

Chapter 551

Analysis of Covariance

Introduction

A common task in research is to compare the averages of two or more populations (groups). We might want to compare the income level of two regions, the nitrogen content of three lakes, or the effectiveness of four drugs. The one-way analysis of variance compares the means of two or more groups to determine if at least one mean is different from the others. The F test is used to determine statistical significance.

Analysis of Covariance (ANCOVA) is an extension of the one-way analysis of variance model that adds quantitative variables (covariates). When used, it is assumed that their inclusion will reduce the size of the error variance and thus increase the power of the design.

Covariates can only be used if the assumption of parallel slopes is viable.

Planned Comparisons

PASS performs power and sample size calculations for user-specified contrasts.

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 planned comparisons.

We call the comparison *planned* because it was determined before the experiment was conducted. We planned to test the comparison.

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. Statisticians call these weights the *contrast coefficients*.

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 M_0 , M_{20} , and M_{40} , we are interested in comparing M_0 with M_{20} and M_{40} . This can be done in two ways. One way is to construct two tests, one comparing M_0 and M_{20} and the other comparing M_0 and M_{40} . Another method is to perform one test comparing M_0 with the average of M_{20} and M_{40} . These tests are conducted using planned comparisons. The coefficients are as follows:

Analysis of Covariance

M0 vs. M20

To compare M0 versus M20, use the coefficients -1,1,0. When applied to the group means, these coefficients result in the comparison $M0(-1)+M20(1)+M40(0)$ which reduces to $M20-M0$. That is, this contrast results in the difference between the 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. M40

To compare M0 versus M40, use the coefficients -1,0,1. When applied to the group means, these coefficients result in the comparison $M0(-1)+M20(0)+M40(1)$ which reduces to $M40-M0$. That is, this contrast results in the difference between the two group means.

M0 vs. Average of M20 and M40

To compare M0 versus the average of M20 and M40, use the coefficients -2,1,1. When applied to the group means, these coefficients result in the comparison $M0(-2)+M20(1)+M40(1)$ which is equivalent to $M40+M20-2(M0)$.

To see how these coefficients were obtained, consider the following manipulations. Beginning with the difference between the average of M20 and M40 and M0, we obtain the coefficients above—aside from a scale factor of one-half.

$$\begin{aligned} \frac{M20 + M40}{2} - M0 &= \frac{M20}{2} + \frac{M40}{2} - \frac{M0}{1} \\ &= \frac{1}{2} M20 + \frac{1}{2} M40 - M0 \\ &= \frac{1}{2} (M20 + M40 - 2M0) \end{aligned}$$

Assumptions

Using the F test requires certain assumptions. One reason for the popularity of the F test is its robustness in the face of assumption violation. However, if an assumption is not even approximately met, the significance levels and the power of the F test are invalidated. Unfortunately, in practice it often happens that several assumptions are not met. This makes matters even worse. Hence, steps should be taken to check the assumptions before important decisions are made.

The following assumptions are needed for a one-way analysis of variance:

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.

Additional assumptions are needed for an analysis of covariance:

1. The covariates have a linear relationship the response variable.
2. The slopes of these linear relationships between the covariate and the response variable are approximately equal across all groups.

Technical Details for ANCOVA

We found two, slightly different, formulations for computing power for analysis of covariance. Keppel (1991) gives results that modify the standard deviation by an amount proportional to its reduction because of the covariate. Borm et al. (2007) give results that use a normal approximation to the noncentral F distribution. We use the Keppel approach in PASS.

Suppose that each observation consists of a response measurement, Y , and one or more covariate measurements: X_1, X_2, \dots, X_p . Further suppose that samples of n_1, n_2, \dots, n_k observations will be obtained from each of k groups. The multiple regression equation relating Y to the X 's within the i^{th} group is

$$Y_i = \beta_{0i} + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_p X_p + \varepsilon$$

The β 's are the regression coefficients or slopes. Analysis of covariance assumes that, except for the intercept β_0 , the slopes are equal across all groups. Thus, the difference between the means of any two groups is equal to the difference between their intercepts.

Let σ^2 denote the common variance of all groups ignoring the covariates and σ_ε^2 the within-group variance after considering the covariates. These values are related according to the formula

$$\sigma_\varepsilon^2 = (1 - \rho^2) \sigma^2$$

where ρ^2 is the coefficient of multiple determination (estimated by R^2).

Given the above terminology, the ratio of the mean square between groups to the mean square within groups follows a central F distribution with two parameters matching the degrees of freedom of the numerator mean square and the denominator mean square. When the null hypothesis of mean equality is rejected, the above ratio has a noncentral F distribution which also depends on the noncentrality parameter, λ . This parameter is calculated as

$$\lambda = \bar{n}k \frac{\sigma_m^2}{\sigma_\varepsilon^2}$$

where

$$\sigma_m^2 = \sum_{i=1}^k \frac{n_i (\mu_i - \bar{\mu})^2}{N},$$

$$\bar{\mu} = \sum_{i=1}^k \frac{n_i \mu_i}{N},$$

$$N = \sum_{i=1}^k n_i,$$

and

$$\bar{n} = \frac{N}{k}.$$

Some authors use the symbol ϕ for the noncentrality parameter. The relationship between the two noncentrality parameters is

$$\phi = \sqrt{\frac{\lambda}{k}}.$$

Analysis of Covariance

The process of planning an experiment should include the following steps:

1. Determine an estimate of the within group standard deviation, σ . This may be done from prior studies, from experimentation with the Standard Deviation Estimation module, from pilot studies, or from crude estimates based on the range of the data. See the chapter on estimating the standard deviation for more details.
2. Determine a set of means that represent the group differences that you want to detect.
3. Determine the R-squared value between the response and the covariates.
4. Determine the appropriate group sample sizes that will ensure desired levels of α and β .

Power Calculations for ANCOVA

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

1. Determine the critical value, $F_{k-1, N-k-p, \alpha}$ where α is the probability of a type-I error and k , p , 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 based on the values of N , k , σ_m , ρ^2 , and σ .
3. Compute the power as the probability of being greater than $F_{k-1, N-k-p, \alpha}$ on a noncentral- F distribution with noncentrality parameter λ .

Technical Details for a Planned Comparison

The terminology of planned comparisons is identical to that of the one-way AOV, so the notation used above will be repeated here.

Suppose you want to test whether the contrast C

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

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

Define

$$\sigma_{mc} = \frac{\left| \sum_{i=1}^k c_i \mu_i \right|}{\sqrt{N \sum_{i=1}^k \frac{c_i^2}{n_i}}}$$

Define the noncentrality parameter λ_c , as

$$\lambda = \bar{n}k \frac{\sigma_{mc}^2}{\sigma_\varepsilon^2}$$

Power Calculations for Planned Comparisons

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

1. Determine the critical value, $F_{1,N-k-p,\alpha}$ where α is the probability of a type-I error and k 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 based on the values of N , k , σ_{mc} , ρ^2 , and σ .
3. Compute the power as the probability of being greater than $F_{1,N-k-p,\alpha}$ 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 *Sm*, *S*, *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*.

Analysis of Covariance

Sample Size / Groups
k (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 Hypothesized Means box and the Group Allocation Ratios box is controlled by this number.

Group Allocation Ratios

A set of positive, numeric values, one for each group, is entered here. 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, k . 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 k 1's will be used. This will result in $N1 = N2 = N3 = n$. That is, all sample sizes are equal to n .

n (Sample Size 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 Allocation Ratios numbers. If the Group Allocation Ratios numbers are represented by $m1, m2, m3, \dots, mk$ and this value is represented by n , the group sample sizes $N1, N2, N3, \dots, Nk$ 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 Allocation Ratios 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

As a second example, suppose there are three groups and the Group Allocation Ratios is $0.2,0.3,0.5$. When fractional allocation ratios values sum to one, n can be interpreted as the total sample size of all groups and the allocation ratios values 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.

Analysis of Covariance

Effect Size – Means

Hypothesized Means

Enter a set of hypothesized means, one for each group. These means represent the group centers under the alternative hypothesis (the null hypothesis is that they are equal). The standard deviation of these means (SM) is used in the power calculations to represent the average size of the differences among the means. The standard deviation of the means is calculated using the formula:

$$\sigma_m = \sqrt{\frac{\sum_{i=1}^k (\mu_i - \bar{\mu})^2}{k}}$$

This quantity gives the magnitude of the differences among the group means. Note that when all means are equal, σ_m is zero.

You should enter a set of means that give the pattern of differences you expect or the pattern that you wish to detect. For example, in a particular study involving three groups, your research might be “meaningful” if either of two treatment means is 50% larger than the control mean. If the control mean is 50, then you would enter 50,75,75 as the three means.

It is usually more intuitive to enter a set of mean values. However, it is possible to enter the standard deviation of the means directly by placing an S in front of the number (see below).

Some might wish to specify the alternative hypothesis as the effect size, f , which is defined as

$$f = \frac{\sigma_m}{\sigma}$$

If so, set $\sigma = 1$ and $\sigma_m = f$. Cohen (1988) has designated values of f less than 0.1 as *small*, values around 0.25 to be *medium*, and values over 0.4 to be *large*.

Entering a List of Means

If a set of numbers is entered without a leading S , they are assumed to be the hypothesized group means under the alternative hypothesis. Their standard deviation will be calculated and used in the calculations. Blanks or commas may separate the numbers. Note that it is not the values of the means themselves that is important, but only their differences. Thus, the mean values 0,1,2 produce the same results as the values 100,101,102.

If too few means are entered to match the number of groups, the last mean is repeated. For example, suppose that four means are needed and you enter 1,2 (only two means). **PASS** will treat this as 1,2,2,2. If too many values are entered, **PASS** will truncate the list to the number of means needed.

Examples:

5 20 60

2,5,7

-4,0,6,9

S Option

If an S is entered before the list of numbers, they are assumed to be values of σ_m , the standard deviations of the group means. A separate power calculation is made for each value. Note that this list can be a TO-BY phrase.

Examples:

S 4.7

S 4.3 5.7 4.2

S 10 to 20 by 2

Analysis of Covariance

Effect Size – Standard Deviation

S (Standard Deviation of Subjects)

This is σ , the standard deviation 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 *SD* 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

Effect Size – Covariates

Number of Covariates

This is the number of covariates (*X*'s) in the study. Because of the stringent assumptions, this value is usually set to 1 or 2.

R2 (R-Squared with Covariates)

This is the average R-squared value that is achieved within a group by the covariates. It must be between 0 and 1.

Planned Comparisons

Contrast Coefficients

If you want to analyze a specific planned comparison, enter a set of contrast coefficients here. The calculations will then refer to the hypothesis that the corresponding 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:

$$-3(\text{Mean1}) - (\text{Mean2}) + (\text{Mean3}) + 3(\text{Mean4})$$

The point to realize is that these numbers (the coefficients) are used to calculate a specific weighted average of the means which is to be compared against zero using a standard *F* (or *t*) test.

- **NONE or blank**

When the box is left blank or the word *None* is entered, this option is ignored.

- **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.

Analysis of Covariance

- **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 (k), 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.

Analysis of Covariance

Example 1 – Finding the Statistical Power

An experiment is being designed to compare the means of four groups using an F test with a significance level of 0.05. A covariate is available that is estimated to have an R-squared of 0.4 with the response. Previous studies have shown that the standard deviation within a group is 18. Note that this value ignores the covariate.

Treatment means of 40, 10, 10, and 10 represent clinically important treatment differences. To better understand the relationship between power, sample size, and R-squared, the researcher wants to compute the power for R-squared's of 0.2, 0.3, 0.4, and 0.5, and for several group sample sizes between 2 and 10. 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 **Analysis of Covariance** procedure window by expanding **Means**, then clicking on **Analysis of Covariance**, and then clicking on **Analysis of Covariance**. 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**.

<u>Option</u>	<u>Value</u>
Design Tab	
Solve For	Power
Alpha.....	0.05
k (Number of Groups)	4
Group Allocation Ratios.....	Equal
n (Sample Size per Group).....	2 to 14 by 2
Hypothesized Means	40 10 10 10
S (Standard Deviation of Subjects)	18
Number of Covariates.....	1
R2 (R-Squared with Covariates).....	0.2 to 0.5 by 0.1
Contrast Coefficients	None

Annotated Output

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

Numeric Results

Numeric Results											
Power	Average n	Groups (k)	Total N	Alpha	Beta	Std Dev of Means (Sm)	Standard Deviation (S)	Effect Size	Cov's	R2	
0.17245	2.0	4	8	0.05000	0.82755	12.99	18.00	0.7217	1	0.200	
0.19041	2.0	4	8	0.05000	0.80959	12.99	18.00	0.7217	1	0.300	
0.21428	2.0	4	8	0.05000	0.78572	12.99	18.00	0.7217	1	0.400	
0.24742	2.0	4	8	0.05000	0.75258	12.99	18.00	0.7217	1	0.500	
0.61111	4.0	4	16	0.05000	0.38889	12.99	18.00	0.7217	1	0.200	
0.67475	4.0	4	16	0.05000	0.32525	12.99	18.00	0.7217	1	0.300	
0.74725	4.0	4	16	0.05000	0.25275	12.99	18.00	0.7217	1	0.400	
0.82656	4.0	4	16	0.05000	0.17344	12.99	18.00	0.7217	1	0.500	
0.86165	6.0	4	24	0.05000	0.13835	12.99	18.00	0.7217	1	0.200	
0.90662	6.0	4	24	0.05000	0.09338	12.99	18.00	0.7217	1	0.300	
0.94625	6.0	4	24	0.05000	0.05375	12.99	18.00	0.7217	1	0.400	
0.97632	6.0	4	24	0.05000	0.02368	12.99	18.00	0.7217	1	0.500	

(report continues)

Analysis of Covariance

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.
- Borm, Franssen, and Lemmens. 2007. 'A simple sample size formula for analysis of covariance in randomized clinical trials.' J of Clinical Epidemiology, 60, 1234-1238.

Report Definitions

- Power is the probability of rejecting a false null hypothesis. It should be close to one.
- n is the average group sample size.
- k is the number of groups.
- Total N is the total sample size of all groups combined.
- 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.
- S_m is the standard deviation of the group means under the alternative hypothesis.
- Standard deviation is the within group standard deviation.
- The Effect Size is the ratio of S_m to standard deviation.
- Cov's is the number of covariates.
- R^2 (R-Squared) gives the strength of the relationship between the response and the covariates.

Summary Statements

In an analysis of covariance study, sample sizes of 2, 2, 2, and 2 are obtained from each of the 4 groups whose means are to be compared. The covariate has an R-squared of 0.200. The total sample of 8 subjects achieves 17% power to detect differences among the means versus the alternative of equal means using an F test with a 0.05000 significance level. The size of the variation in the means is represented by their standard deviation which is 12.99. The common standard deviation within a group is assumed to be 18.00.

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.

Average n

The average of the group sample sizes.

k

The number of groups.

Total N

The total sample size of the study.

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 that S_m is zero when S_m is actually equal to the value shown in the next column.

Std Dev of Means (S_m)

This is the standard deviation of the hypothesized means. It was computed from the hypothesized means. It is roughly equal to the average difference between the group means and the overall mean.

Once you have computed this, you can enter a range of values to determine the effect of the hypothesized means on the power.

Standard Deviation (S)

This is the within-group standard deviation. It was set in the Data window.

Analysis of Covariance

Effect Size

The effect size is the ratio of SM to S . It is an index of relative difference between the means that can be compared from study to study.

Cov's

This is the number of covariates.

R2

This is the value of R-squared.

Detailed Results Report

Details when Alpha = 0.05000, Power = 0.17245, SM = 12.99, S = 18.00, Cov's = 1, R2 = 0.20					
Group	Ni	Percent Ni of Total Ni	Mean	Deviation From Mean	Ni Times Deviation
1	2	25.00	40.00	22.50	45.00
2	2	25.00	10.00	7.50	15.00
3	2	25.00	10.00	7.50	15.00
4	2	25.00	10.00	7.50	15.00
ALL	8	100.00	17.50		

These reports show details of each row of the previous report.

Group

The number of the group shown on this line. The last line, labeled *ALL*, gives the average or the total as appropriate.

Ni

This is the sample size of each group. This column is especially useful when the sample sizes are unequal.

Percent Ni of Total Ni

This is the percentage of the total sample that is allocated to each group.

Mean

This is the value of the Hypothesized Mean. The final row gives the average for all groups.

Deviation From Mean

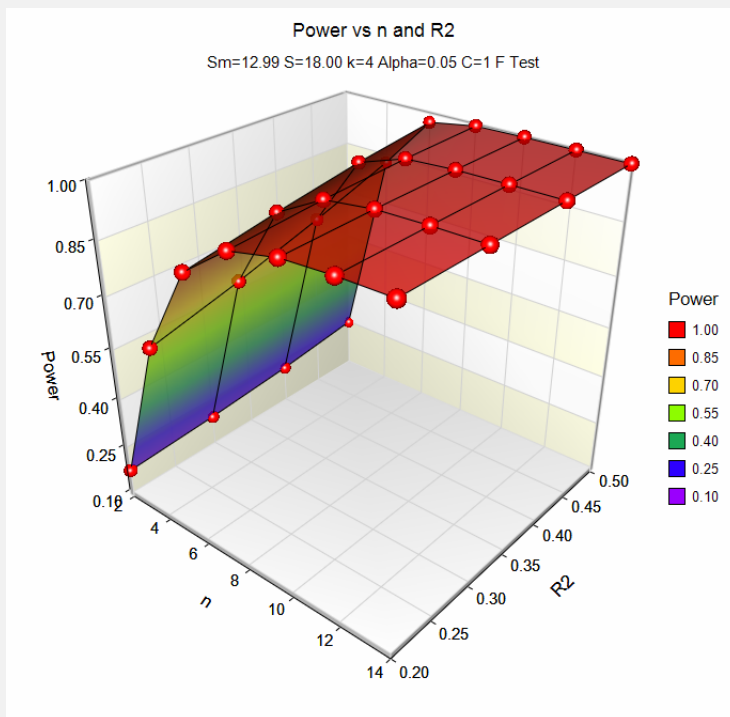
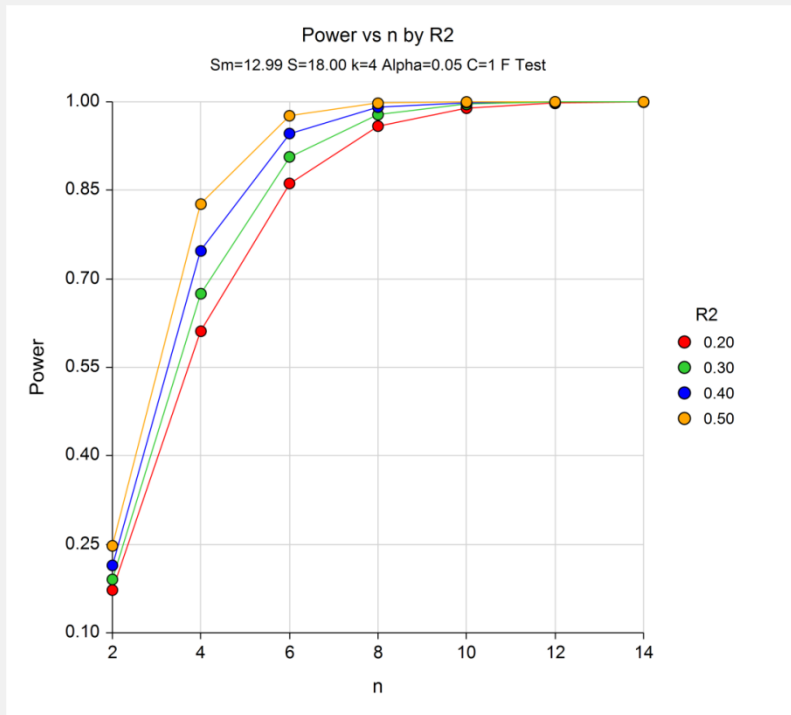
This is the absolute value of the mean minus the overall mean. Since Sm is the sum of the squared deviations, these values show the relative contribution to Sm .

Ni Times Deviation

This is the group sample size times the absolute deviation. It shows the combined influence of the size of the deviation and the sample size on Sm .

Analysis of Covariance

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 increasing the R-squared.

Note that the value of R-squared has a large impact on power for small sample sizes, but low for larger ones.

Analysis of Covariance

Example 2 – Validation using Borm, et al.

Borm, Fransen, and Lemmens (2007) page 1237 presents an example of determining a sample size in an experiment with 2 groups, mean difference of 0.6, standard deviation of 1.2, alpha of 0.05, one covariate with an R-squared of 0.25, and power of 0.80. They find a total sample size of 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 **Analysis of Covariance** procedure window by expanding **Means**, then clicking on **Analysis of Covariance**, and then clicking on **Analysis of Covariance**. 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**.

<u>Option</u>	<u>Value</u>
Design Tab	
Solve For	Sample Size
Power.....	0.80
Alpha.....	0.05
k (Number of Groups).....	2
Group Allocation Ratios.....	Equal
Hypothesized Means	0 0.6
S (Standard Deviation of Subjects)	1.2
Number of Covariates.....	1
R2 (R-Squared with Covariates).....	0.25
Contrast Coefficients	None

Output

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

Numeric Results

Numeric Results										
Power	Average n	Groups (k)	Total N	Alpha	Beta	Std Dev of Means (Sm)	Standard Deviation (S)	Effect Size	Cov's	R2
0.80752	49.0	2	98	0.05000	0.19248	0.30	1.20	0.2500	1	0.250

PASS also found $N = 98$. Note that Borm (2007) used calculations based on a normal approximation and obtained $N = 95$, but PASS uses exact calculations based on the non-central F distribution.