

Chapter 560

Factorial Analysis of Variance

Introduction

A common task in research is to compare the average response across levels of one or more factor variables. Examples of factor variables are income level of two regions, nitrogen content of three lakes, or drug dosage. The factorial analysis of variance compares the means of two or more factors. F tests are used to determine statistical significance of the factors and their interactions. The tests are non-directional in that the null hypothesis specifies that all means are equal, and the alternative hypothesis simply states that at least one mean is different. This **PASS** module performs power analysis and sample size estimation for an analysis of variance design with up to three fixed factors.

In the following example, the responses of a weight loss experiment are arranged in a two-factor, fixed-effect, design. The first factor is diet (D1 and D2) and the second factor is dose level of a dietary drug (low, medium, and high). The twelve individuals available for this study were assigned at random to one of the six treatment groups (cells) so that there were two per group. The response variable was an individual's weight loss after four months.

Table of Individual Weight Losses			
	Dietary Drug Dose Level		
Diet	Low	Medium	High
D1	14, 16	15, 18	23, 28
D2	18, 21	18, 22	38, 39

Important features to note are that each table entry represents a different individual, and that the response variable (weight loss) is continuous, while the factors (Diet and Dose) are discrete.

Means can be calculated for each cell of the table. These means are shown in the table below. Note that we have added an additional row and column for the row, column, and overall means. The six means in the interior of this table are called the *cell means*.

Table of Means				
	Dietary Drug Dose Level			
Diet	Low	Medium	High	Total
D1	15.00	16.50	25.50	19.00
D2	19.50	20.00	38.50	26.00
Total	17.25	18.25	32.00	22.50

The Linear Model

A mathematical model may be formulated that underlies this experimental design. This model expresses each cell mean, μ_{ij} , as the sum of parameters called *effects*. A common linear model for a two-factor experiment is

$$\mu_{ij} = m + a_i + b_j + (ab)_{ij}$$

where $i = 1, 2, \dots, I$ and $j = 1, 2, \dots, J$. This model expresses the value of a cell mean as the sum of four components:

m the grand mean.

a_i the effect of the i^{th} level of factor A. Note that $\sum a_i = 0$.

b_j the effect of the j^{th} level of factor B. Note that $\sum b_j = 0$.

ab_{ij} the combined effect of the i^{th} level of factor A and the j^{th} level of factor B. Note that $\sum (ab)_{ij} = 0$.

Another way of stating this model for the two-factor case is

$$\text{Cell Mean} = \text{Overall Effect} + \text{Row Effect} + \text{Column Effect} + \text{Interaction Effect}.$$

Since this model is the sum of various constants, it is called a *linear model*.

Calculating the Effects

We will now calculate the effects for our example. We will let Drug Dose correspond to factor A and Diet correspond to factor B.

Step 1 – Remove the Grand Mean

Remove the grand mean from the table of means by subtracting 22.50 from each entry. The values in the margins are the *effects* of the corresponding factors.

Table of Mean Weight Losses After Subtracting the Grand Mean				
	Dietary Drug Dose Level			
Diet	Low	Medium	High	Overall
D1	-7.50	-6.00	3.00	-3.50
D2	-3.00	-2.50	16.00	3.50
Overall	-5.25	-4.25	9.50	22.50

Factorial Analysis of Variance

Step 2 – Remove the Effects of Factor B (Diet)

Subtract the Diet effects (-3.50 and 3.50) from the entries in those rows.

Table of Mean Weight Losses After Subtracting the Diet Effects				
	Dietary Drug Dose Level			
Diet	Low	Medium	High	Overall
D1	-4.00	-2.50	6.50	-3.50
D2	-6.50	-6.00	12.50	3.50
Overall	-5.25	-4.25	9.50	22.50

Step 3 – Remove the Effects of Factor A (Drug Dose)

Subtract the Drug Dose effects (-5.25, -4.25, and 9.50) from the rest of the entries in those columns. This will result in a table of effects.

Table of Effects				
	Dietary Drug Dose Level			
Diet	Low	Medium	High	Overall
D1	1.25	1.75	-3.00	-3.50
D2	-1.25	-1.75	3.00	3.50
Overall	-5.25	-4.25	9.50	22.50

We have calculated a table of effects for the two-way linear model. Each cell mean can be calculated by summing the appropriate entries from this table.

The estimated linear effects are:

$$m = 22.50$$

$$a_1 = -5.25 \quad a_2 = -4.25 \quad a_3 = 9.50$$

$$b_1 = -3.50 \quad b_2 = 3.50$$

$$ab_{11} = 1.25 \quad ab_{21} = 1.75 \quad ab_{31} = -3.00$$

$$ab_{12} = -1.25 \quad ab_{22} = -1.75 \quad ab_{32} = 3.00$$

The six cell means are calculated from these effects as follows:

$$15.00 = 22.50 - 5.25 - 3.50 + 1.25$$

$$19.50 = 22.50 - 5.25 + 3.50 - 1.25$$

$$16.50 = 22.50 - 4.25 - 3.50 + 1.75$$

$$20.00 = 22.50 - 4.25 + 3.50 - 1.75$$

$$25.50 = 22.50 + 9.50 - 3.50 - 3.00$$

$$38.50 = 22.50 + 9.50 + 3.50 + 3.00$$

Analysis of Variance Hypotheses

The hypotheses that are tested in an analysis of variance table concern the effects, so in order to conduct a power analysis you must have a firm grasp of their meaning. For example, we would usually test the following hypotheses:

1. Are there differences in weight loss among the three drug doses? That is, are the drug dose effects all zero? This hypothesis is tested by the F test for factor A , which tests whether the standard deviation of the α_i is zero.
2. Is there a difference in weight loss between the two diets? That is, are the diet effects all zero? This hypothesis is tested by the F test for factor B , which tests whether the standard deviation of the b_j is zero.
3. Are there any diet-dose combinations that exhibit a weight loss that cannot be explained by diet and/or drug dose singly? This hypothesis is tested by the F test for the AB interaction, which tests whether the standard deviation of the $(ab)_{ij}$ is zero.

Each of these hypotheses can be tested at a different alpha level and different precision. Hence each can have a different power. One of the tasks in planning such an experiment is to determine a sample size that yields necessary power values for each of these hypothesis tests. This is accomplished using this program module.

Definition of Terms

Factorial designs evaluate the effect of two or more categorical variables (called *factors*) on a response variable by testing hypotheses about various averages. These designs are popular because they allow experimentation across a wide variety of conditions and because they evaluate the *interaction* of two or more factors. Interaction is the effect that may be attributed to a combination of two or more factors, but not to one factor singly.

A *factor* is a variable that relates to the response. Either the factor is discrete by nature (as in location or gender) or has been made discrete by collapsing a continuous variable (as in income level or age group). The term *factorial* implies that all possible combinations of the factors being studied are included in the design.

A *fixed* factor is one in which all possible *levels* (categories) are considered. Examples of fixed factors are gender, dose level, and country of origin. They are different from *random* factors which represent a random selection of individuals from the population described by the factor. Examples of random factors are people living within a region, a sample of schools in a state, or a selection of labs. Again, a fixed factor includes the range of interest while a random factor includes only a sample of all possible levels.

A factorial design is analyzed using the analysis of variance. When only fixed factors are used in the design, the analysis is said to be a *fixed-effects analysis of variance*.

Suppose a group of individuals have agreed to be in a study involving six treatments. In a *completely randomized factorial design*, each individual is assigned at random to one of the six groups and then the treatments are applied. In some situations, the randomization occurs by randomly selecting individuals from the populations defined by the treatment groups. The designs analyzed by this module are completely randomized factorial designs.

Power Calculations

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

1. Determine the critical value, $F_{df1,df2,\alpha}$ where $df1$ is the numerator degrees of freedom, $df2$ is the denominator degrees of freedom, and α is the probability of a type-I error (significance level). Note that the F test is a two-tailed test as no logical direction is assigned in the alternative hypothesis.
2. Calculate the standard deviation of the hypothesized effects, using the formula:

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

where the e_i are effect values and k is the number of effects. Note that the average effect will be zero by construction, so this formula reduces to

$$\sigma_m = \sqrt{\frac{\sum_{i=1}^k (e_i)^2}{k}}$$

3. Compute the noncentrality parameter λ using the relationship:

$$\lambda = N \frac{\sigma_m^2}{\sigma^2}$$

where N is the total number of subjects.

4. Compute the power as the probability of being greater than $F_{df1,df2,\alpha}$ on a noncentral-F distribution with noncentrality parameter λ .

Example

In the example discussed earlier, the standard deviation of the dose effects is

$$\begin{aligned}\sigma_m(A) &= \sqrt{\frac{(-5.25)^2 + (-4.25)^2 + 9.50^2}{3}} \\ &= 6.729908\end{aligned}$$

the standard deviation of the diet effects is

$$\begin{aligned}\sigma_m(B) &= \sqrt{\frac{(-3.5)^2 + 3.5^2}{2}} \\ &= 3.5\end{aligned}$$

Factorial Analysis of Variance

and the standard deviation of the interaction effects is

$$\begin{aligned}\sigma_m(AB) &= \sqrt{\frac{1.25^2 + (-1.25)^2 + 1.75^2 + (-1.75)^2 + (-3.00)^2 + 3.00^2}{6}} \\ &= 2.131119\end{aligned}$$

Standard Deviation of Effects (of Means)

In the two-sample t-test case, the alternative hypothesis was represented as the difference between two group means. Unfortunately, for three or more groups, there is no simple extension of the two group difference. Instead, you must hypothesize a set of effects and calculate the value of σ_m .

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

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

where σ is the standard deviation of values within a cell (see Sigma below). If you want to use f , set $\sigma = 1$ and then f is always equal to σ_m so that the values you enter for σ_m will be the values of 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*. You should come up with your own cutoff values for low, medium, and high.

When you are analyzing the power of an existing analysis of variance table, you can compute the values of σ_m for each term from its mean square or F ratio using the following formulas:

$$\sigma_m = \sqrt{\frac{df_{\text{numerator}} MS_{\text{numerator}}}{N}}$$

or

$$\sigma_m = \sqrt{\frac{df_{\text{numerator}}(F)(MSE)}{N}}$$

where N is the total number of observations, MSE is the mean square error, df is the numerator degrees of freedom, MS is the mean square of the term, and F is the F ratio of the term. If you do this, you are setting the sample effects equal to the population effects for the purpose of computing the power.

Example 1 – Power after a Study

This example will explain how to calculate the power of F tests from data that have already been collected and analyzed.

Analyze the power of the experiment that was given at the beginning of this chapter. These data were analyzed using the analysis of variance procedure in **NCSS** and the following results were obtained.

Analysis of Variance Table

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level	Power (Alpha=0.05)
A: Dose	2	543.5	271.75	50.95	0.000172*	1.000000
B: Diet	1	147	147	27.56	0.001920*	0.990499
AB	2	54.5	27.25	5.11	0.050629	0.588884
S	6	32	5.333333			
Total (Adjusted)	11	777				
Total	12					

* Term significant at alpha = 0.05

Means and Standard Error Section

Term	Count	Mean	Standard Error	Effect
All	12	22.5		22.50
A: Dose				
High	4	32	1.154701	9.50
Low	4	17.25	1.154701	-4.25
Medium	4	18.25	1.154701	-5.25
B: Diet				
D1	6	19	0.942809	-3.50
D2	6	26	0.942809	3.50
AB: Dose, Diet				
High, D1	2	25.5	1.632993	-3.00
High, D2	2	38.5	1.632993	3.00
Low, D1	2	15	1.632993	1.25
Low, D2	2	19.5	1.632993	-1.25
Medium, D1	2	16.5	1.632993	1.75
Medium, D2	2	20	1.632993	-1.75

To analyze these data, we can enter the means for factors A and B as well as the AB interaction effects.

Alternatively, we could have calculated the standard deviation of the interaction. This can be done in either of two ways.

Factorial Analysis of Variance

Using mean square for AB (27.25), the degrees of freedom for AB (2), and the total sample size (12), the standard deviation of the AB -interaction effects is calculated as follows

$$\sigma_m(AB) = \sqrt{\frac{2(27.25)}{12}} = 2.1311$$

Using the formula based on the effects, the standard deviation of the AB -interaction effects is calculated as follows

$$\sigma_m(AB) = \sqrt{\frac{3^2 + 3^2 + 1.25^2 + 1.25^2 + 1.75^2 + 1.75^2}{6}} = 2.1311$$

The value of σ is estimated from the square root of the mean square error:

$$\sigma = \sqrt{5.333333} = 2.3094$$

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 for All Terms	0.05
Number of Factors	2
A - Levels	3
A - Means	List of Means
A - List of Means	17.25 18.25 32
B - Levels	2
B - Means	List of Means
B - List of Means	19 26
A*B	Checked
Effects	Std Dev of Effects
Std Dev of Effects	2.1311
σ (Std Dev of Subjects)	2.3094
n (Size Per Group)	2

Factorial Analysis of Variance

Output

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

Numeric Results

Solve For: **Power**

Number of Groups: 6

Term	Power	Sample Size		Degrees of Freedom		Standard Deviation		Effect Size σ_m/σ	Alpha
		Group n	Total N	df1	df2	Means σ_m	Subjects σ		
A(3)	1.00000	2	12	2	6	6.72991	2.3094	2.91414	0.05
B(2)	0.99050	2	12	1	6	3.50000	2.3094	1.51555	0.05
A*B	0.58888	2	12	2	6	2.13110	2.3094	0.92279	0.05

Power The probability of rejecting a false null hypothesis when the alternative hypothesis is true.

n The average sample size of per group (treatment combination).

N The total sample size of all groups combined.

df1 The numerator degrees of freedom.

df2 The denominator degrees of freedom.

σ_m The standard deviation of the group means or effects.

σ The subject-to-subject standard deviation.

σ_m/σ The Effect Size (f).

Alpha The probability of rejecting a true null hypothesis.

Summary Statements

A (fixed factor) factorial design with two factors (one with 3 levels and the other with 2 levels, resulting in 6 treatment combinations) will be used to test whether there are differences among the levels of the factors. Each term will be tested using a linear model analysis of variance F-test with a Type I error rate (α) of 0.05. The standard deviation within each treatment combination is assumed to be 2.3094. With a sample size of 2 subjects per treatment combination, the total number of subjects is 12. For factor A, to detect a standard deviation among level means of 6.72991 (an effect size of 2.91414), the power is 1. For factor B, to detect a standard deviation among level means of 3.5 (an effect size of 1.51555), the power is 0.9905. For the A*B interaction effect, to detect a standard deviation among effects of 2.1311 (an effect size of 0.92279), the power is 0.58888.

Dropout-Inflated Sample Size

Group	Dropout Rate	Sample Size n	Dropout-Inflated Enrollment Sample Size n'	Expected Number of Dropouts D
1 - 6	20%	2	3	1
Total		12	18	6

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.

n The evaluable sample size for each group 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 number of subjects that should be enrolled in each group 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 in each group. $D = n' - n$.

Factorial Analysis of Variance

Dropout Summary Statement

Anticipating a 20% dropout rate, group sizes of 3, 3, 3, 3, 3, and 3 subjects should be enrolled to obtain final group sample sizes of 2, 2, 2, 2, 2, and 2 subjects.

References

Neter, J., Kutner, M., Nachtsheim, C., and Wasserman, W. 1996. Applied Linear Statistical Models. Richard D. Irwin, Inc. Chicago, Illinois.
 Winer, B.J. 1991. Statistical Principles in Experimental Design. Third Edition. McGraw-Hill. New York, NY.

This report shows the power for each of the three factors.

It is important to emphasize that these power values are for the case when the effects associated with the alternative hypotheses are equal to those given by the data. It will usually be more informative to calculate the power for other values as well.

Term

This is the term (main effect or interaction) from the analysis of variance model being displayed on this line.

Power

This is the power of the F test for this term. Note that since adding and removing terms changes the denominator degrees of freedom (df_2), the power depends on which other terms are included in the model.

n

This is the sample size per group (treatment combination). Fractional values indicate an unequal allocation among the cells.

N

This is the total sample size for the complete design.

df1

This is the numerator degrees of freedom of the F test.

df2

This is the denominator degrees of freedom of the F test. This value depends on which terms are included in the AOV model.

Standard Deviation of Means (σ_m)

This is the standard deviation of the means (or effects). It represents the size of the differences among the effects that is to be detected by the analysis. If you have entered hypothesized means, only their standard deviation is displayed here.

Factorial Analysis of Variance

Standard Deviation of Subjects (σ)

This is the standard deviation of the means (or effects). It represents the size of the differences among the effects that is to be detected by the analysis. If you have entered hypothesized means, only their standard deviation is displayed here.

Effect Size (σ_m/σ)

This is the standard deviation of the means divided by the standard deviation of subjects. It provides an index of the magnitude of the difference among the means that can be detected by this design.

Alpha

This is the significance level of the F test. This is the probability of a type-I error given the null hypothesis of equal means and zero effects.

Example 2 – Finding the Sample Size

In this example, we will investigate the impact of increasing the sample size on the power of each of the seven tests in the analysis of variance table of a three-factor experiment. The first factor (A) has two levels, the second factor (B) has three levels, and the third factor (C) has four levels. This creates a design with $2 \times 3 \times 4 = 24$ treatment combinations.

All values of σ_m will be set equal to 0.2, σ is set equal to 1.0, alpha is set to 0.05, and minimum power is 0.90. Determine the required sample size so that the F-tests of all terms meet the minimum power requirement.

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
Based on Term	All
Minimum Power	0.90
Alpha for All Terms	0.05
Number of Factors	3
A - Levels	2
A - Means	As a Std Dev
A - Std Dev	0.2
B - Levels	3
B - Means	As a Std Dev
B - Std Dev	0.2
C - Levels	4
C - Means	As a Std Dev
C - Std Dev	0.2
A*B	Checked
Effects	Std Dev of Effects
Std Dev of Effects	0.2
<i>All other interactions are set the same as A*B</i>	
σ (Std Dev of Subjects)	1

Factorial Analysis of Variance

Output

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

Numeric Results

Solve For: **Sample Size**

Number of Groups: 24

Term	Power	Sample Size		Degrees of Freedom		Standard Deviation		Effect Size σ_m/σ	Alpha
		Group n	Total N	df1	df2	Means σ_m	Subjects σ		
A(2)	0.98931	19	456	1	432	0.2	1	0.2	0.05
B(3)	0.97523	19	456	2	432	0.2	1	0.2	0.05
C(4)	0.95982	19	456	3	432	0.2	1	0.2	0.05
A*B	0.97523	19	456	2	432	0.2	1	0.2	0.05
A*C	0.95982	19	456	3	432	0.2	1	0.2	0.05
B*C	0.91028	19	456	6	432	0.2	1	0.2	0.05
A*B*C	0.91028	19	456	6	432	0.2	1	0.2	0.05

The necessary sample size is 19 per group for a total of 456 subjects.

Example 3 – Latin Square Design

This example shows how to study the power of a complicated experimental design like a Latin square. Suppose you want to run a five-level Latin square design. This design consists of three factors each at five levels. One factor is associated with the columns of the square, a second factor is associated with the rows of the square, and a third factor is associated with the letters of the square. There are only $5 \times 5 = 25$ observations used instead of the $5 \times 5 \times 5 = 125$ that would normally be required. The Latin square design has reduced the number of observations by 80%.

The 80% decrease in observations comes at a price—the interaction terms must be ignored. If you can legitimately assume that the interactions are zero, the Latin square (or some other design which reduces the number of observations) is an efficient design to use.

We will now show you how to analyze the power of the F tests from such a design. The key is to enter 0.2 (which is $25/125$) for n and set all the interaction indicators off.

Since all three factors have five levels, the power of the three F tests will be the same if σ_m is the same. Hence, we can try three different sets of hypothesized means. The first set will be five means 0.1 units apart. The second set will be five means 0.5 units apart. The third set will be five means 1.0 unit apart. The standard deviation will be set to 1.0. All alpha levels will be set at 0.05.

The sample size per cell is set at 0.2 and 0.4. This will result in total sample sizes of 25 (one replication) and 50 (two replications).

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 3** 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 for All Terms	0.05
Number of Factors	3
A - Levels	5
A - Means	List of Means
A - List of Means	1.0 1.1 1.2 1.3 1.4
B - Levels	5
B - Means	List of Means
B - List of Means	1.0 1.5 2.0 2.5 3.0
C - Levels	5
C - Means	List of Means
C - List of Means	1 2 3 4 5
<i>All interactions are unchecked</i>	
σ (Std Dev of Subjects)	1
n (Size Per Group)	0.2 0.4

Factorial Analysis of Variance

Output

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

Numeric Results

Solve For: **Power**

Number of Groups: 125

Term	Power	Sample Size		Degrees of Freedom		Standard Deviation		Effect Size σ_m/σ	Alpha
		Group n	Total N	df1	df2	Means σ_m	Subjects σ		
A(5)	0.0681	0.2	25	4	12	0.141	1	0.141	0.05
B(5)	0.6367	0.2	25	4	12	0.707	1	0.707	0.05
C(5)	0.9987	0.2	25	4	12	1.414	1	1.414	0.05
A(5)	0.0984	0.4	50	4	37	0.141	1	0.141	0.05
B(5)	0.9774	0.4	50	4	37	0.707	1	0.707	0.05
C(5)	1.0000	0.4	50	4	37	1.414	1	1.414	0.05

In the first design in which $N = 25$, only the power of the test for C is greater than 0.9. Of course, this power value also depends on the value of the standard deviation of subjects.

It is interesting to note that doubling the sample size did not double the power!

Example 4 – Validation using Winer (1991)

Winer (1991) pages 428-429 presents the power calculations for a two-way design in which factor *A* has two levels and factor *B* has three levels. Winer provides estimates of the sum of squared *A* effects (1.0189), sum of squared *B* effects (5.06), and sum of squared interaction effects (42.11). The mean square error is 8.83 and the per cell sample size is 3. All alpha levels are set to 0.05.

Winer's results are approximate because he has to interpolate in the tables that he is using. He finds the power of the *F* test for factor *A* to be between 0.10 and 0.26. He estimates it as 0.17. The exact power of the *F* test for factor *B* is not given. Instead, the range is found to be between 0.26 and 0.36. The power of the *F* test for the *A*B* interaction is "approximately" 0.86.

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 4** 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 for All Terms	0.05
Number of Factors	2
A - Levels	2
A - Means	As a Std Dev
A - Std Dev of Effects.....	0.714
B - Levels	3
B - Means	As a Std Dev
B - Standard Deviation of Effects.....	1.3
A*B.....	Checked
Effects.....	Std Dev of Effects
Std Dev of Effects	2.65
σ (Std Dev of Subjects).....	2.97
n (Size Per Group)	3

Factorial Analysis of Variance

Output

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

Numeric Results

Solve For: **Power**

Number of Groups: 6

Term	Power	Sample Size		Degrees of Freedom		Standard Deviation		Effect Size σ_m/σ	Alpha
		Group n	Total N	df1	df2	Means σ_m	Subjects σ		
A(2)	0.1558	3	18	1	12	0.714	2.97	0.240	0.05
B(3)	0.2918	3	18	2	12	1.300	2.97	0.438	0.05
A*B	0.8534	3	18	2	12	2.650	2.97	0.892	0.05

The power of the test for factor A is 0.16 which is between 0.10 and 0.26. It is close to the interpolated 0.17 that Winer obtained from his tables. The power of the test for factor B is 0.29 which is between 0.26 and 0.36. The power of the test for the A*B interaction is 0.85 which is close to the interpolated 0.86 that Winer obtained from his tables.

Example 5 – Validation using Prihoda (1983)

Prihoda (1983) pages 7-8 presents the power calculations for a two-way design with the following pattern of means:

		Factor B				
		1	2	3	4	All
Factor A	1	41	34	30	27	33
	2	33	24	22	29	27
All		37	29	26	28	30

The means may be manipulated to show the overall mean, the main effects, and the interaction effects:

		Factor B				
		1	2	3	4	All
Factor A	1	1	2	1	-4	3
	2	-1	-2	-1	4	-3
All		7	-1	-4	-2	30

Based on the above effects, Prihoda calculates the power of the interaction test when the sample size per group is 6, 8, 10, 12, and 14 to be 0.34, 0.45, 0.56, 0.65, and 0.73. The σ_m for the interaction is 2.345208. The mean square error is 64 and the alpha level is 0.05.

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 5** 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 for All Terms	0.05
Number of Factors	2
A - Levels	2
A - Means	List of Means
A - Std Dev of Effects.....	33 27
B - Levels	4
B - Means	List of Means
B - Standard Deviation of Effects.....	37 29 26 28
A*B.....	Checked
Effects	Std Dev of Effects
Std Dev of Effects	2.345208
σ (Std Dev of Subjects).....	8
n (Size Per Group)	6 8 10 12 14

Factorial Analysis of Variance

Output

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

Numeric Results

Solve For: **Power**

Number of Groups: 8

Term	Power	Sample Size		Degrees of Freedom		Standard Deviation		Effect Size σ_m/σ	Alpha
		Group n	Total N	df1	df2	Means σ_m	Subjects σ		
A(2)	0.7175	6	48	1	40	3.000	8	0.375	0.05
B(4)	0.8368	6	48	3	40	4.183	8	0.523	0.05
A*B	0.3372	6	48	3	40	2.345	8	0.293	0.05
A(2)	0.8385	8	64	1	56	3.000	8	0.375	0.05
B(4)	0.9387	8	64	3	56	4.183	8	0.523	0.05
A*B	0.4510	8	64	3	56	2.345	8	0.293	0.05
A(2)	0.9113	10	80	1	72	3.000	8	0.375	0.05
B(4)	0.9792	10	80	3	72	4.183	8	0.523	0.05
A*B	0.5556	10	80	3	72	2.345	8	0.293	0.05
A(2)	0.9529	12	96	1	88	3.000	8	0.375	0.05
B(4)	0.9935	12	96	3	88	4.183	8	0.523	0.05
A*B	0.6475	12	96	3	88	2.345	8	0.293	0.05
A(2)	0.9757	14	112	1	104	3.000	8	0.375	0.05
B(4)	0.9981	14	112	3	104	4.183	8	0.523	0.05
A*B	0.7254	14	112	3	104	2.345	8	0.293	0.05

Prihoda only presents the power for the interaction test at each sample size. You can check to see that the results match Prihoda's exactly.

Example 6 – Validation using Neter, Kutner, Nachtsheim, and Wasserman (1996)

Neter, Kutner, Nachtsheim, and Wasserman (1996) page 1057 presents a power analysis of a two-factor experiment in which factor *A* has three levels and factor *B* has two levels. The significance level is 0.05, the standard deviation is 3.0, and *n* is 2. They calculate a power of about 0.89 for the test of factor *A* when the three means are 50, 55, and 45.

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 6** 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 for All Terms	0.05
Number of Factors	2
A - Levels	3
A - Means	List of Means
A - List of Means	50 55 45
B - Levels	2
B - Means	As a Std Dev
B - Standard Deviation of Effects	1
A*B	Checked
Effects	Std Dev of Effects
Std Dev of Effects	1
σ (Std Dev of Subjects)	3
n (Size Per Group)	2

Output

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

Numeric Results

Solve For: **Power**

Number of Groups: 6

Term	Power	Sample Size		Degrees of Freedom		Standard Deviation		Effect Size σ_m/σ	Alpha
		Group n	Total N	df1	df2	Means σ_m	Subjects σ		
A(3)	0.9016	2	12	2	6	4.082	3	1.361	0.05
B(2)	0.1648	2	12	1	6	1.000	3	0.333	0.05
A*B	0.1178	2	12	2	6	1.000	3	0.333	0.05

Note that the power of 0.90 that **PASS** has calculated is within rounding of the 0.89 that Neter *et al.* (1996) calculated.