Chapter 549

One-Way Analysis of Variance Contrasts

Introduction

The one-way (multiple group) design allows the means of two or more populations (groups) to be compared to determine if at least one mean is different from the others. The *F* test is used to determine statistical significance.

The usual *F* test tests the hypothesis that all means are equal versus the alternative that at least one mean is different from the rest. Often, a more specific alternative is desired. For example, you might want to test whether the treatment means are different from the control mean, the low dose is different from the high dose, a linear trend exists across dose levels, and so on. These questions are tested using specific contrasts.

A *comparison* is a weighted average of the means, in which the weights may be negative. When the weights sum to zero, the comparison is called a *contrast*. **PASS** provides results for contrasts. To specify a contrast, we need only specify the weights.

For example, suppose an experiment conducted to study a drug will have three dose levels: none (control), 20 mg, and 40 mg. The first question is whether the drug made a difference. If it did, the average response for the two groups receiving the drug should be different from the control. If we label the group means M0, M2, and M4, we are interested in comparing M0 with M2 and M4. This can be done in two ways. One way is to construct two tests, one comparing M0 and M2 and the other comparing M0 and M4. Another method is to perform one test comparing M0 with the average of M2 and M4. These tests are conducted using contrasts. The coefficients are as follows:

M0 vs. M2

To compare M0 versus M2, use the coefficients -1, 1, 0. When applied to the group means, these coefficients result in the comparison M0(-1) + M2(1) + M4(0) which reduces to M2-M0. That is, this contrast results in the difference between two group means. We can test whether this difference is non-zero using the *t* test (or *F* test since the square of the *t* test is an *F* test).

M0 vs. M4

To compare M0 versus M4, use the coefficients -1, 0, 1. When applied to the group means, these coefficients result in the comparison M0(-1) + M2(0) + M4(1) which reduces to M4 - M0. That is, this contrast results in the difference between the two group means.

M0 vs. Average of M2 and M4

To compare M0 versus the average of M2 and M4, use the coefficients -2, 1, 1. When applied to the group means, these coefficients result in the comparison M0(-2) + M2(1) + M4(1) which is equivalent to M4 + M2 - 2(M0).

Assumptions

This test requires certain assumptions.

1. The data are continuous (not discrete).
2. The data follow the normal probability distribution. Each group is normally distributed about the group mean.
3. The variances of the populations are equal.
4. The groups are independent. There is no relationship among the individuals in one group as compared to another.
5. Each group is a simple random sample from its population. Each individual in the population has an equal probability of being selected in the sample.

Technical Details for a Planned Comparison

Suppose *G* groups each have a normal distribution and equal means ($\mu_1 = \mu_2 = \ldots = \mu_G$). Let $n_1, n_2, \ldots, n_G$ denote the number of subjects in each group and let *N* denote the total sample size of all groups. Let μ denote the weighted mean of all groups. That is

$$\mu = \sum_{i=1}^{G} \frac{n_i}{N} \mu_i$$

Let σ denote the common standard deviation of all groups.

Suppose you want to test whether the contrast *C*

$$C = \sum_{i=1}^{G} c_i \mu_i$$

is significantly different from zero. Here the c_i's are the contrast coefficients.

Define

$$\sigma_C = \left| \sum_{i=1}^{G} c_i \mu_i \right| / \sqrt{N \sum_{i=1}^{G} \frac{c_i^2}{n_i}}$$

Define the noncentrality parameter λ_C, as

$$\lambda_C = N \sigma_C^2 / \sigma^2$$

Power Calculations for Contrasts

The calculation of the power of a particular test proceeds as follows:

1. Determine the critical value, $F_{1,N-G,\alpha}$, where α is the probability of a type-I error and *G* and *N* are defined above. Note that this is a two-tailed test as no direction is assigned in the alternative hypothesis.
2. From a hypothesized set of μ_i's, calculate the noncentrality parameter λ_C.
3. Compute the power as the probability of being greater than $F_{1,N-G,\alpha}$ on a noncentral-*F* distribution with noncentrality parameter λ_C.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Design tab. For more information about the options of other 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.

Solve For

Solve For

This option specifies the parameter to be solved for from the other parameters. The parameters that may be selected are *Sample Size*, *Alpha*, and *Power*. Under most situations, you will select either *Power* for a power analysis or *Sample Size* for sample size determination.

Power and Alpha

Power

This option specifies one or more values for power. 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

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 – Groups

G (Number of Groups)

This is the number of group means being compared. It must be greater than or equal to two.

Note that the number of items used in the Means box and the Group Sample Size Pattern box is controlled by, and must match, this number.

Group Allocation Ratios

Enter a list of positive, numeric values, one for each group. The sample size of group *i* is found by multiplying the i^{th} number from this list times the value of n and rounding up to the next whole number. The number of values must match the number of groups, *G*. When too few numbers are entered, 1's are added. When too many numbers are entered, the extras are ignored.

Equal

If all sample sizes are to be equal, enter "Equal" here and the desired sample size in n. A set of *G* 1's will be used. This will result in $n1 = n2 = n3 = n$. That is, all sample sizes are equal to *n*.

n (Subjects per Group)

This is the base, per group, sample size. One or more values, separated by blanks or commas, may be entered. A separate analysis is performed for each value listed here.

The group samples sizes are determined by multiplying this number by each of the Group Sample Size Pattern numbers. If the Group Sample Size Pattern numbers are represented by *m1, m2, m3, …, mG* and this value is represented by *n*, the group sample sizes *n1, n2, n3, ..., nG* are calculated as follows:

n1=[n(m1)]

n2=[n(m2)]

n3=[n(m3)]

etc.

where the operator, [*X*] means the next integer after *X*, e.g. [3.1]=4.

For example, suppose there are three groups and the Group Sample Size Pattern is set to *1,2,3*. If n is 5, the resulting sample sizes will be 5, 10, and 15. If n is 50, the resulting group sample sizes will be 50, 100, and 150. If n is set to *2,4,6,8,10*, five sets of group sample sizes will be generated and an analysis run for each. These sets are:

2	4	6
4	8	12
6	12	18
8	16	24
10	20	30

Fractional Allocation Ratios

As a second example, suppose there are three groups and the Group Sample Size Pattern is *0.2,0.3,0.5*. When the fractional Pattern values sum to one, *n* can be interpreted as the total sample size of all groups (N) and the allocation ratios as the proportion of the total in each group.

If n is 10, the three group sample sizes would be 2, 3, and 5.

If n is 20, the three group sample sizes would be 4, 6, and 10.

If n is 12, the three group sample sizes would be

(0.2)12 = 2.4 which is rounded up to the next whole integer, 3.

(0.3)12 = 3.6 which is rounded up to the next whole integer, 4.

(0.5)12 = 6.

Note that in this case, 3+4+6 does not equal n (which is 12). This can happen because of rounding.

Effect Size

Means (μ1, μ2, ..., μG)

Enter G group means assumed by the alternative hypothesis. These are combined with the contrast coefficients to form σc which is used in the power analysis. The formula for this is

$$\sigma c = |\Sigma(c_i \mu_i)| / \sqrt{[N\, \Sigma(c_i^2/n_i)]}$$

If not enough means are entered, the last mean is repeated. If too many values are entered, PASS will ignore the extra values.

Examples:

5 20 60

2,5,7

-4,0,6,9

K (Means Multiplier)

Enter one or more values for K, the means multiplier. A separate power calculation is conducted for each value of K. In each analysis, all means (μi's) are multiplied by K. In this way, you can determine how sensitive the power values are to the magnitude of the means without the need to change them individually.

For example, if the original means are '0 1 2', setting this option to '1 2' results in two sets of means used in separate analyses: '0 1 2' in the first analysis and '0 2 4' in the second analysis.

If you want to ignore this option, enter '1'.

Examples

1

0.5 1 1.5

0.8 to 1.2 by 0.1

Contrast Coefficients

Enter a set of contrast coefficients which are for testing whether this contrast of the means is zero versus the alternative that it is non-zero (two-sided test). These are often called *Planned Comparisons*.

A contrast is a weighted average of the means in which the weights sum to zero. For example, suppose you are studying four groups and that the main hypothesis of interest is whether there is a linear trend across the groups. You would enter *-3, -1, 1, 3* here. This would form the weighted average of the means:

$-3\mu 1 - \mu 2 + \mu 3 + 3\mu 4$

These coefficients are used to form a specific weighted average of the means, $\Sigma ci\mu i$, which is compared against zero using a standard F test. The power analysis uses σc which measures the variation in the contrast. The formula for σc is

$\sigma c = |\Sigma(ci\mu i)| / \sqrt{[N\ \Sigma(ci^2/ni)]}$

Options

- **Linear Trend**
  A set of coefficients is generated appropriate for testing the alternative hypothesis that there is a linear (straight-line) trend across the means. These coefficients assume that the means are equally spaced across the trend variable.

- **Quadratic**
  A set of coefficients is generated appropriate for testing the alternative hypothesis that the means follow a quadratic model. These coefficients assume that the means are equally spaced across the implicit *X* variable.

- **Cubic**
  A set of coefficients is generated appropriate for testing the alternative hypothesis that the means follow a cubic model. These coefficients assume that the means are equally spaced across the implicit *X* variable.

- **First Against Others**
  A set of coefficients is generated appropriate for testing the alternative hypothesis that the first mean is different from the average of the remaining means. For example, if there were four groups, the generated coefficients would be *-3, 1, 1, 1*.

- **List of Coefficients**

  A list of coefficients, separated by commas or blanks, may be entered. If the number of items in the list does not match the number of groups (*G*), zeros are added or extra coefficients are truncated.

  Remember that these coefficients must sum to zero. Also, the scale of the coefficients does not matter. That is *0.5,0.25,0.25*; *-2,1,1*; and *-200,100,100* will yield the same results.

  To avoid rounding problems, it is better to use *-3,1,1,1* than the equivalent *-1,0.333,0.333,0.333*. The second set does not sum to zero.

Effect Size – Standard Deviation of the Data

σ (Standard Deviation)

This is standard deviation between subjects within a group. It represents the variability from subject to subject that occurs when the subjects are treated identically. It is assumed to be the same for all groups. This value is approximated in an analysis of variance table by the square root of the mean square error.

Since they are positive square roots, the numbers must be strictly greater than zero. You can press the σ button to obtain further help on estimating the standard deviation.

Note that if you are using this procedure to test a factor (such as an interaction) from a more complex design, the value of standard deviation is estimated by the square root of the mean square of the term that is used as the denominator in the *F* test.

You can enter a list of values separated by blanks or commas, in which case, a separate analysis will be calculated for each value.

Examples of valid entries:

1,4,7,10

1 4 7 10

1 to 10 by 3

Example 1 – Finding Power

An experiment is being designed to compare the means of four groups using contrast test with a significance level of 0.05. The first group is a control group. The other three groups will be have slightly different treatments applied. The researchers are mainly interested in whether the three treatment groups are different from the control group. Hence, they want to test the contrast represented by the coefficients -3, 1, 1, 1. Treatment means of 40, 10, 10, and 10 represent clinically important group differences.

Previous studies have had standard deviations between 18 and 24. To better understand the relationship between power and sample size, the researcher wants to compute the power for several group sample sizes between 2 and 14. The sample sizes will be equal across all groups.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **One-Way Analysis of Variance Contrasts** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance Contrasts**. You may then make the appropriate entries as listed below, or open **Example 1** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Power**
Alpha	**0.05**
G (Number of Groups)	**4**
Group Allocation Ratios	**Equal**
n (Subjects per Group)	**2 to 14 by 2**
Means	**40 10 10 10**
K (Means Multiplier)	**1**
Contrast Coefficients	**-3 1 1 1**
σ (Standard Deviation)	**18 21 24**
Reports Tab	
Show Details	***Checked***

Annotated Output

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

Numeric Results

Numeric Results
Means: 40 10 10 10
Contrast: -3 1 1 1

Power	Average n	G	Total N	K	Contrast Std Dev σc	Data Std Dev σ	Effect Size	Alpha
0.3471	2.00	4	8	1.00	12.99	18.00	0.7217	0.0500
0.2713	2.00	4	8	1.00	12.99	21.00	0.6186	0.0500
0.2201	2.00	4	8	1.00	12.99	24.00	0.5413	0.0500
0.7550	4.00	4	16	1.00	12.99	18.00	0.7217	0.0500
0.6231	4.00	4	16	1.00	12.99	21.00	0.6186	0.0500
0.5123	4.00	4	16	1.00	12.99	24.00	0.5413	0.0500
0.9194	6.00	4	24	1.00	12.99	18.00	0.7217	0.0500
0.8218	6.00	4	24	1.00	12.99	21.00	0.6186	0.0500
0.7132	6.00	4	24	1.00	12.99	24.00	0.5413	0.0500
0.9761	8.00	4	32	1.00	12.99	18.00	0.7217	0.0500
0.9218	8.00	4	32	1.00	12.99	21.00	0.6186	0.0500
0.8402	8.00	4	32	1.00	12.99	24.00	0.5413	0.0500
0.9934	10.00	4	40	1.00	12.99	18.00	0.7217	0.0500
0.9676	10.00	4	40	1.00	12.99	21.00	0.6186	0.0500
0.9148	10.00	4	40	1.00	12.99	24.00	0.5413	0.0500
0.9983	12.00	4	48	1.00	12.99	18.00	0.7217	0.0500
0.9871	12.00	4	48	1.00	12.99	21.00	0.6186	0.0500
0.9561	12.00	4	48	1.00	12.99	24.00	0.5413	0.0500
0.9996	14.00	4	56	1.00	12.99	18.00	0.7217	0.0500
0.9951	14.00	4	56	1.00	12.99	21.00	0.6186	0.0500
0.9780	14.00	4	56	1.00	12.99	24.00	0.5413	0.0500

References
Desu, M. M. and Raghavarao, D. 1990. Sample Size Methodology. Academic Press. New York.
Fleiss, Joseph L. 1986. The Design and Analysis of Clinical Experiments. John Wiley & Sons. New York.
Kirk, Roger E. 1982. Experimental Design: Procedures for the Behavioral Sciences. Brooks/Cole. Pacific Grove, California.

Report Definitions
Power is the probability of rejecting a false null hypothesis.
n is the average group sample size.
G is the number of groups.
Total N is the total sample size of all groups combined.
K is the group means multiplier.
σc is the standard deviation of the contrast of the group means.
σ is the subject-to-subject standard deviation.
The Effect Size is the ratio of σc and σ.
Alpha is the probability of rejecting a true null hypothesis.

Summary Statements
In a one-way ANOVA study, sample sizes of 2, 2, 2, and 2 are obtained from the 4 groups whose means are to be compared using a planned comparison (contrast). The total sample of 8 subjects achieves 35% power to detect a non-zero contrast of the means versus the alternative that the contrast is zero using an F test with a 0.0500 significance level. The value of the contrast of the means is -90.00. The common standard deviation of the data is assumed to be 18.00.

This report shows the numeric results of this power study.

Detailed Results Report

Details when Alpha = 0.0500, Power = 0.3471, σc = 12.99, σ = 18.00

Group	Ni	Percent Ni of Total N	Mean	Contrast Coefficient	Mean Times Coefficient
1	2.00	25.00	40.00	-3.000	-120.00
2	2.00	25.00	10.00	1.000	10.00
3	2.00	25.00	10.00	1.000	10.00
4	2.00	25.00	10.00	1.000	10.00
ALL	8	100.00	17.50	0.00	-90.00

(report continues)

This report shows the details of each row of the previous report.

Plots Section

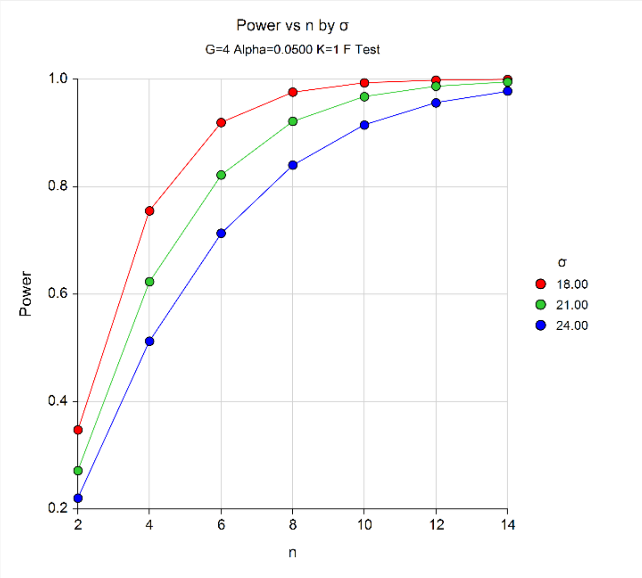

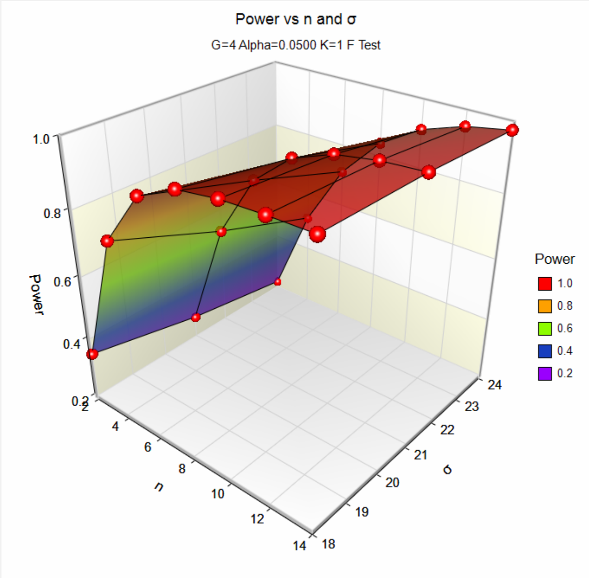

These plots give a visual presentation to the results in the Numeric Report. We can quickly see the impact on the power of increasing the sample size and the increase in σ.

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 – Validation using Hand Calculations

We will compute the following example by hand and then compare that with the results that **PASS** obtains. Here are the settings:

Alpha	0.05
G	3
Allocation	Equal
n	5
Means	1, 2, 3
K	1
Coefficients	-2, 1, 1
σ	5

Using these values, we find the following

C'μ	3
σ_c^2	9/18 = 0.5
λ_C	15 x 0.5/(25) = 0.3
$F_{0.95,1,12}$	4.747225

Power = 0.0797

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **One-Way Analysis of Variance Contrasts** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance Contrasts**. You may then make the appropriate entries as listed below, or open **Example 2** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Power**
Alpha	**0.05**
G (Number of Groups)	**3**
Group Allocation Ratios	**Equal**
n (Subjects per Group)	**5**
Means	**1 2 3**
K (Means Multiplier)	**1**
Contrast Coefficients	**-2 1 1**
σ (Standard Deviation)	**5**

Output

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

Numeric Results Report

Power	Average n	G	Total N	K	Contrast Std Dev σc	Data Std Dev σ	Effect Size	Alpha
0.0797	5.00	3	15	1.00	0.71	5.00	0.1414	0.0500

PASS has also calculated the power to be 0.0797.