Chapter 591

Analysis of Covariance (ANCOVA)

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.

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 with 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 G groups each have a normal distribution and with means $\mu_1, \mu_2, \ldots, \mu_G$. Further suppose that each observation consists of a response measurement, Y, and one or more covariate measurements: $X_1, X_2, \ldots, X_p$. Further suppose that samples of $n_1, n_2, \ldots, n_G$ observations will be obtained from each of G groups. The multiple regression equation relating Y to the X's within the i^{th} group is

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

The $\beta's$ are the regression coefficients or slopes. Analysis of covariance assumes that, except for the intercepts, β_{0i}, 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

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

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

$$\lambda = \bar{N} G \frac{\sigma_m^2}{\sigma_\varepsilon^2}$$

where

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

$$\bar{\mu} = \sum_{i=1}^{G} \left(\frac{n_i}{N}\right) \mu_i$$

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

$$\bar{N} = \frac{N}{G}$$

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, $F_{G-1,N-G-p,\alpha}$ where α is the probability of a type-I error and G, 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 $\mu_i's$, calculate the noncentrality parameter λ based on the values of N, G, σ_m, ρ^2, and σ.

3. Compute the power as the probability of being greater than $F_{G-1,N-G-p,\alpha}$ from a noncentral-F distribution with noncentrality parameter λ.

Example 1 – Finding the 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^2 of 0.3 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^2, the researcher wants to compute the power for R^2's of 0.2, 0.3, and 0.4, and for several group sample sizes between 2 and 10. The sample sizes will be equal across all groups.

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 ....**Power**
Alpha...**0.05**
G (Number of Groups) **4**
Group Allocation Input Type ....**Equal to ni (Sample Size per Group)**
ni (Sample Size Per Group) **2 4 6 8 10**
µ1 Input Type...**Enter µ1 (Group Means|H1)**
µ1 (Group Means|H1) **40 10 10 10**
Standard Deviation Input Type ....**Enter σ (SD of Subjects), R² (with Covariates)**
σ (SD of Subjects) **18**
R² (R-Squared with Covariates).....................**0.2 0.3 0.4**
Number of Covariates....................................**1**

Output

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

Numeric Reports

Numeric Results

Solve For: Power
Number of Groups: 4
Number of Covariates: 1

	Sample Size		Group Means|H1		Standard Deviation		R-Squared with Covariates	
Power	**Total N**	**Group ni**	**Means $\mu 1$**	**SD σm**	**Subjects σ**	**Residuals σe**	**R²**	**Alpha**
0.17245	8	2	$\mu 1(1)$	12.99038	18	16.09969	0.2	0.05
0.19041	8	2	$\mu 1(1)$	12.99038	18	15.05988	0.3	0.05
0.21428	8	2	$\mu 1(1)$	12.99038	18	13.94274	0.4	0.05
0.61111	16	4	$\mu 1(1)$	12.99038	18	16.09969	0.2	0.05
0.67475	16	4	$\mu 1(1)$	12.99038	18	15.05988	0.3	0.05
0.74725	16	4	$\mu 1(1)$	12.99038	18	13.94274	0.4	0.05
0.86165	24	6	$\mu 1(1)$	12.99038	18	16.09969	0.2	0.05
0.90662	24	6	$\mu 1(1)$	12.99038	18	15.05988	0.3	0.05
0.94625	24	6	$\mu 1(1)$	12.99038	18	13.94274	0.4	0.05
0.95878	32	8	$\mu 1(1)$	12.99038	18	16.09969	0.2	0.05
0.97817	32	8	$\mu 1(1)$	12.99038	18	15.05988	0.3	0.05
0.99102	32	8	$\mu 1(1)$	12.99038	18	13.94274	0.4	0.05
0.98919	40	10	$\mu 1(1)$	12.99038	18	16.09969	0.2	0.05
0.99559	40	10	$\mu 1(1)$	12.99038	18	15.05988	0.3	0.05
0.99873	40	10	$\mu 1(1)$	12.99038	18	13.94274	0.4	0.05

Item	Values
$\mu 1(1)$	40, 10, 10, 10

Power   The probability of rejecting a false null hypothesis when the alternative hypothesis is true.
N   The total number of subjects in the study.
ni   The number of items sampled from each group.
$\mu 1$ The Group Means \| H1 is the set name and number of the group means under the alternative hypothesis. This is the set of means at which the power is calculated.
σm The standard deviation of the group means. Note that this value also depends on the group sample sizes.
σ The common standard deviation of the responses within each group.
σe The standard deviation of the residuals from a regression of the response on both groups and covariates.
R²   The R-Squared when the response is regressed on only the covariates.
Alpha   The probability of rejecting a true null hypothesis.

Group Sample Size Details

n	N	Group Sample Sizes	Group Allocation Proportions
n(1)	8	2, 2, 2, 2	0.25, 0.25, 0.25, 0.25
n(2)	16	4, 4, 4, 4	0.25, 0.25, 0.25, 0.25
n(3)	24	6, 6, 6, 6	0.25, 0.25, 0.25, 0.25
n(4)	32	8, 8, 8, 8	0.25, 0.25, 0.25, 0.25
n(5)	40	10, 10, 10, 10	0.25, 0.25, 0.25, 0.25

Summary Statements
──

An analysis of covariance design with 4 groups and 1 covariate will be used to test whether there is a difference among the 4 group means. The comparison will be made using an ANCOVA F-test with a Type I error rate (α) of 0.05. The standard deviation of the responses is assumed to be 18, with an R^2 between the response and the covariate of 0.2 (this corresponds to a standard deviation of the residuals of 16.09969). To detect the means 40, 10, 10, 10 (standard deviation of group means = 12.99038), with group sample sizes of 2, 2, 2, 2 (for a total of 8 subjects), the power is 0.17245.

──

Dropout-Inflated Sample Size
──

Dropout Rate	Sample Size N	Dropout-Inflated Enrollment Sample Size N'	Expected Number of Dropouts D
20%	8	10	2
20%	16	20	4
20%	24	30	6
20%	32	40	8
20%	40	50	10

──

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.
N	The evaluable sample size at which power is computed (as entered by the user). If N subjects are evaluated out of the N' subjects that are enrolled in the study, the design will achieve the stated power.
N'	The total number of subjects that should be enrolled in the study in order to obtain N evaluable subjects, based on the assumed dropout rate. N' is calculated by inflating N using the formula N' = N / (1 - DR), with N' 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.)
D	The expected number of dropouts. D = N' - N.

Dropout Summary Statements
──

Anticipating a 20% dropout rate, 10 subjects should be enrolled to obtain a final sample size of 8 subjects.

──

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

Keppel, Geoffrey. 1991. Design and Analysis - A Researcher's Handbook. Third Edition. Prentice Hall. Englewood Cliffs, New Jersey. See pages 323 - 324.

──

This report shows the numeric results of this power study.

Plots Section

Plots
_____
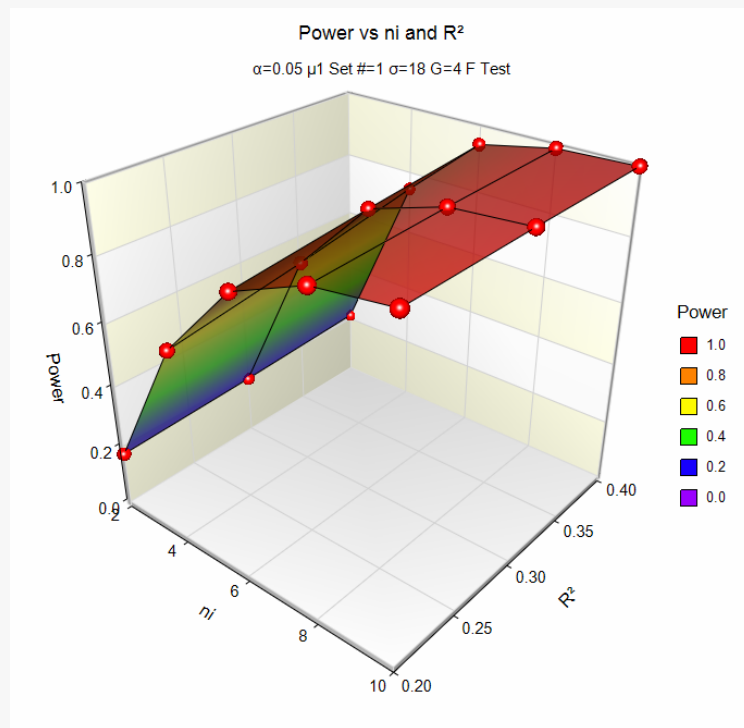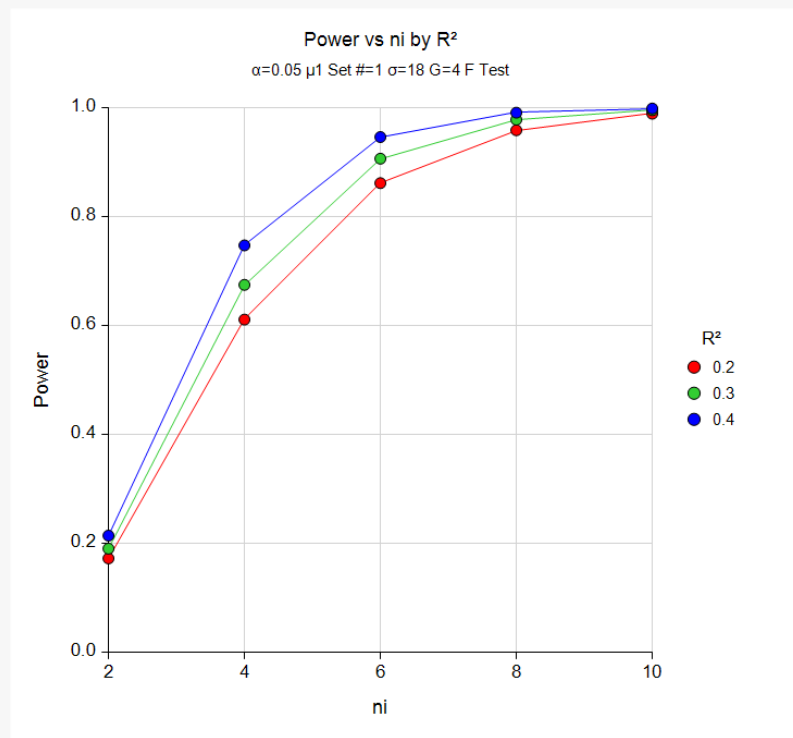
These plots give a visual presentation to the results in the Numeric Report. We can quickly see the impact on the power of increasing the sample size and increasing the R-squared.

Example 2 – Validation using Borm, et al. (2007)

Borm, Fransen, and Lemmens (2007) page 1237 present 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 found a total sample size of 95 using a normal approximation.

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**
Power..**0.8**
Alpha...**0.05**
G (Number of Groups) **2**
Group Allocation Input Type ....**Equal (n1 = ··· = nG)**
μ1 Input Type...**Enter μ1 (Group Means|H1)**
μ1 (Group Means|H1) **0 0.6**
Standard Deviation Input Type ....**Enter σ (SD of Subjects), R² (with Covariates)**
σ (SD of Subjects) **1.2**
R² (R-Squared with Covariates)....................**0.25**
Number of Covariates....................................**1**

Output

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

Numeric Results

Solve For: Sample Size
Number of Groups: 2
Number of Covariates: 1

	Sample Size		Group Means|H1		Standard Deviation		R-Squared with Covariates	
Power	Total N	Group ni	Means μ1	SD σm	Subjects σ	Residuals σe	R²	Alpha
0.80752	98	49	μ1(1)	0.3	1.2	1.03923	0.25	0.05

Item	Values
μ1(1)	0, 0.6

Group Sample Size Details

n	N	Group Sample Sizes	Group Allocation Proportions
n(1)	98	49, 49	0.5, 0.5

PASS calculated $N = 98$. 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.