Chapter 327

Multi-Arm Tests for the Difference Between Treatment and Control Means Assuming Equal Variance

Introduction

This module computes power and sample size for multiple comparisons of treatment means versus a control mean based on the results in Machin, Campbell, Tan, and Tan (2018). In this 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 of the difference between two means.

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.

The Dunnett comparison-with-a-control procedure is a direct competitor to this procedure. If you can meet all of the assumptions of Dunnett's procedure, it will usually be a more powerful test and require a smaller sample size. We have added this procedure for completeness and because it provides an alternative to Dunnett's test that may be useful in some situations.

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. For example, 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_i. The total sample size is $N = N_1 + N_2 + \cdots + N_k + N_C$.

The statistical hypotheses for two-sided tests are

$$H_{0i}: \mu_i = \mu_C \text{ vs. } H_{1i}: \mu_i \neq \mu_C$$

and for one-sided tests the hypotheses are

$$H_{0i}: \mu_i \leq \mu_C \text{ vs. } H_{1i}: \mu_i > \mu_C$$

or

$$H_{0i}: \mu_i \geq \mu_C \text{ vs. } H_{1i}: \mu_i < \mu_C$$

If we define $\delta_i = \mu_i - \mu_C$, these are equivalent to

$$H_{0i}: \delta_i = 0 \text{ vs. } H_{1i}: \delta_i \neq 0 \quad \text{for } i = 1, 2, \ldots, k$$

$$H_{0i}: \delta_i \leq 0 \text{ vs. } H_{1i}: \delta_i > 0 \quad \text{for } i = 1, 2, \ldots, k$$

$$H_{0i}: \delta_i \geq 0 \text{ vs. } H_{1i}: \delta_i < 0 \quad \text{for } i = 1, 2, \ldots, k$$

For convenience, these hypotheses are collectively referred to as

$$H_0: \delta = 0 \text{ vs. } H_1: \delta \neq 0$$

$$H_0: \delta \leq 0 \text{ vs. } H_1: \delta > 0$$

$$H_0: \delta \geq 0 \text{ vs. } H_1: \delta < 0$$

Test Statistic

A suitable Type I error probability is chosen for the test, the data are collected, and a t-statistic is generated using the formula

$$t = \frac{\bar{x}_i - \bar{x}_C}{\sqrt{\frac{(N_i - 1)s_i^2 + (N_C - 1)s_C^2}{N_i + N_C - 2}\left(\frac{1}{N_i} + \frac{1}{N_C}\right)}}$$

This *t*-statistic follows a *t* distribution with $N_i + N_C - 2$ degrees of freedom.

Power Calculation

The power of this test is computed using the noncentral t distribution with $N_i + N_C - 2$ degrees of freedom and non-centrality parameter

$$\lambda = \frac{\mu_i - \mu_C}{\sigma\sqrt{\frac{1}{N_i} + \frac{1}{N_C}}}$$

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 the 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. Suppose the standard therapy has a mean response of 9.3 with a standard deviation of 2.5. The investigators would like a sample size large enough to find statistical significance at the 0.05 level if the actual mean responses of the three treatments are 7.3, 7.6, and 8.1 and the power is 0.80 in each test. They want to consider a range of standard deviations from 2.0 to 3.0. The tests will be two-sided.

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**
Alternative Hypothesis	**Two-Sided (H1: δ ≠ 0)**
Power of Each Test	**0.8**
Overall Alpha	**0.05**
Bonferroni Adjustment	**Standard Bonferroni**
Group Allocation	**Enter Group Allocation Pattern, solve for group sample sizes**
Control Mean	**9.3**
Control Sample Size Allocation	**1.732**
Set A Number of Groups	**1**
Set A Mean	**7.3**
Set A Sample Size Allocation	**1**
Set B Number of Groups	**1**
Set B Mean	**7.6**
Set B Sample Size Allocation	**1**
Set C Number of Groups	**1**
Set C Mean	**8.1**
Set C Sample Size Allocation	**1**
Set D Number of Groups	**0**
More	**Unchecked**
σ (Standard Deviation)	**2 2.5 3.0**

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:	T-Test
Hypotheses:	H0: δ = 0 vs. H1: δ ≠ 0
Number of Groups:	4
Bonferroni Adjustment:	Standard Bonferroni (Divisor = 3)

	Power		Sample Size					Alpha	
Comparison	Target	Actual	Ni	Allocation	Mean μi	Difference δi	Standard Deviation σ	Overall	Bonferroni-Adjusted
Control			83	1.732	9.3		2.0		
vs A	0.8	0.99889	48	1.000	7.3	-2.0	2.0	0.05	0.01667
vs B	0.8	0.98749	48	1.000	7.6	-1.7	2.0	0.05	0.01667
vs C	0.8	0.81003	48	1.000	8.1	-1.2	2.0	0.05	0.01667
Total			227						
Control			126	1.732	9.3		2.5		
vs A	0.8	0.99867	73	1.000	7.3	-2.0	2.5	0.05	0.01667
vs B	0.8	0.98593	73	1.000	7.6	-1.7	2.5	0.05	0.01667
vs C	0.8	0.80111	73	1.000	8.1	-1.2	2.5	0.05	0.01667
Total			345						
Control			182	1.732	9.3		3.0		
vs A	0.8	0.99873	105	1.000	7.3	-2.0	3.0	0.05	0.01667
vs B	0.8	0.98633	105	1.000	7.6	-1.7	3.0	0.05	0.01667
vs C	0.8	0.80333	105	1.000	8.1	-1.2	3.0	0.05	0.01667
Total			497						

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 difference.
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.
μi	The mean of the ith group at which the power is computed. The first row contains μc, the control group mean.
δi	The difference between the ith treatment mean and the control mean (μi - μc) at which the power is computed.
σ	The standard deviation of the responses within each group.
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 different from the control group mean (H0: δ = 0 versus H1: δ ≠ 0, δ = μi - μc). The hypotheses will be evaluated using 3 two-sided, two-sample, Bonferroni-adjusted, equal-variance t-tests, with an overall (experiment-wise) Type I error rate (α) of 0.05. The common standard deviation for all groups is assumed to be 2. The control group mean is assumed to be 9.3. To detect the treatment means 7.3, 7.6, and 8.1 with at least 80% power for each test, the control group sample size needed will be 83 and the number of needed subjects for the treatment groups will be 48, 48, and 48 (totaling 227 subjects overall).

Dropout-Inflated Sample Size

Group	Dropout Rate	Sample Size Ni	Dropout-Inflated Enrollment Sample Size Ni'	Expected Number of Dropouts Di
1	20%	83	104	21
2	20%	48	60	12
3	20%	48	60	12
4	20%	48	60	12
Total		227	284	57
1	20%	126	158	32
2	20%	73	92	19
3	20%	73	92	19
4	20%	73	92	19
Total		345	434	89
1	20%	182	228	46
2	20%	105	132	27
3	20%	105	132	27
4	20%	105	132	27
Total		497	624	127

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 104, 60, 60, and 60 subjects should be enrolled to obtain final group sample sizes of 83, 48, 48, and 48 subjects.

References

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.

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.

Julious, Steven A. 2004. 'Tutorial in Biostatistics. Sample sizes for clinical trials with Normal data.' Statistics in Medicine, 23:1921-1986.

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

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

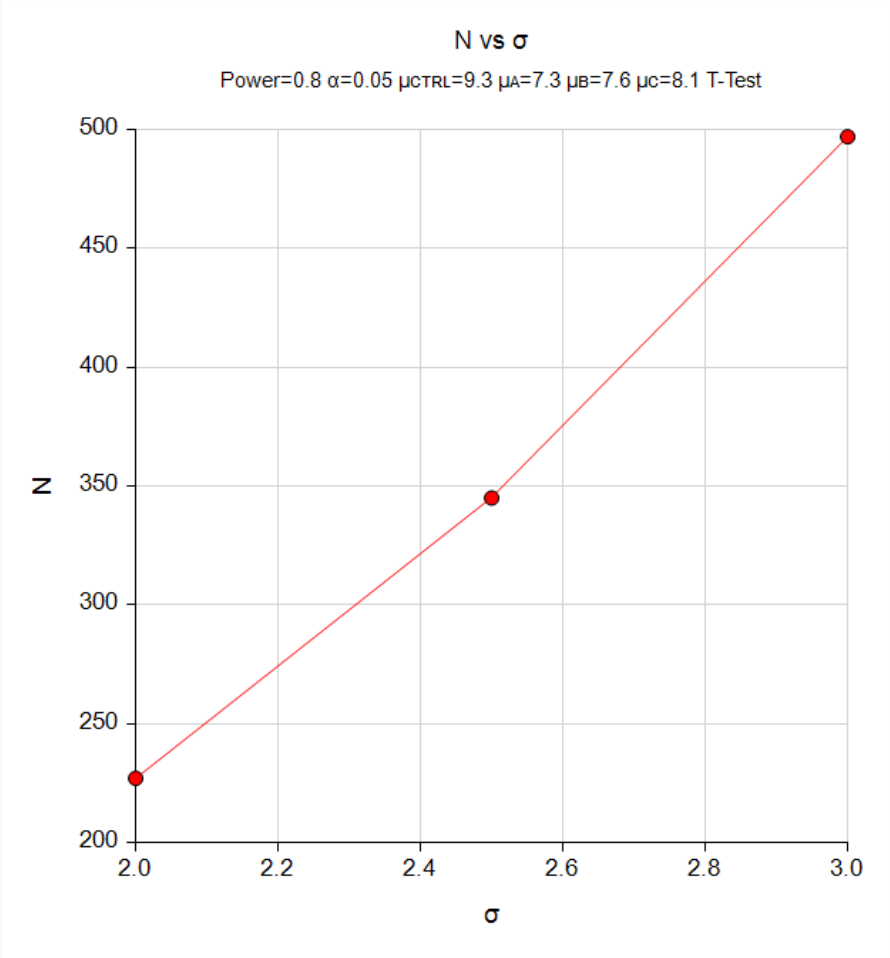

This plot gives a visual presentation to the results in the Numeric Report. We can quickly see the impact on the sample size of changing the standard deviation.

Example 2 – Validation using a Previously Validated Procedure

We could not find a validation result in the statistical literature, so we will use a previously validated **PASS** procedure (**Two-Sample T-Tests Assuming Equal Variance**) 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. Suppose the standard therapy has a mean response of 9.3 with a standard deviation of 2.5. The investigators would like a sample size large enough to find statistical significance at the 0.05 level if the actual mean responses of the three treatments are 7.3, 7.6, and 8.1 and the power is 0.80 in each test. The tests will be two-sided.

The sample sizes of all groups will be equal.

The **Two-Sample T-Tests Assuming Equal Variance** procedure is set up as follows.

Design Tab

Solve For	**Sample Size**
Alternative Hypothesis	**Two-Sided**
Power	**0.8**
Alpha	**0.016667** (which is Alpha / k)
Group Allocation	**Equal (N1 = N2)**
Input Type	**Means**
μ1	**7.3 7.6 8.1**
μ2	**9.3**
σ	**2.5**

This set of options generates the following report.

Numeric Results

Solve For: Sample Size
Test Type: Two-Sample Equal-Variance T-Test
Difference: δ = μ1 - μ2
Hypotheses: H0: δ = 0 vs. H1: δ ≠ 0

Power		Sample Size			Mean				
							Difference	Standard Deviation	
Target	Actual	N1	N2	N	μ1	μ2	δ	σ	Alpha
0.8	0.81105	35	35	70	7.3	9.3	-2.0	2.5	0.01667
0.8	0.80259	47	47	94	7.6	9.3	-1.7	2.5	0.01667
0.8	0.80335	93	93	186	8.1	9.3	-1.2	2.5	0.01667

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 93. This table contains the validation values. We will now run these values through the current procedure and compare the results with these values.

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**
Alternative Hypothesis	**Two-Sided (H1: δ ≠ 0)**
Power of Each Test	**0.8**
Overall Alpha	**0.05**
Bonferroni Adjustment	**Standard Bonferroni**
Group Allocation	**Equal (Nc = N1 = N2 = ...)**
Control Mean	**9.3**
Set A Number of Groups	**1**
Set A Mean	**7.3**
Set B Number of Groups	**1**
Set B Mean	**7.6**
Set C Number of Groups	**1**
Set C Mean	**8.1**
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:	T-Test
Hypotheses:	H0: δ = 0 vs. H1: δ ≠ 0
Number of Groups:	4
Bonferroni Adjustment:	Standard Bonferroni (Divisor = 3)

	Power		Sample Size				Alpha	
Comparison	Target	Actual	Ni	Mean μi	Difference δi	Standard Deviation σ	Overall	Bonferroni-Adjusted
Control			93	9.3		2.5		
vs A	0.8	0.99873	93	7.3	-2.0	2.5	0.05	0.01667
vs B	0.8	0.98633	93	7.6	-1.7	2.5	0.05	0.01667
vs C	0.8	0.80335	93	8.1	-1.2	2.5	0.05	0.01667
Total			372					

As you can see, the sample sizes are all 93, which match the largest sample size found in the validation run above. The procedure is validated.