means of the log-transformed data.

The Bonferroni adjustment of the type I error rate may be optionally made because several comparisons are being tested using the same data. Making a multiplicity adjustment is usually recommended, but not always. In fact, Saville (1990) advocates not applying it and Machin, Campbell, Tan, and Tan (2018) include omitting it as a possibility.

Background

Whether you want to test several doses of a single treatment or several types of treatments, good research practice requires that each treatment be compared with a control. A popular three-arm design consists of three groups: control, treatment A, and treatment B. Two tests are run: treatment A versus control and treatment B versus the same control. This design avoids having to obtain a second control group for treatment B. Besides the obvious efficiency in subjects, it may be easier to recruit subjects if their chances of receiving the new treatment are better than 50-50.

Technical Details

Suppose you want to compare k treatment groups with means μ_i and sample sizes N_i and one control group with mean μ_c and sample size N_c . The total sample size is $N = N_1 + N_2 + \cdots + N_k + N_c$.

Non-Inferiority Tests

A *non-inferiority test* tests that the treatment mean is not worse than the control mean by more than the non-inferiority margin. The actual direction of the hypothesis depends on the response variable being studied. Define $R = \mu_i/\mu_c$.

Case 1: High Values Good

In this case, higher response values are better. The hypotheses are arranged so that rejecting the null hypothesis implies that the treatment mean is no less than a small amount below the control mean. This results in a non-inferiority boundary called R_L . The null and alternative hypotheses with are

$$H_0: R \leq R_L \quad \text{vs.} \quad H_1: R > R_L$$

where $R_L < 1$.

Case 2: High Values Bad

In this case, lower values are better. The hypotheses are arranged so that rejecting the null hypothesis implies that the treatment mean is no more than a small amount above the control mean. This results in a non-inferiority boundary called R_U . The null and alternative hypotheses with are

$$H_0: R \geq R_U \quad \text{vs.} \quad H_1: R < R_U$$

where $R_U > 1$.

Log-Transformation

In many cases, hypotheses stated in terms of ratios are more convenient than hypotheses stated in terms of differences. This is because ratios can be interpreted as percentages, but differences must be interpreted as actual amounts in their original scale. Hence, it has become a common practice to take the following steps in hypothesis testing.

1. State the statistical hypotheses in terms of the ratio of the means.
2. Transform this into hypotheses about a difference by taking logarithms.
3. Analyze the logged data—that is, do the analysis in terms of the difference.
4. Draw the conclusion in terms of the ratio.

Multi-Arm Non-Inferiority Tests for the Ratio of Treatment and Control Means (Log-Normal Data)

The details of step 2 for the two-sided null hypothesis are as follows:

$$H_0: R = R_0 \Rightarrow H_0: \frac{\mu_i}{\mu_c} = R_0 \Rightarrow H_0: \ln(\mu_i) - \ln(\mu_c) = \ln(R_0)$$

Thus, a hypothesis about the ratio of the means on the original scale can be translated into a hypothesis about the difference of two means on the logged scale.

Coefficient of Variation

The coefficient of variation (COV) is the ratio of the standard deviation to the mean. This parameter can be used to represent the variation in the data because of a unique relationship that it has in the case of log-normal data.

Suppose the variable X is the logarithm of the original variable Y . That is, $X = \ln(Y)$ and $Y = \exp(X)$. Label the mean and variance of X as μ_X and σ_X^2 , respectively. Similarly, label the mean and variance of Y as μ_Y and σ_Y^2 , respectively. If X is normally distributed, then Y is log-normally distributed. Julious (2004) presents the following well-known relationships between these two variables

$$\begin{aligned}\mu_Y &= e^{\mu_X + \frac{\sigma_X^2}{2}} \\ \sigma_Y^2 &= \mu_Y^2 (e^{\sigma_X^2} - 1)\end{aligned}$$

From this relationship, the coefficient of variation of Y can be found to be

$$\begin{aligned}COV_Y &= \frac{\sqrt{\mu_Y^2 (e^{\sigma_X^2} - 1)}}{\mu_Y} \\ &= \sqrt{e^{\sigma_X^2} - 1}\end{aligned}$$

Solving this relationship for σ_X^2 , the standard deviation of X can be stated in terms of the coefficient of variation of Y as

$$\sigma_X = \sqrt{\ln(COV_Y^2 + 1)}$$

Similarly, the mean of X is

$$\mu_X = \ln\left(\frac{\mu_Y}{\sqrt{COV_Y^2 + 1}}\right)$$

Thus, the hypotheses can be stated in the original (Y) scale and then the power can be analyzed in the transformed (X) scale. For parallel-group designs, $\sigma_X^2 = \sigma_d^2$, the average variance used in the t-test of the logged data.

Power Calculation

As is shown above, the hypotheses can be stated in the original (Y) scale using ratios or the logged (X) scale using differences. In either case, the power and sample size calculations are made using the formulas for testing the difference in two means. These formulas are presented in another chapter and are not duplicated here.

Multiplicity Adjustment

Because k t-tests between treatment groups and the control group are run when analyzing the results of this study, many statisticians recommend that the Bonferroni adjustment be applied. This adjustment is easy to apply: the value of alpha that is used in the test is found by dividing the original alpha by the number of tests. For example, if the original alpha is set at 0.05 and the number of treatment (not including the control) groups is five, the individual tests will be conducted using an alpha of 0.01.

The main criticism of this procedure is that if there are many tests, the value of alpha becomes very small. To mitigate against this complaint, some statisticians recommend separating the treatment groups into those that are of primary interest and those that are of secondary interest. The Bonferroni adjustment is made by using the number of primary treatments rather than the total number of treatments.

There are some who advocate ignoring the adjustment entirely in the case of randomized clinical trials. See for example Saville (1990) and the discussion in chapter 14 of Machin, Campbell, Tan, and Tan (2018).

Size of the Control Group

Because the control group is used over and over, some advocate increasing the number of subjects in this group. The standard adjustment is to include \sqrt{k} subjects in the control group for each subject in one of the treatment groups. See Machin, Campbell, Tan, and Tan (2018, pages 231-232). Note that often, the treatment groups all have the same size.

Example 1 – Finding the Sample Size

A parallel-group clinical trial is being designed to compare three treatment therapies against the standard therapy. Higher values of the response are desirable. Suppose the standard therapy has a mean response of 10 with a standard deviation of 2.5. The investigators would like a sample size large enough to find statistical significance at the 0.025 level (since this is a one-sided test) if the actual mean responses of the three treatments are 9.5, 10.0, and 10.5, the power of each test is 0.80, and the non-inferiority limit is 0.8. They want to consider a range of standard deviations from 2.0 to 3.0.

Following standard procedure, the control group multiplier will be set to $\sqrt{k} = \sqrt{3} = 1.732$ since the control group is used for three comparisons in this design.

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 1** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For	Sample Size
Higher Means Are	Better (H1: R > RL)
Power of Each Test	0.8
Overall Alpha	0.025
Bonferroni Adjustment	Standard Bonferroni
Group Allocation	Enter Group Allocation Pattern, solve for group sample sizes
RL (Lower Non-Inferiority Limit)	0.8
Control Mean	10
Control Sample Size Allocation	1.732
Set A Number of Groups	1
Set A Mean	9.5
Set A Sample Size Allocation	1
Set B Number of Groups	1
Set B Mean	10
Set B Sample Size Allocation	1
Set C Number of Groups	1
Set C Mean	10.5
Set C Sample Size Allocation	1
Set D Number of Groups	0
More	Unchecked
σ (Standard Deviation)	2 2.5 3

Multi-Arm Non-Inferiority Tests for the Ratio of Treatment and Control Means (Log-Normal Data)

Output

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

Numeric Reports

Numeric Results

Solve For: [Sample Size](#)
 Group Allocation: Enter Group Allocation Pattern, solve for group sample sizes
 Test Type: Two-Sample T-Test of Means of Log-Transformed Data
 Higher Means Are: Better
 Hypotheses: $H_0: R \leq RL$ vs. $H_1: R > RL$
 Number of Groups: 4
 Bonferroni Adjustment: Standard Bonferroni (Divisor = 3)

Comparison	Power		Sample Size		Mean μ_i	Mean Ratio R_i	Lower Non-Inferiority Limit RL	Standard Deviation σ	Coefficient of Variation $COVi$	Alpha	
	Target	Actual	Ni	Allocation						Overall	Bonferroni-Adjusted
Control			40	1.732	10.0			2.0	0.20000		
vs A	0.8	0.80059	23	1.000	9.5	0.95	0.8	2.0	0.21053	0.025	0.00833
vs B	0.8	0.96481	23	1.000	10.0	1.00	0.8	2.0	0.20000	0.025	0.00833
vs C	0.8	0.99681	23	1.000	10.5	1.05	0.8	2.0	0.19048	0.025	0.00833
Total			109								
Control			61	1.732	10.0			2.5	0.25000		
vs A	0.8	0.80154	35	1.000	9.5	0.95	0.8	2.5	0.26316	0.025	0.00833
vs B	0.8	0.96518	35	1.000	10.0	1.00	0.8	2.5	0.25000	0.025	0.00833
vs C	0.8	0.99687	35	1.000	10.5	1.05	0.8	2.5	0.23810	0.025	0.00833
Total			166								
Control			87	1.732	10.0			3.0	0.30000		
vs A	0.8	0.80773	50	1.000	9.5	0.95	0.8	3.0	0.31579	0.025	0.00833
vs B	0.8	0.96737	50	1.000	10.0	1.00	0.8	3.0	0.30000	0.025	0.00833
vs C	0.8	0.99720	50	1.000	10.5	1.05	0.8	3.0	0.28571	0.025	0.00833
Total			237								

Comparison	The group that is involved in the comparison between the treatment and control displayed on this report line. The comparison is made using the ratio.
Target Power	The power desired. Power is probability of rejecting a false null hypothesis for this comparison. This power is of the comparison shown on this line only.
Actual Power	The power actually achieved.
Ni	The number of subjects in the ith group. The total sample size shown below the groups is equal to the sum of all individual group sample sizes.
Allocation	The group sample size allocation ratio of the ith group. The value on each row represents the relative number of subjects assigned to the group.
μ_i	The mean of the ith group at which the power is computed. The first row contains μ_c , the control group mean.
R_i	The ratio of the means at which the power is calculated. $R_i = \mu_i / \mu_c$.
RL	The lower non-inferiority boundary for the mean ratio. $RL < 1$.
σ	The standard deviation of the responses within each group.
$COVi$	The coefficient of variation of the ith group. $COVi = \sigma / \mu_i$.
Overall Alpha	The probability of rejecting at least one of the comparisons in this experiment when each null hypothesis is true.
Bonferroni Alpha	The adjusted significance level at which each individual comparison is made.

Summary Statements

A parallel, 4-group design (with one control group and 3 treatment groups) will be used to test whether the mean for each treatment group is non-inferior to the control group mean, with a non-inferiority ratio limit of 0.8 ($H_0: R \leq 0.8$ versus $H_1: R > 0.8$, $R = \mu_i / \mu_c$). In this study, higher means are considered to be better. The non-inferiority hypotheses will be evaluated using 3 one-sided, two-sample, Bonferroni-adjusted, equal-variance t-tests calculated on the log-transformed data, with an overall (experiment-wise) Type I error rate (α) of 0.025. The common standard deviation for all groups is assumed to be 2. The control group mean is assumed to be 10. To detect the treatment means 9.5, 10, and 10.5 with at least 80% power for each test, the control group sample size needed will be 40 and the number of needed subjects for the treatment groups will be 23, 23, and 23 (totaling 109 subjects overall).

Multi-Arm Non-Inferiority Tests for the Ratio of Treatment and Control Means (Log-Normal Data)

Dropout-Inflated Sample Size

Group	Dropout Rate	Sample Size Ni	Dropout- Inflated Enrollment Sample Size Ni'	Expected Number of Dropouts Di
1	20%	40	50	10
2	20%	23	29	6
3	20%	23	29	6
4	20%	23	29	6
Total		109	137	28
1	20%	61	77	16
2	20%	35	44	9
3	20%	35	44	9
4	20%	35	44	9
Total		166	209	43
1	20%	87	109	22
2	20%	50	63	13
3	20%	50	63	13
4	20%	50	63	13
Total		237	298	61

Group	Lists the group numbers.
Dropout Rate	The percentage of subjects (or items) that are expected to be lost at random during the course of the study and for whom no response data will be collected (i.e., will be treated as "missing"). Abbreviated as DR.
Ni	The evaluable sample size for each group at which power is computed (as entered by the user). If Ni subjects are evaluated out of the Ni' subjects that are enrolled in the study, the design will achieve the stated power.
Ni'	The number of subjects that should be enrolled in each group in order to obtain Ni evaluable subjects, based on the assumed dropout rate. Ni' is calculated by inflating Ni using the formula $Ni' = Ni / (1 - DR)$, with Ni' always rounded up. (See Julious, S.A. (2010) pages 52-53, or Chow, S.C., Shao, J., Wang, H., and Lokhnygina, Y. (2018) pages 32-33.)
Di	The expected number of dropouts in each group. $Di = Ni' - Ni$.

Dropout Summary Statements

Anticipating a 20% dropout rate, group sizes of 50, 29, 29, and 29 subjects should be enrolled to obtain final group sample sizes of 40, 23, 23, and 23 subjects.


References

- Blackwelder, W.C. 1998. 'Equivalence Trials.' In Encyclopedia of Biostatistics, John Wiley and Sons. New York. Volume 2, 1367-1372.
- Chow, S.C., Shao, J., Wang, H., and Lokhnygina, Y. 2018. Sample Size Calculations in Clinical Research, 3rd Edition. Chapman & Hall/CRC. Boca Raton, FL. Pages 86-88.
- Julious, Steven A. 2004. 'Tutorial in Biostatistics. Sample sizes for clinical trials with Normal data.' Statistics in Medicine, 23:1921-1986.
- Machin, D., Campbell, M.J., Tan, S.B, and Tan, S.H. 2018. Sample Sizes for Clinical, Laboratory, and Epidemiology Studies, 4th Edition. Wiley Blackwell.

This report shows the numeric results of this power study. Notice that the results are shown in blocks of three rows at a time. Each block represents a single design.

Plots Section

Plots



This plot gives a visual presentation to the results in the Numeric Report. We can quickly see the impact on the sample size of changing the standard deviation.

Example 2 – Validation using a Previously Validated Procedure

We could not find a validation result in the statistical literature, so we will use the previously validated **PASS** procedure **Non-Inferiority Tests for the Ratio of Two Means (Log-Normal Data)** to produce the results for the following example.

A parallel-group clinical trial is being designed to compare three treatment therapies against the standard therapy. Higher values of the response are desirable. Suppose the standard therapy has a mean response of 10 with a standard deviation of 2.5. The investigators would like a sample size large enough to find statistical significance at the 0.025 level (since this a one-sided test) if the actual mean responses of the three treatments are 9.5, 10.0, and 10.5, the power of each test is 0.80, and the non-inferiority limit is 0.8 (NIM = 0.2). The standard deviation of 2.5 leads to a coefficient of variation in the control group of 0.25.

The sample sizes of all groups will be equal.

The **Non-Inferiority Tests for the Ratio of Two Means (Log-Normal Data)** procedure is set up as follows.

Design Tab

Solve For **Sample Size**
 Higher Means Are **Better (H1: $R > 1 - \text{NIM}$)**
 Power **0.8**
 Alpha **0.00833333** (which is Alpha / k)
 Group Allocation **Equal (N1 = N2)**
 NIM (Non-Inferiority Margin) **0.2**
 R1 (Actual Ratio) **0.95 1 1.05**
 COV (Coefficient of Variation) **0.25**

This set of options generates the following report.

Numeric Results

Solve For: [Sample Size](#)
 Groups: 1 = Treatment, 2 = Reference
 Ratio: R = Treatment Mean / Reference Mean
 Higher Means Are: Better
 Hypotheses: H0: $R \leq 1 - \text{NIM}$ vs. H1: $R > 1 - \text{NIM}$

Power		Sample Size			Non-Inferiority		Actual Ratio R1	Coefficient of Variation COV	Alpha
Target	Actual	N1	N2	N	Margin -NIM	Bound R0			
0.8	0.80582	45	45	90	-0.2	0.8	0.95	0.25	0.00833
0.8	0.80054	27	27	54	-0.2	0.8	1.00	0.25	0.00833
0.8	0.80876	19	19	38	-0.2	0.8	1.05	0.25	0.00833

In order to maintain a power of 80% for all three groups, it is apparent that the groups will all need to have a sample size of 45. This table contains the validation values. We will now run these values through the current procedure and compare the results with these values.

Multi-Arm Non-Inferiority Tests for the Ratio of Treatment and Control Means (Log-Normal Data)

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 2** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Solve For **Sample Size**
 Higher Means Are **Better (H1: R > RL)**
 Power of Each Test **0.8**
 Overall Alpha **0.025**
 Bonferroni Adjustment **Standard Bonferroni**
 Group Allocation **Equal (Nc = N1 = N2 = ...)**
 RL (Lower Non-Inferiority Limit) **0.8**
 Control Mean **10**
 Set A Number of Groups..... **1**
 Set A Mean **9.5**
 Set B Number of Groups..... **1**
 Set B Mean **10**
 Set C Number of Groups **1**
 Set C Mean..... **10.5**
 Set D Number of Groups **0**
 More..... **Unchecked**
 σ (Standard Deviation)..... **2.5**

Output

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

Numeric Results

Solve For: [Sample Size](#)
 Group Allocation: Equal (Nc = N1 = N2 = ...)
 Test Type: Two-Sample T-Test of Means of Log-Transformed Data
 Higher Means Are: Better
 Hypotheses: H0: R ≤ RL vs. H1: R > RL
 Number of Groups: 4
 Bonferroni Adjustment: Standard Bonferroni (Divisor = 3)

Comparison	Power		Sample Size Ni	Mean μ_i	Mean Ratio Ri	Lower Non-Inferiority Limit RL	Standard Deviation σ	Coefficient of Variation COVi	Alpha	
	Target	Actual							Overall	Bonferroni-Adjusted
Control			45	10.0			2.5	0.25000		
vs A	0.8	0.80582	45	9.5	0.95	0.8	2.5	0.26316	0.025	0.00833
vs B	0.8	0.96669	45	10.0	1.00	0.8	2.5	0.25000	0.025	0.00833
vs C	0.8	0.99710	45	10.5	1.05	0.8	2.5	0.23810	0.025	0.00833
Total			180							

As you can see, the sample sizes are all 45, which match the largest sample size found in the validation run above. The procedure is validated.