Chapter 550

One-Way Analysis of Variance F-Tests

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. *F* 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.

The methods described here are usually applied to the one-way experimental design. This design is an extension of the design used for the two-sample *t* test. Instead of two groups, there are three or more groups. With careful modifications, this procedure may be used to test interaction terms as well.

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 assumptions of the one-way analysis of variance are:

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 the One-Way ANOVA

Suppose k groups each have a normal distribution and equal means $(\mu_1 = \mu_2 = \cdots = \mu_k)$. Let $n_1 = n_2 = \cdots = n_k$ denote the number of subjects in each group and let N denote the total sample size of all groups. Let $\overline{\mu}_W$ denote the weighted mean of all groups. That is

$$\overline{\mu}_W = \sum_{i=1}^{k}\left(\frac{n_i}{N}\right)\mu_i$$

Let σ denote the common standard deviation of all groups.

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

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

where

$$\sigma_m = \sqrt{\sum_{i=1}^{k}\frac{n_i\left(\mu_i - \overline{\mu}_W\right)^2}{N}}$$

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

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

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 appropriate group sample sizes that will ensure desired levels of α and β. Although it is tempting to set all group sample sizes equal, it is easy to show that putting more subjects in some groups than in others may have better power than keeping group sizes equal (see Example 4).

Power Calculations for One-Way ANOVA

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

1. Determine the critical value, $F_{k-1,N-k,\alpha}$ 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 based on the values of N, k, σ_m, and σ.

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

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 *σm*, *σ*, *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 – σm

Input σm Using

Select the method used to specify *σm*, the standard deviation of the hypothesized means (*μ1, μ2, ..., μG*). This value represents the amount of variation in the means that you want to detect with the F-test. It is analogous to the detectable difference in the two-sample t-test.

There are two ways in which you can enter *σm*: enter one or more values of σm directly or enter a list of means, one for each group, from which PASS calculates σm.

Means (μ1, μ2, ..., μG)

Enter *G* group means assumed by the alternative hypothesis. The standard deviation of these means (*σm*) is used in the power analysis to represent the magnitude of the differences among the means that is detectable by the design.

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*.

When entering the list, blanks or commas separate the items in 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, 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 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

σm (√[Σ(μi-μ)²]/G])

Enter one or more values of *σm*, the standard deviation of the group means. This value approximates the average size of the differences among the means that is to be detected. By detected, we mean that if *σm* is this large, the null hypothesis of mean equality will likely be rejected.

The value of *σm* is calculated using the formula $\sigma m = \sqrt{(\Sigma(\mu i-\mu)^2)/G}$, where *G* is the number of levels (categories), *μ* is the mean of the group means, and *μi* is the mean of the ith group.

Range

Since these are standard deviations, they must be greater than zero. They must be in the same scale as σ, the between subject standard deviaton.

Effect Size – Standard Deviation

σ (Standard Deviation)

This is *σ*, the 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 an *F* test with a significance level of either 0.01 or 0.05. Previous studies have shown that the standard deviation is 18. Treatment means of 40, 10, 10, and 10 represent clinically important treatment differences. 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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. 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	**0.01 0.05**
G (Number of Groups)	**4**
Group Allocation Ratios	**Equal**
n (Subjects per Group)	**2 to 14 by 2**
Input σm Using	**List of means (μi's) from which σm is calculated**
Means	**40 10 10 10**
K (Means Multiplier)	**1**
σ (Standard Deviation)	**18**
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

Power	Average n	G	Total N	K	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.0424	2.00	4	8	1.00	12.99	18.00	0.7217	0.0100
0.2389	4.00	4	16	1.00	12.99	18.00	0.7217	0.0100
0.5058	6.00	4	24	1.00	12.99	18.00	0.7217	0.0100
0.7269	8.00	4	32	1.00	12.99	18.00	0.7217	0.0100
0.8670	10.00	4	40	1.00	12.99	18.00	0.7217	0.0100
0.9414	12.00	4	48	1.00	12.99	18.00	0.7217	0.0100
0.9762	14.00	4	56	1.00	12.99	18.00	0.7217	0.0100
0.1751	2.00	4	8	1.00	12.99	18.00	0.7217	0.0500
0.5216	4.00	4	16	1.00	12.99	18.00	0.7217	0.0500
0.7733	6.00	4	24	1.00	12.99	18.00	0.7217	0.0500
0.9064	8.00	4	32	1.00	12.99	18.00	0.7217	0.0500
0.9651	10.00	4	40	1.00	12.99	18.00	0.7217	0.0500
0.9880	12.00	4	48	1.00	12.99	18.00	0.7217	0.0500
0.9961	14.00	4	56	1.00	12.99	18.00	0.7217	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. It should be close to one.
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.
σm is the standard deviation of the group means under the alternative hypothesis.
σ is the within group standard deviation.
The Effect Size is the ratio of σm and σ.
Alpha is the probability of rejecting a true null hypothesis. It should be small.

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. The total sample of 8 subjects achieves 4% power to detect differences among the means versus the alternative of equal means using an F test with a 0.0100 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.

Detailed Results Report

Details when Alpha = 0.01000, Power = 0.0424, σm = 12.99, σ = 18.00

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		

(report continues)

This report shows the 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

The 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 σm is the sum of the squared deviations, these values show the relative contribution to σm.

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 σm.

Plots Section

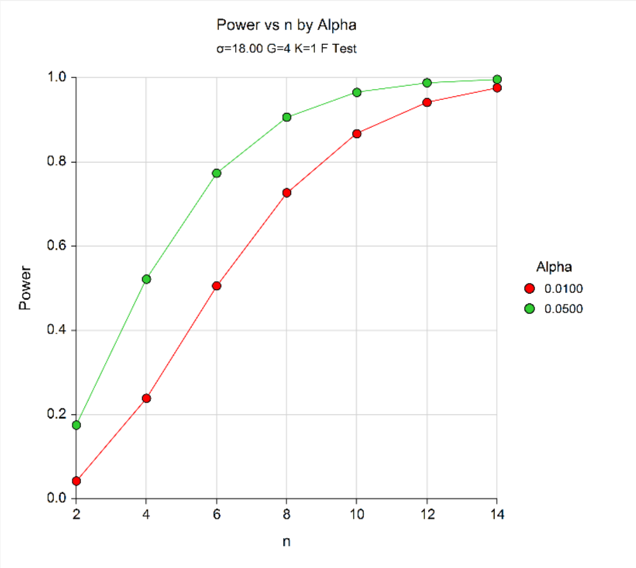

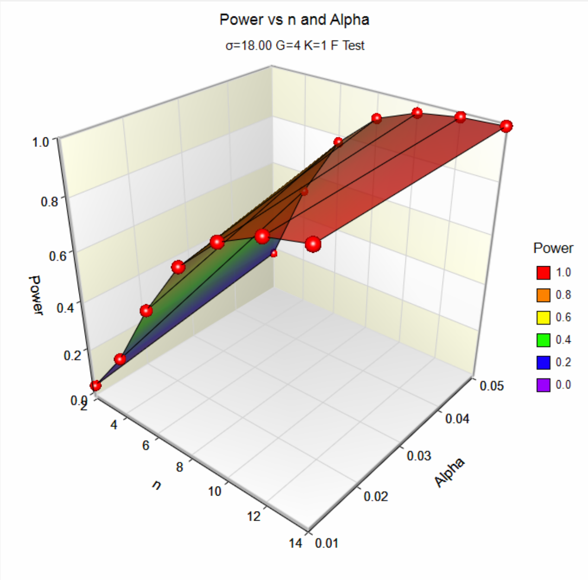

These plots give a visual presentation to the results in the Numeric Report. We can quickly see the impact on the power of increasing the sample size and the increase in the significance level.

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 – Power after a Study

This example will cover the situation in which you are calculating the power of a one-way analysis of variance *F* test on data that have already been collected and analyzed.

An experiment included a control group and two treatment groups. Each group had seven individuals. A single response was measured for each individual and recorded in the following table.

Control	T1	T2
452	646	685
674	547	658
554	774	786
447	465	536
356	759	653
654	665	669
558	767	557

When analyzed using the one-way analysis of variance procedure in NCSS, the following results were obtained.

Analysis of Variance Table

Source Term	DF	Sum of Squares	Mean Square	F-Ratio	Prob Level
A (...)	2	75629.8	37814.9	3.28	0.061167
S(A)	18	207743.4	11541.3		
Total (Adjusted)	20	283373.3			
Total	21				

Means Section

Group	Count	Mean
Control	7	527.8571
T1	7	660.4286
T2	7	649.1429

The significance level (Prob Level) was 0.061—not enough for statistical significance. The researcher had hoped to show that the treatment groups had higher response levels than the control group. He could see that the group means followed this pattern since the mean for *T1* was about 25% higher than the control mean and the mean for *T2* was about 23% higher than the control mean. He decided to calculate the power of the experiment using these values of the means. (We do not recommend this approach because the power should be calculated for the minimum difference among the means that is of interest, not at the values of the sample means.)

The data entry for this problem is simple. The only entry that is not straight forward is finding an appropriate value for the standard deviation. Since the standard deviation is estimated by the square root of the mean square error, it is calculated as $\sqrt{11541.3} = 107.4304$.

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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. 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)	**7**
Input σm Using	**List of means (μi's) from which σm is calculated**
Means	**527.8571 660.4286 649.1429**
K (Means Multiplier)	**1**
σ (Standard Deviation)	**107.4304**

Output

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

Numeric Results

Numeric Results
Means: 527.8571 660.4286 649.1429

Power	Average n	G	Total N	K	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.5479	7.00	3	21	1.00	60.01	107.43	0.5586	0.0500

The power is only 0.55. That is, there was only a 55% chance of rejecting a false null hypothesis. It is important to understand this power statement is conditional, so we will state it in detail. Given that the population means are equal to the sample means (that σm is 60.01) and the population standard deviation is equal to 107.43, the probability of rejecting the false null hypothesis is 0.55. If the population means are different from the sample means (which they must be), the power is different.

Example 3 – Finding the Sample Size Necessary to Reject

Continuing with the last example, we will determine how large the sample size would need to have been for alpha = 0.05 and beta = 0.20.

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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. You may then make the appropriate entries as listed below, or open **Example 3** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Sample Size**
Power	**0.8**
Alpha	**0.05**
G (Number of Groups)	**3**
Group Allocation Ratios	**Equal**
Input σm Using	**List of means (μi's) from which σm is calculated**
Means	**527.8571 660.4286 649.1429**
K (Means Multiplier)	**1**
σ (Standard Deviation of Subjects)	**107.4304**

Output

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

Numeric Results

Numeric Results
Means: 527.8571 660.4286 649.1429

Power	Average n	G	Total N	K	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.8251	12.00	3	36	1.00	60.01	107.43	0.5586	0.0500

The required sample size is 12 per group or 36 subjects.

Example 4 – Using Unequal Sample Sizes

Continuing with the last example, consider the impact of allowing the group sample sizes to be unequal. Since the control group is being compared to two treatment groups, the mean of the control group is assumed to be different from those of the treatment groups. In this situation, experience has shown that adding extra subjects to the control group can increase the power. In a separate analysis, the power with 11 subjects per group was found to be 0.7851—not quite the required 0.80.

We will try moving two subjects from each treatment group into the control group. This will give an experimental design with 15 in the control group and 9 in each of the treatment 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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. You may then make the appropriate entries as listed below, or open **Example 4** by going to the **File** menu and choosing **Open Example Template**.

Pay particular attention to how the sample size parameters were changed. The value of n is set to one so that it is essentially ignored. The Group Sample Size Pattern contains the three unequal sample sizes.

Option	Value
Design Tab	
Solve For	**Power**
Alpha	**0.05**
G (Number of Groups)	**3**
Group Allocation Ratios	**15 9 9**
n (Subjects per Group)	**1**
Input σm Using	**List of means (μi's) from which σm is calculated**
Means	**527.8571 660.4286 649.1429**
K (Means Multiplier)	**1**
σ (Standard Deviation)	**107.4304**

Output

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

Numeric Results

Numeric Results
Means: 527.8571 660.4286 649.1429

Power	Average n	G	Total N	K	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.82967	11.00	3	33	0.05000	0.17033	63.34	107.43	0.5896

The power of 0.82967 achieved with the 33 subjects in this design is slightly higher than the power of 0.82511 that was achieved with the 36 subjects in the equal group size design. Apparently, unequal sample allocation can achieve better power!

We suggest that you try several different sample allocations. You will find that the optimum sample allocation depends on the values of the hypothesized means.

You should keep in mind that power may not be the only goal of the experiment. Other goals may include finding confidence intervals for each of the group means. And the narrowness of the width of the confidence interval is directly related to the sample size.

Example 5 – Minimum Detectable Difference

It may be useful to determine the minimum detectable difference among the means that can be found at the experimental conditions. This amounts to finding *σm*.

Continuing with the previous example, find *σm* for a wide range of sample sizes when alpha is 0.05 and power is 0.80 or 0.90.

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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. You may then make the appropriate entries as listed below, or open **Example 5** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**σm (Std Dev of Means)**
Power	**0.8 0.9**
Alpha	**0.05**
G (Number of Groups)	**3**
Group Allocation Ratios	**Equal**
n (Subjects per Group)	**2 3 5 8 10 15 20 40 60 80 100**
σ (Standard Deviation)	**107.4304**

Output

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

Numeric Results and Plots

Numeric Results

Power	Average n	G	Total N	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.8000	2.00	3	6	244.31	107.43	2.2741	0.0500
0.8000	3.00	3	9	145.82	107.43	1.3573	0.0500
0.8000	5.00	3	15	98.08	107.43	0.9130	0.0500
0.8000	8.00	3	24	73.23	107.43	0.6817	0.0500
0.8000	**10.00**	**3**	**30**	**64.42**	**107.43**	**0.5997**	**0.0500**
0.8000	15.00	3	45	51.54	107.43	0.4797	0.0500
0.8000	20.00	3	60	44.21	107.43	0.4115	0.0500
0.8000	40.00	3	120	30.83	107.43	0.2870	0.0500
0.8000	60.00	3	180	25.07	107.43	0.2333	0.0500
0.8000	80.00	3	240	21.66	107.43	0.2016	0.0500
0.8000	100.00	3	300	19.35	107.43	0.1801	0.0500

(report continues)

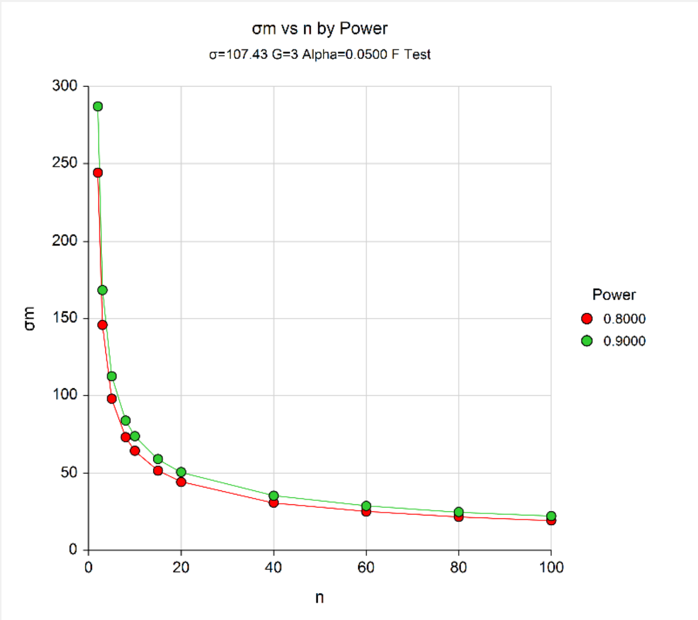

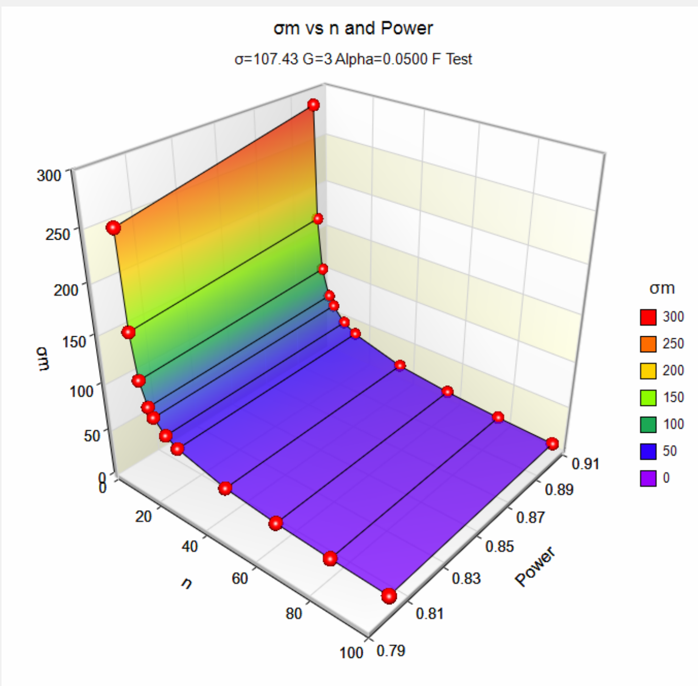

These plots show the relationships between power, sample size, and detectable difference. Several conclusions are possible, but the most impressive is the sharp elbow in the curve that occurs near n = 10 when *σm* is about 64.

How to interpret a *σm* of 64? One way is to find a set of means that have a standard deviation of 64. To do this, press the σ button in the lower right corner of the panel to load the Standard Deviation Estimator module. Under the *Data* tab, enter the three values 0, 0, and 64 as a starting point. These values have a standard deviation of 30. Doubling the 64 to 128 results in an SD of 60. Increasing the 128 to 136 results in the desired SD of 64. The difference between the minimum and the maximum of these three values is 136. Hence the minimum detectable difference is about 136 for a *σm* of 64 when n is 9 and the power is 80%.

Example 6 – Validation using Fleiss (1986)

Fleiss (1986) page 374 presents an example of determining a sample size in an experiment with 4 groups; means of 9.775, 12, 12, and 14.225; standard deviation of 3; alpha of 0.05, and beta of 0.20. He finds a sample size of 11 per group.

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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. You may then make the appropriate entries as listed below, or open **Example 6** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Sample Size**
Power	**0.8**
Alpha	**0.05**
G (Number of Groups)	**4**
Group Allocation Ratios	**Equal**
Input σm Using	**List of means (μi's) from which σm is calculated**
Means	**9.775 12 12 14.225**
K (Means Multiplier)	**1**
σ (Standard Deviation)	**3**
Reports Tab	
Show Details	***Checked***

Output

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

Numeric Results

Numeric Results
Means: 9.775 12 12 14.225

Power	Average n	G	Total N	K	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.8027	11.00	4	44	1.00	1.57	3.00	0.5244	0.0500

Details when Alpha = 0.05000, Power = 0.80273, σm = 1.57, σ = 3.00

Group	Ni	Percent Ni of Total Ni	Mean	Deviation From Mean	Ni Times Deviation
1	11	25.00	9.78	2.23	24.48
2	11	25.00	12.00	0.00	0.00
3	11	25.00	12.00	0.00	0.00
4	11	25.00	14.23	2.23	24.48
ALL	44	100.00	12.00		

PASS also found n = 11. Note that Fleiss used calculations based on a normal approximation, but **PASS** uses exact calculations based on the non-central F distribution.

Example 7 – Validation using Desu (1990)

Desu (1990) page 48 presents an example of determining a sample size in an experiment with 3 groups; means of 0, -0.2553, and 0.2553; standard deviation of 1; alpha of 0.05, and beta of 0.10. He finds a sample size of 99 per group.

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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. You may then make the appropriate entries as listed below, or open **Example 7** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Sample Size**
Power	**0.9**
Alpha	**0.05**
G (Number of Groups)	**3**
Group Allocation Ratios	**Equal**
Input σm Using	**List of means (μi's) from which σm is calculated**
Hypothesized Means	**0 -0.2553 0.2553**
K (Means Multiplier)	**1**
σ (Standard Deviation)	**1**

Output

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

Numeric Results

Numeric Results
Means: 0 -0.2553 0.2553

Power	Average n	G	Total N	K	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.9028	99.00	3	297	1.00	0.21	1.00	0.2085	0.0500

PASS also found *n* = 99.

Example 8 – Validation using Kirk (1982)

Kirk (1982) pages 140-144 presents an example of determining a sample size in an experiment with 4 groups; means of 2.75, 3.50, 6.25, and 9.0; standard deviation of 1.20995; alpha of 0.05, and beta of 0.05. He finds a sample size of 3 per group.

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 F-Tests** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then **ANOVA F-Test**, and then clicking on **One-Way Analysis of Variance F-Tests**. You may then make the appropriate entries as listed below, or open **Example 8** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Sample Size**
Power	**0.95**
Alpha	**0.05**
G (Number of Groups)	**4**
Group Allocation Ratios	**Equal**
Input σm Using	**List of means (μi's) from which σm is calculated**
Hypothesized Means	**2.75 3.5 6.25 9**
K (Means Multiplier)	**1**
σ (Standard Deviation)	**1.20995**

Output

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

Numeric Results

Numeric Results
Means: 9.775 12 12 14.225

Power	Average n	G	Total N	K	Std Dev of Means σm	Standard Deviation σ	Effect Size	Alpha
0.9977	3.00	4	12	1.00	2.47	1.21	2.0376	0.0500

PASS also found *n* = 3.