

Chapter 559

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Introduction

This procedure analyzes the power and significance level of tests which may be used to test statistical hypotheses about location (means and medians) in a one-way experimental design. For each scenario that is set up, two simulations are run. One simulation estimates the significance level and the other estimates the power.

The four tests that are compared in this procedure are

1. One-Way ANOVA F-Test
2. Kruskal-Wallis Test of Ranks
3. Terry-Hoeffding Normal Scores Test
4. Van der Waerden Normal Quantiles Test

Technical Details

Computer simulation allows us to estimate the power and significance level that is actually achieved by a test procedure in situations that are not mathematically tractable. Computer simulation was once limited to mainframe computers. But, in recent years, as computer speeds have increased, simulation studies can be completed on desktop and laptop computers in a reasonable period of time.

The steps to a simulation study are

1. Specify how the test is carried out. This includes indicating how the test statistic is calculated and how the significance level is specified.
2. Generate random samples from the distributions specified by the alternative hypothesis. Calculate the test statistics from the simulated data and determine if the null hypothesis is accepted or rejected. Tabulate the number of rejections and use this to calculate the test's power.
3. Generate random samples from the distributions specified by the null hypothesis. Calculate each test statistic from the simulated data and determine if the null hypothesis is accepted or rejected. Tabulate the number of rejections and use this to calculate the test's significance level.

Power Comparison of Tests of Means in One-Way Designs (Simulation)

- Repeat steps 2 and 3 several thousand times, tabulating the number of times the simulated data leads to a rejection of the null hypothesis. The power is the proportion of simulated samples in step 2 that lead to rejection. The significance level is the proportion of simulated samples in step 3 that lead to rejection.

Generating Random Distributions

Two methods are available in PASS to simulate random samples. The first method generates the random variates directly, one value at a time. The second method generates a large pool (over 10,000) of random values and then draws the random numbers from this pool. This second method can cut the running time of the simulation by 70%.

The second method begins by generating a large pool of random numbers from the specified distributions. Each of these pools is evaluated to determine if its mean is within a small relative tolerance (0.0001) of the target mean. If the actual mean is not within the tolerance of the target mean, individual members of the population are replaced with new random numbers if the new random number moves the mean towards its target. Only a few hundred such swaps are required to bring the actual mean to within tolerance of the target mean. This population is then sampled with replacement using the uniform distribution. We have found that this method works well as long as the size of the pool is at least the maximum of twice the number of simulated samples desired and 10,000.

Tests

The details of the four tests that are compared here are found in the chapters that discuss each procedure and they are not duplicated here.

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Design and Simulation tabs. For more information about the options of other tabs, go to the Procedure Window chapter.

Design Tab

The Design tab contains the parameters and options needed to describe the experimental design, the data distributions, the error rates, and sample sizes.

Alpha and Simulations

Alpha

This option specifies one or more values for the probability of a type-I error, alpha. 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.

Alpha values must be between zero and one. Historically, 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 single value or a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Simulations

This option specifies the number of simulations, M . The larger the number of simulations, the longer the running time and the more accurate the results.

The precision of the simulated power estimates are calculated from the binomial distribution. Thus, confidence intervals may be constructed for various power values using the binomial distribution. The following table gives an estimate of the precision that is achieved for various simulation sizes when the power is either 0.50 or 0.95.

Power Comparison of Tests of Means in One-Way Designs (Simulation)

The table values are interpreted as follows: a 95% confidence interval of the true power is given by the power reported by the simulation plus and minus the 'Precision' amount given in the table.

Simulation Size M	Precision when Power = 0.50	Precision when Power = 0.95
100	0.100	0.044
500	0.045	0.019
1000	0.032	0.014
2000	0.022	0.010
5000	0.014	0.006
10000	0.010	0.004
50000	0.004	0.002
100000	0.003	0.001

Notice that a simulation size of 1000 gives a precision of plus or minus 0.01 when the true power is 0.95. Also note that as the simulation size is increased beyond 5000, there is only a small amount of additional accuracy achieved.

Number of Groups and Sample Size Allocation

Number of Groups

Enter the number of groups in the design. The program can handle between 2 and 20 groups.

Group Allocation

Select the method used to allocate subjects to groups. The choices are

- **Equal ($n = n_1 = n_2 = \dots$)**

All group sample sizes are the same. The value(s) of n are entered in the n (*Group Size*) box that will appear immediately below this box. See also n (*Group Size*) below.

- **Enter base group size and group multipliers**

The group sample sizes are found by multiplying the corresponding *Multiplier* value (found in the Power Simulation section) times the *Base Group Size* that will appear immediately below this box. See also *Base Group Size* and *Multiplier* below.

- **Enter total sample size and group percentages**

The group sample sizes are found by taking the corresponding percentage (found in the Power Simulation section) of the N (*Total Size*) value that will appear immediately below this box. See also N (*Total Size*) and *Percent of N* below.

Note that the *Percent of N* values are adjusted so they sum to 100%.

- **Enter individual group sample sizes**

Enter each group's sample size directly in the corresponding *Sample Size* box in the Power Simulation section. See also *Sample Size* below.

n (Group Size)

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

Base Group Size

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

Power Comparison of Tests of Means in One-Way Designs (Simulation)

The individual group samples sizes are determined by multiplying this value times the corresponding *Multiplier* value entered in the Power Simulation section.

If the Multiplier numbers are represented by

$$m_1, m_2, m_3, \dots$$

and this value is represented by n , the group sample sizes are calculated as follows:

$$n_1 = [n(m_1)]$$

$$n_2 = [n(m_2)]$$

$$n_3 = [n(m_3)]$$

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

For example, suppose there are three groups and the multipliers are set to 1, 2, and 3. If n is 5, the resulting group sample sizes will be 5, 10, and 15.

N (Total Size)

This is the total 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 individual group samples sizes are determined by multiplying this value times the corresponding *Percent of N* value entered in the Power Simulation section.

If the Percent values are represented by

$$p_1, p_2, p_3, \dots$$

and this value is represented by N , the group sample sizes are calculated as follows:

$$n_1 = [N(p_1)]$$

$$n_2 = [N(p_2)]$$

$$n_3 = [N(p_3)]$$

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

For example, suppose there are three groups and the percentages are set to 25, 25, and 50. If N is 36, the resulting group sample sizes will be 9, 9, and 18.

Power Simulation

These options specify the distributions to be used in the power simulation, one row per group. The first option specifies distribution. The second option, if visible, specifies the sample size of that group.

Distribution

Specify the distribution of each group under the alternative hypothesis, H_1 . This distribution is used in the simulation that determines the power.

A fundamental quantity in a power analysis is the amount of variation among the group means. In fact, in classical power analysis formulas, this variation is summarized as the standard deviation of the means. You must pay particular attention to the values you give to the means of these distributions because they are fundamental to the interpretation of the simulation.

For convenience in specifying a range of values, the parameters of the distribution can be specified using numbers or letters. If letters are used, their values are specified in the *Parameter Values* boxes below.

Following is a list of the distributions that are available and the syntax used to specify them. Each of the parameters should be replaced with a number or parameter name.

Power Comparison of Tests of Means in One-Way Designs (Simulation)**Distributions with Common Parameters**

Beta(Shape1, Shape2, Min, Max)

Binomial(P,N)

Cauchy(Mean,Scale)

Constant(Value)

Exponential(Mean)

Gamma(Shape,Scale)

Gumbel(Location,Scale)

Laplace(Location,Scale)

Logistic(Location,Scale)

Lognormal(Mu,Sigma)

Multinomial(P1,P2,P3,...,Pk)

Normal(Mean,Sigma)

Poisson(Mean)

TukeyGH(Mu,S,G,H)

Uniform(Min, Max)

Weibull(Shape,Scale)

Distributions with Mean and SD Parameters

BetaMS(Mean,SD,Min,Max)

BinomialMS(Mean,N)

GammaMS(Mean,SD)

GumbelMS(Mean,SD)

LaplaceMS(Mean,SD)

LogisticMS(Mean,SD)

LognormalMS(Mean,SD)

UniformMS(Mean,SD)

WeibullMS(Mean,SD)

Details of writing mixture distributions, combined distributions, and compound distributions are found in the chapter on *Data Simulation* and will not be repeated here.

Finding the Value of the Mean of a Specified Distribution

The mean of a distribution created as a linear combination of other distributions is found by applying the linear combination to the individual means. However, the mean of a distribution created by multiplying or dividing other distributions is not necessarily equal to applying the same function to the individual means. For example, the mean of $4 \text{ Normal}(4, 5) + 2 \text{ Normal}(5, 6)$ is $4*4 + 2*5 = 26$, but the mean of $4 \text{ Normal}(4, 5) * 2 \text{ Normal}(5, 6)$ is not exactly $4*4*2*5 = 160$ (although it is close).

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Multiplier

The group allocation *Multiplier* for this group is entered here. Typical values of this parameter are 0.5, 1, and 2.

- **Solve For = Power**

The group sample size is found by multiplying this value times the *Base Group Size* value and rounding up to the next whole number.

- **Solve For = Sample Size**

The group sample size is found by multiplying this value times the *base group size* that is being searched for and rounding up to the next whole number.

Percent of N

This is the percentage of the total sample size that is allocated to this group. The individual group samples sizes are determined by multiplying this value times the *N (Total Size)* value entered above.

If these values are represented by p_1, p_2, p_3, \dots the group sample sizes are calculated as follows:

$$n_1 = [N(p_1)]$$

$$n_2 = [N(p_2)]$$

$$n_3 = [N(p_3)]$$

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

For example, suppose there are three groups and these percentages are set to 25, 25, and 50. If N is 36, the resulting group sample sizes will be 9, 9, and 18.

Sample Size

This option allows you to enter the sample size of this group directly.

Power Simulation – Parameter Values for Group Distributions

These options specify the names and values for the parameters used in the distributions.

Name

Up to six sets of named parameter values may be used in the simulation distributions. This option lets you select an appropriate name for each set of values. Possible names are

- **M1 to M5**

You might use M1 to M5 for means.

- **SD**

You might use SD for standard deviations.

- **A to T**

You might use these letters for location, shape, and scale parameters.

Value(s)

These values are substituted for the *Parameter Name* (M1, M2, SD, A, B, C, etc.) in the simulation distributions. If more than one value is entered, a separate calculation is made for each value.

You can enter a single value such as “2” or a series of values such as “0 2 3” or “0 to 3 by 1”.

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Alpha Simulation

These options specify the distributions to be used in the alpha simulation, one row per group. The alpha simulation generates an estimate of the actual alpha value that will be achieved by the test procedure. The magnitude of the differences of the means of these distributions, which is often summarized as the standard deviation of the means, represents the magnitude of the mean differences specified under H_0 . Usually, the means are assumed to be equal under H_0 , so their standard deviation should be zero except for rounding.

Specify Alpha Distributions

This option lets you choose how you will specify the alpha distributions. Possible choices are

- **All equal to the group 1 distribution of the power simulation**
Often, the first group will represent a neutral group such as the control group. This option indicates that the first group of the power distributions should be used for all of the alpha simulation groups. Typically, this is what you want since the null hypothesis assumes that all distributions are the same.
- **All equal to a custom distribution**
Set all group distributions in the alpha simulation equal to the custom alpha distribution specified below.
- **Enter each distribution separately**
Each group's simulation distribution is entered individually below.

Custom Alpha Distribution

Specify the alpha distribution to be used for all groups. The syntax is the same as that for the Power Simulation Distribution (see above) and will not be repeated here.

Individual Alpha Distributions

Specify a separate distribution for each group. The syntax is the same as that for the Power Simulation Distribution (see above) and will not be repeated here.

Keep in mind that the null hypothesis assumes that all distributions are the same.

Simulations Tab

The option on this tab controls the generation of the random numbers. For complicated distributions, a large pool of random numbers from the specified distributions is generated. Each of these pools is evaluated to determine if its mean is within a small relative tolerance (0.0001) of the target mean. If the actual mean is not within the tolerance of the target mean, individual members of the population are replaced with new random numbers if the new random number moves the mean towards its target. Only a few hundred such swaps are required to bring the actual mean to within tolerance of the target mean. This population is then sampled with replacement using the uniform distribution. We have found that this method works well as long as the size of the pool is at least the maximum of twice the number of simulated samples desired and 10,000.

Random Number Pool Size

This is the size of the pool of values from which the random samples will be drawn. Pools should be at least the minimum of 10,000 and twice the number of simulations. You can enter *Automatic* and an appropriate value will be calculated.

If you do not want to draw numbers from a pool, enter 0 here.

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Example 1 – Power at Various Sample Sizes

This example repeats Example 1 of four individual procedures that are compared in this procedure.

An experiment is being designed to compare the means of four groups using an F test with a significance level of 0.05. Previous studies have shown that the standard deviation within a group 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 group sample sizes of 4, 8, and 12. The group sample sizes are equal.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Power Comparison of Tests of Means in One-Way Designs (Simulation)** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then clicking on **ANOVA F-Test**, and then clicking on **Power Comparison of Tests of Means in One-Way Designs (Simulation)**. 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	
Alpha.....	0.05
Simulations	5000
Number of Groups	4
Group Allocation	Equal
n (Group Size)	4 8 12
Group 1 Distribution.....	Normal(M1 SD)
Group 2 Distribution.....	Normal(M2 SD)
Group 3 Distribution.....	Normal(M2 SD)
Group 4 Distribution.....	Normal(M2 SD)
M1	40
M2	10
SD	18
Specify Alpha Distributions.....	All equal to the group 1 distribution of the power simulation

Annotated Output

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

Numeric Results Report

Simulation Summary		
Group	Power Distribution	Alpha Distribution
1	Normal(M1 SD)	Normal(M1 SD)
2	Normal(M2 SD)	Normal(M1 SD)
3	Normal(M2 SD)	Normal(M1 SD)
4	Normal(M2 SD)	Normal(M1 SD)
Number of Groups		4
Random Number Pool Size		10000
Number of Simulations Per Run		5000

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Numeric Results

Row	Total Sample Size N	AOV F Test Power	Kruskal Wallis Test Power	Terry Hoeffding Test Power	Van der Waer'n Test Power	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	SD
1	16	0.517	0.372	0.354	0.358	0.050	13.0	18.0	40.0	10.0	18.0
2	32	0.911	0.868	0.879	0.877	0.050	13.0	17.9	40.0	10.0	18.0
3	48	0.988	0.982	0.983	0.984	0.050	13.0	18.0	40.0	10.0	18.0

Definitions of the Power Comparison of Various Tests Report

Row identifies this line of the report. It will be used cross-reference this line in other reports.
 H0 is the null hypothesis that all σ 's are equal.
 H1 is the alternative hypothesis that at least one σ is different from the others.
 Total Sample Size (N) is the total sample size of all groups combined.
 Power is the probability of rejecting H0 when it is false calculated by the power simulation. It is usually different from the Target Power because of the discrete nature of the sample size, the violation of test assumptions, and the approximate nature of the test.
 AOV F Test is the usual analysis of variance F-test.
 Kruskal Wallis Test is the nonparametric analog of the F-test based on ranks.
 Terry Hoeffding Test is the Terry Hoeffding expected normal scores test.
 Van der Waer'n Test is the Van der Waerden normal quantiles test.
 Target Alpha is the desired probability of rejecting a true null hypothesis at which the tests were run.
 Std Dev of H1 μ 's is the standard deviation of the group means assumed by the power simulation. This is a measure of the effect size.
 Mean of H1 σ 's is the mean of the within-group standard deviations assumed during the power simulation. This measures the variation.

Alpha Comparison of Location Tests

Row	Total Sample Size N	AOV F Test Alpha	Kruskal Wallis Test Alpha	Terry Hoeffding Test Alpha	Van der Waer'n Test Alpha	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	SD
1	16	0.049	0.034	0.031	0.031	0.050	0.0	18.0	40.0	10.0	18.0
2	32	0.051	0.042	0.040	0.040	0.050	0.0	18.0	40.0	10.0	18.0
3	48	0.051	0.045	0.045	0.046	0.050	0.0	18.0	40.0	10.0	18.0

Definitions of the Alpha Comparison of Various Tests Report

Row identifies this line of the report. It will be used cross-reference this line in other reports.
 H0 is the null hypothesis that all σ 's are equal.
 H1 is the alternative hypothesis that at least one σ is different from the others.
 Total Sample Size (N) is the total sample size of all groups combined.
 Alpha is the probability of rejecting H0 when it is true calculated by the alpha simulation.
 AOV F Test is the usual analysis of variance F-test.
 Kruskal Wallis Test is the nonparametric analog of the F-test based on ranks.
 Terry Hoeffding Test is the Terry Hoeffding expected normal scores test.
 Van der Waer'n Test is the Van der Waerden normal quantiles test.
 Target Alpha is the desired probability of rejecting a true null hypothesis at which the tests were run.
 Std Dev of H1 μ 's is the standard deviation of the group means assumed by the power simulation. This is a measure of the effect size.
 Mean of H1 σ 's is the mean of the within-group standard deviations assumed during the power simulation. This measures the variation.

These reports show the output for this run. The definitions of each column are shown below the report..

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Detailed Results Reports

Details for Row 1 of the Numeric Results Report

Group	ni	Pct of N	H0 μ	H1 μ	H0 σ	H1 σ
1	4	25.0	40.0	40.0	18.3	18.2
2	4	25.0	40.0	10.0	18.2	18.1
3	4	25.0	40.0	10.0	17.7	17.9
4	4	25.0	40.0	10.0	17.9	17.8
Group	N		Std Dev of H0 μ 's	Std Dev of H1 μ 's	Mean of H0 σ 's	Mean of H1 σ 's
All	16	100.0	0.0	13.0	18.0	18.0

Details for Row 2 of the Numeric Results Report

Group	ni	Pct of N	H0 μ	H1 μ	H0 σ	H1 σ
1	8	25.0	40.0	40.0	17.9	17.6
2	8	25.0	40.0	10.0	18.0	18.1
3	8	25.0	40.0	10.0	18.1	17.8
4	8	25.0	40.0	10.0	17.8	17.9
Group	N		Std Dev of H0 μ 's	Std Dev of H1 μ 's	Mean of H0 σ 's	Mean of H1 σ 's
All	32	100.0	0.0	13.0	18.0	17.9

Details for Row 3 of the Numeric Results Report

Group	ni	Pct of N	H0 μ	H1 μ	H0 σ	H1 σ
1	12	25.0	40.0	40.0	18.1	18.2
2	12	25.0	40.0	10.0	18.0	17.9
3	12	25.0	40.0	10.0	17.9	18.0
4	12	25.0	40.0	10.0	18.1	17.8
Group	N		Std Dev of H0 μ 's	Std Dev of H1 μ 's	Mean of H0 σ 's	Mean of H1 σ 's
All	48	100.0	0.0	13.0	18.0	18.0

These reports show the details of each scenario.

Row (in Title)

This is the row number of the Numeric Results report about which this report gives the details.

Group

This is the number of the group shown on this line.

ni

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

Percent of N

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

H0 μ and H1 μ

These are the means that were used in the alpha and power simulations, respectively.

H0 σ and H1 σ

These are the standard deviations that were obtained by the alpha and power simulations, respectively. Note that they often are not exactly equal to what was specified because of the error introduced by simulation.

Std Dev of H0 (and H1) μ 's


These are the standard deviations of the means that were obtained by the alpha and power simulations, respectively. Under H0, this value should be near zero. The H0 value lets you determine if your alpha simulation was correctly specified. The H1 value represents the magnitude of the effect size (when divided by an appropriate measure of the standard deviation).

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Mean of H0 (and H1) σ 's

These are the average of the individual group standard deviations that were obtained by the alpha and power simulations, respectively.

Plots Section



This plot gives a visual presentation to the results in the Numeric Report. We can quickly see the impact on the power of increasing the sample size. We can also see that the F-test is much more power for sample sizes below 8.

Example 2 – Validation using Fleiss

Fleiss (1986) page 374 presents an example of determining an appropriate sample size when using an F-test 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. The validation runs for the three nonparametric tests found a sample size of 12 per group. Therefore, we generate the results for both 11 and 12 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 **Power Comparison of Tests of Means in One-Way Designs (Simulation)** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then clicking on **ANOVA F-Test**, and then clicking on **Power Comparison of Tests of Means in One-Way Designs (Simulation)**. 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**.

<u>Option</u>	<u>Value</u>
Design Tab	
Alpha	0.05
Simulations	5000
Number of Groups	4
Group Allocation	Equal
n (Group Size)	11 12
Group 1 Distribution	Normal(M1 SD)
Group 2 Distribution	Normal(M2 SD)
Group 3 Distribution	Normal(M3 SD)
Group 4 Distribution	Normal(M4 SD)
M1	9.775
M2	12
M3	12
M4	14.225
SD	3
Specify Alpha Distributions	All equal to the group 1 distribution of the power simulation

Output

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

Numeric Results

Power Comparison of Location Tests													
Row	Total Sample Size	AOV F Test Power	Kruskal Wallis Test Power	Terry Hoeffding Test Power	Van der Waer'n Test Power	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	M3	M4	SD
1	44	0.803	0.757	0.766	0.767	0.050	1.6	3.0	9.8	12.0	12.0	14.2	3.0
2	48	0.851	0.815	0.830	0.828	0.050	1.6	3.0	9.8	12.0	12.0	14.2	3.0

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Alpha Comparison of Location Tests													
Row	Total Sample Size N	AOV F Test Alpha	Kruskal Wallis Test Alpha	Terry Hoeffding Test Alpha	Vander Waer'n Test Alpha	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	M3	M4	SD
1	44	0.053	0.046	0.047	0.045	0.050	0.0	3.0	9.8	12.0	12.0	14.2	3.0
2	48	0.047	0.041	0.044	0.043	0.050	0.0	3.0	9.8	12.0	12.0	14.2	3.0

Note that this procedure has also found the power of the F-test to be 0.8 for $n = 11$ and the power of the three nonparametric tests to be about 0.8 for $n = 12$.

Example 3 – Selecting a Test Statistic when the Data Contain Outliers

The F-test is known to be robust to the violation of some assumptions, but it is susceptible to inaccuracy when the data contain outliers. This example will compare the impact of outliers on the power and alpha of the F-test and the three nonparametric tests.

A mixture of two normal distributions will be used to randomly generate outliers. The mixture will draw 95% of the data from a normal distribution with mean 0 and variance 1. The other 5% of the data will come from a normal distribution with mean 0 and variance that ranges from 1 to 10. In the alternative distributions, two will have a mean of 0 and one will have a mean of 1.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Power Comparison of Tests of Means in One-Way Designs (Simulation)** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then clicking on **ANOVA F-Test**, and then clicking on **Power Comparison of Tests of Means in One-Way Designs (Simulation)**. 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**.

<u>Option</u>	<u>Value</u>
Design Tab	
Alpha.....	0.05
Simulations.....	5000
Number of Groups.....	3
Group Allocation.....	Equal
n (Group Size).....	10 20
Group 1 Distribution.....	Normal(M1 SD)[95];Normal(M1 A)[5]
Group 2 Distribution.....	Normal(M1 SD)[95];Normal(M1 A)[5]
Group 3 Distribution.....	Normal(M2 SD)[95];Normal(M2 A)[5]
M1.....	0
M2.....	1
SD.....	1
A.....	1 5 10
Specify Alpha Distributions.....	All equal to the group 1 distribution of the power simulation

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Output

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

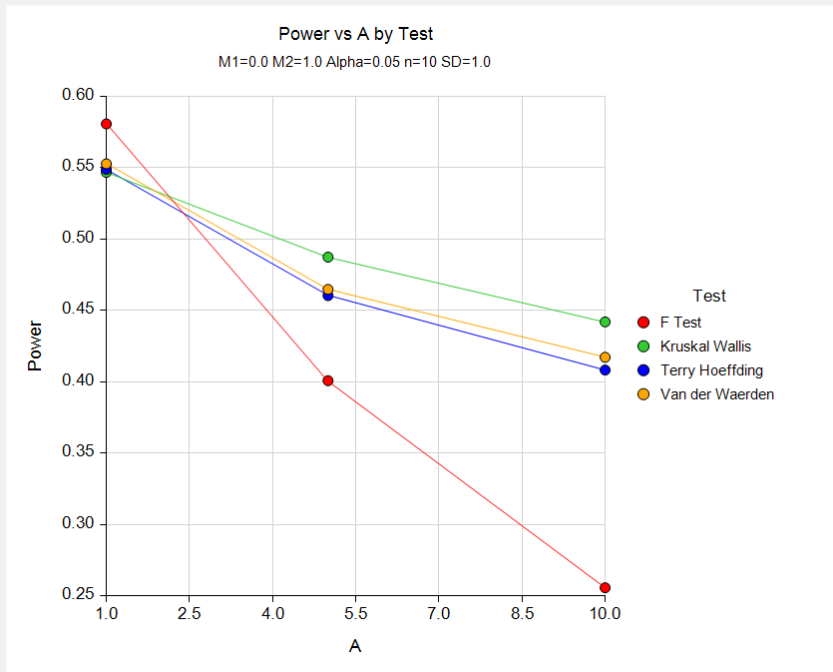
Numeric Results

Power Comparison of Location Tests

Row	Total Sample Size N	AOV F Test Power	Kruskal Wallis Test Power	Terry Hoeffding Test Power	Van der Waer'n Test Power	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	M3	M4	SD
1	30	0.581	0.547	0.549	0.553	0.050	0.5	1.0	0.0	1.0	1.0	1.0	
2	60	0.894	0.872	0.882	0.883	0.050	0.5	1.0	0.0	1.0	1.0	1.0	
3	30	0.401	0.487	0.460	0.465	0.050	0.5	1.5	0.0	1.0	5.0	1.0	
4	60	0.614	0.791	0.765	0.770	0.050	0.5	1.5	0.0	1.0	5.0	1.0	
5	30	0.256	0.442	0.408	0.417	0.050	0.5	2.5	0.0	1.0	10.0	1.0	
6	60	0.420	0.828	0.784	0.793	0.050	0.5	2.4	0.0	1.0	10.0	1.0	

Alpha Comparison of Location Tests

Row	Total Sample Size N	AOV F Test Alpha	Kruskal Wallis Test Alpha	Terry Hoeffding Test Alpha	Van der Waer'n Test Alpha	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	A	SD
1	30	0.051	0.046	0.046	0.047	0.050	0.0	1.0	0.0	1.0	1.0	1.0
2	60	0.049	0.045	0.048	0.047	0.050	0.0	1.0	0.0	1.0	1.0	1.0
3	30	0.042	0.045	0.046	0.045	0.050	0.0	1.5	0.0	1.0	5.0	1.0
4	60	0.037	0.049	0.045	0.045	0.050	0.0	1.5	0.0	1.0	5.0	1.0
5	30	0.028	0.052	0.050	0.051	0.050	0.0	2.5	0.0	1.0	10.0	1.0
6	60	0.026	0.053	0.052	0.052	0.050	0.0	2.4	0.0	1.0	10.0	1.0



We note that when the variances are equal ($A = 1$), the F-Test is slightly better than the nonparametric tests. However, as the number of outliers is increased, the F-test does increasingly worse than the nonparametric tests, both in terms of power and significance.

Example 4 – Selecting a Test Statistic when the Data are Skewed

The F-test is known to be robust to the violation of some assumptions, but it is susceptible to inaccuracy when the underlying distributions are skewed. This example will investigate the impact of skewness on the power and alpha of the F-test and the three nonparametric tests.

Tukey's G-H distribution will be used because it allows the amount of skewness to be gradually increased.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Power Comparison of Tests of Means in One-Way Designs (Simulation)** procedure window by expanding **Means**, then **One-Way Designs (ANOVA)**, then clicking on **ANOVA F-Test**, and then clicking on **Power Comparison of Tests of Means in One-Way Designs (Simulation)**. 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**.

<u>Option</u>	<u>Value</u>
Design Tab	
Alpha.....	0.05
Simulations.....	5000
Number of Groups.....	3
Group Allocation.....	Equal
n (Group Size).....	10 20
Group 1 Distribution.....	TukeyGH(M1 SD G 0)
Group 2 Distribution.....	TukeyGH(M1 SD G 0)
Group 3 Distribution.....	TukeyGH(M2 SD G 0)
M1.....	0
M2.....	1
SD.....	1
G.....	0 0.5 0.9
Specify Alpha Distributions.....	All equal to the group 1 distribution of the power simulation

Output

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

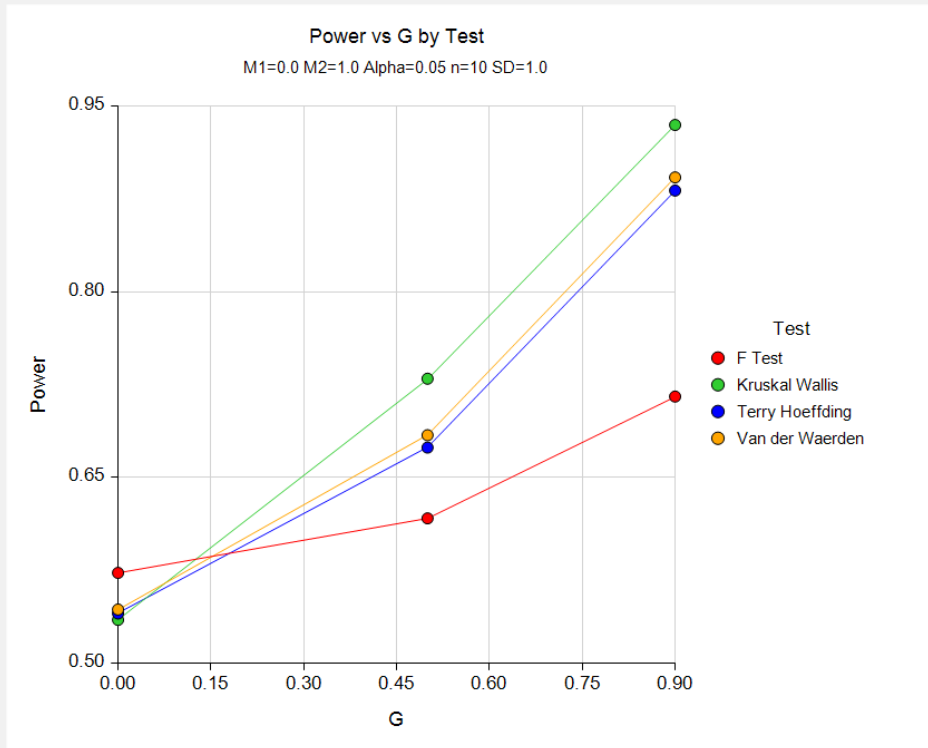
Numeric Results

Power Comparison of Location Tests												
Row	Total Sample Size N	AOV F Test Power	Kruskal Wallis Test Power	Terry Hoeffding Test Power	Van der Waer'n Test Power	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	G	SD
1	30	0.573	0.535	0.540	0.543	0.050	0.5	1.0	0.0	1.0	0.0	1.0
2	60	0.895	0.870	0.885	0.885	0.050	0.5	1.0	0.0	1.0	0.0	1.0
3	30	0.617	0.730	0.674	0.684	0.050	0.5	1.0	0.0	1.0	0.5	1.0
4	60	0.901	0.982	0.971	0.972	0.050	0.5	1.0	0.0	1.0	0.5	1.0
5	30	0.715	0.935	0.882	0.892	0.050	0.5	1.0	0.0	1.0	0.9	1.0
6	60	0.897	0.999	0.998	0.998	0.050	0.5	1.0	0.0	1.0	0.9	1.0

Power Comparison of Tests of Means in One-Way Designs (Simulation)

Alpha Comparison of Location Tests

Row	Total Sample Size N	AOV F Test Alpha	Kruskal Wallis Test Alpha	Terry Hoeffding Test Alpha	Van der Waer'n Test Alpha	Target Alpha	Std Dev of H1 μ 's	Mean of H1 σ 's	M1	M2	G	SD
1	30	0.051	0.043	0.044	0.044	0.050	0.0	1.0	0.0	1.0	0.0	1.0
2	60	0.049	0.051	0.048	0.049	0.050	0.0	1.0	0.0	1.0	0.0	1.0
3	30	0.045	0.046	0.048	0.047	0.050	0.0	1.0	0.0	1.0	0.5	1.0
4	60	0.049	0.052	0.047	0.048	0.050	0.0	1.0	0.0	1.0	0.5	1.0
5	30	0.034	0.048	0.044	0.043	0.050	0.0	1.0	0.0	1.0	0.9	1.0
6	60	0.042	0.045	0.049	0.049	0.050	0.0	1.0	0.0	1.0	0.9	1.0



We note that as the skewness increases, the power of the nonparametric tests increase substantially as compared to the F-test. The Kruskal Wallis test is always the best.