

Chapter 806

Power Comparison of Correlation Tests (Simulation)

Introduction

This procedure analyzes the power and significance level of tests which may be used to test statistical hypotheses about the correlation between two variables. For each scenario, two simulations are run. One simulation estimates the significance level and the other estimates the power.

The three tests that are compared in this procedure are

1. Pearson's Correlation Test
2. Spearman's Rank Correlation Test
3. Kendall's Tau Correlation 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 the test procedure and the test statistic. This includes the significance level, sample size, and underlying data distributions.
2. Generate a random sample of points (X, Y) from the bivariate distribution specified by the alternative hypothesis. Calculate the test statistic from the simulated data and determine if the null hypothesis is accepted or rejected. These samples are used to calculate the power of the test. In the case of paired data, the individual values are simulated are constructed so that they exhibit the specified amount of correlation.
3. Generate a second random sample of points (X, Y) from the bivariate distribution specified by the null hypothesis. Calculate the test statistic from the simulated data and determine if the null hypothesis is accepted or rejected. These samples are used to calculate the significance-level of the test. In the case of paired data, the individual values are simulated so that they exhibit the specified amount of correlation.
4. 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.

Simulating Paired Distributions

In this routine, paired data may be generated from the bivariate normal distribution or from two specified marginal distributions. In the latter case, the simulation should mimic the actual data generation process as closely as possible, including the marginal distributions and the correlation between the two variables.

Obtaining paired samples from arbitrary distributions with a set correlation is difficult because the joint, bivariate distribution must be specified and simulated. Rather than specify the bivariate distribution, *PASS* requires the specification of the two marginal distributions and the correlation between them.

Monte Carlo samples with given marginal distributions and correlation are generated using the method suggested by Gentle (1998). The method begins by generating a large population of random numbers from the two distributions. Each of these populations is evaluated to determine if their means are 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.

The next step is to obtain the target correlation. This is accomplished by permuting one of the populations until they have the desired correlation.

The above steps provide a large pool (ten thousand to one million items) of random number pairs that exhibit the desired characteristics. This pool is then sampled at random using the uniform distribution to obtain the random number pairs used in the simulation.

This algorithm may be stated as follows.

1. Draw individual samples of size M from the two distributions where M is a large number over 10,000. Adjust these samples so that they have the specified mean and standard deviation. Label these samples A and B . Create an index of the values of A and B according to the order in which they are generated. Thus, the first value of A and the first value of B are indexed as one, the second values of A and B are indexed as two, and so on up to the final set which is indexed as M .
2. Compute the product-moment correlation between the two generated variates.
3. If the computed product-moment correlation is within a small tolerance (usually less than 0.001) of the specified correlation, go to step 7.
4. Select two indices (I and J) at random using uniform random numbers.
5. Determine what will happen to the product-moment correlation if B_I is swapped with B_J . If the swap will result in a correlation that is closer to the target value, swap the indices and proceed to step 6. Otherwise, go to step 4.
6. If the computed product-moment correlation is within the desired tolerance of the target correlation, go to step 7. Otherwise, go to step 4.
7. End with a population with the required marginal distributions and correlation.

A population created by this procedure tends to exhibit more variation in the tails of the distribution than in the center. Hence, the results using the bivariate normal option will be slightly different from those obtained when a custom model with two normal distributions is used.

Tests

The technical details of the three 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.

Hypothesis and Simulations

Alternative Hypothesis

This option specifies the alternative hypothesis. This implicitly specifies the direction of the hypothesis test. The null hypothesis is that the correlation is zero.

Possible selections are:

- $\rho_1 \neq \rho_0$ or $\rho_s \neq 0$ or $\tau \neq 0$
This is the most common selection. It yields the *two-tailed* test. Use this option when you are testing whether the correlation values are different, but you do not want to specify beforehand which correlation is larger.
- $\rho_1 > \rho_0$ or $\rho_s > 0$ or $\tau > 0$
This option yields a *one-tailed* test.
- $\rho_1 < \rho_0$ or $\rho_s < 0$ or $\tau < 0$
This option also yields a *one-tailed* test.

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. 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.

Power Comparison of Correlation Tests (Simulation)

Alpha

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*.

Sample Size

N (Sample Size)

The number of observations in the sample. Each observation is made up of two values: one for X and one for Y .

Correlations

ρ_1 (Correlation|H1)

Specify the value of ρ_1 , the population product-moment correlation under the alternative hypothesis. Note that the range of the correlation is between plus and minus one. The difference between 0 and ρ_1 is being tested by this significance test.

You can enter a range of values separated by blanks or commas.

Note that this is the value of the Pearson product-moment correlation, not Spearman's rank or Kendall's tau correlation.

Power and Alpha Simulations

Distribution

Specify whether a bivariate normal or a custom bivariate distribution is used to generate the data. The possible choices are:

- **Bivariate Normal**

Bivariate normal distributions are simulated with zero means, unit variances, and correlations ρ_0 and ρ_1 to create the alpha and power simulations, respectively.

- **Custom**

A bivariate distribution is specified by setting the two marginal distributions and their correlation (ρ_0 for the alpha simulation and ρ_1 for the power simulation).

The algorithm proceeds by generating a pool of several thousand random pairs from the X and Y distributions. To achieve the desired correlation, pairs of points are randomly selected and their X values are swapped if the switch results in an overall correlation closer to the desired value.

When applied to normal data, this algorithm results in a bivariate population that is narrower in the center and wider in the tails than the standard bivariate normal.

Power Comparison of Correlation Tests (Simulation)

Custom Power and Alpha Simulations

These options specify the distributions of X and Y to be used in the power simulation (on the left) and the alpha simulation (on the right). All distributions are specified with the same syntax.

Distribution

Specify the distribution of each variable under the alternative hypothesis, H_1 , on the left and under the null hypothesis, H_0 , on the right.

For convenience in specifying a range of values, the parameters of a 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.

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)

Power Comparison of Correlation Tests (Simulation)

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

Simulations Tab

The options 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 Numbers

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.

Correlation

Maximum Switches

This option specifies the maximum number of index switches that can be made while searching for a permutation of Y that yields a correlation within the specified tolerance.

A value near 5000000 is needed when the correlation is large.

Correlation Tolerance

Specify the amount above and below the target correlation that will still let a particular index-permutation to be selected as the bivariate population.

For example, if you have selected a correlation of 0.3 and you set this tolerance to 0.001, then permutations will continue until one is found which yields a correlation between 0.299 and 0.301.

Recommended: 0.001

Valid Values: 0 to 0.99.

Power Comparison of Correlation Tests (Simulation)

Example 1 – Power at Various Sample Sizes

Suppose a study will be run to test whether the correlation between forced vital capacity (X) and forced expiratory value (Y) in a particular population is non-zero. Find the power if $\rho_1 = 0.20$ and 0.30 , alpha is 0.05 , and $N = 20, 60, 100$.

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 Correlation Tests (Simulation)** procedure window by expanding **Correlation**, then **Correlation**, then clicking on **Test (Inequality)**, and then clicking on **Power Comparison of Correlation Tests (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	
Alternative Hypothesis	$\rho_1 \neq \rho_0$ or $\rho_s \neq 0$ or $\tau \neq 0$
Alpha.....	0.05
N (Sample Size).....	20 60 100
ρ_1 (Correlation H0)	0.2 0.3
Distribution	Bivariate Normal

Annotated Output

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

Numeric Results Report

Simulation Summary					
Variables	Power (H1) Simulation Distribution	Alpha (H0) Simulation Distribution			
X and Y	Bivariate Normal	Bivariate Normal			
Random Number Pool Size	10000				
Number of Simulations	5000				
Power Comparison for Testing Correlation Hypotheses: H1: $\rho_0 \neq \rho_1$ or $\rho_s \neq 0$ or $\tau \neq 0$					
N	H1 Corr ρ_1	Target Alpha	Pearson Corr Power	Spearman Rho Power	Kendall Tau b Power
20	0.200	0.050	0.151	0.132	0.122
60	0.200	0.050	0.359	0.344	0.336
100	0.200	0.050	0.556	0.492	0.491
20	0.300	0.050	0.247	0.222	0.210
60	0.300	0.050	0.625	0.576	0.569
100	0.300	0.050	0.799	0.756	0.751
Run Time: 8.62 seconds.					

Power Comparison of Correlation Tests (Simulation)

Alpha Comparison for Testing Correlation Hypotheses: $H_1: \rho_0 \neq \rho_1$ or $\rho_s \neq 0$ or $\tau \neq 0$

N	H1 Corr ρ_1	Target Alpha	Pearson Corr Alpha	Spearman Rho Alpha	Kendall Tau b Alpha
20	0.200	0.050	0.057	0.051	0.048
60	0.200	0.050	0.048	0.049	0.047
100	0.200	0.050	0.051	0.053	0.051
20	0.300	0.050	0.050	0.052	0.047
60	0.300	0.050	0.052	0.056	0.055
100	0.300	0.050	0.050	0.055	0.052

Definitions of the Power and Alpha Comparison Reports

N is the size of the sample drawn from the population. It is the number of X-Y data points in a sample.

ρ_0 is the Pearson correlation coefficient assuming the null hypothesis, H_0 . It is set to zero which results in a test of non-correlation between X and Y.

ρ_1 is the Pearson correlation coefficient assuming the alternative hypothesis, H_1 . This is the value at which the power is computed.

Target Alpha is the desired probability of rejecting a true null hypothesis at which the tests were run.

Power is the probability of rejecting H_0 when it is false. This is the actual value calculated by the power simulation.

Alpha is the alpha achieved by the test as calculated by the alpha simulation.

Pearson Corr Power (Alpha) is the power (alpha) of the test whether Pearson product-moment correlation coefficient is zero.

Spearman Rho Power (Alpha) is the power (alpha) of the test whether Spearman's Rho is zero.

Kendall Tau b Power (Alpha) is the power (alpha) of the test whether Kendall's Tau b is zero.

References

Graybill, Franklin. 1961. An Introduction to Linear Statistical Models. McGraw-Hill. New York, New York.

Guenther, William C. 1977. 'Desk Calculation of Probabilities for the Distribution of the Sample Correlation Coefficient', The American Statistician, Volume 31, Number 1, pages 45-48.

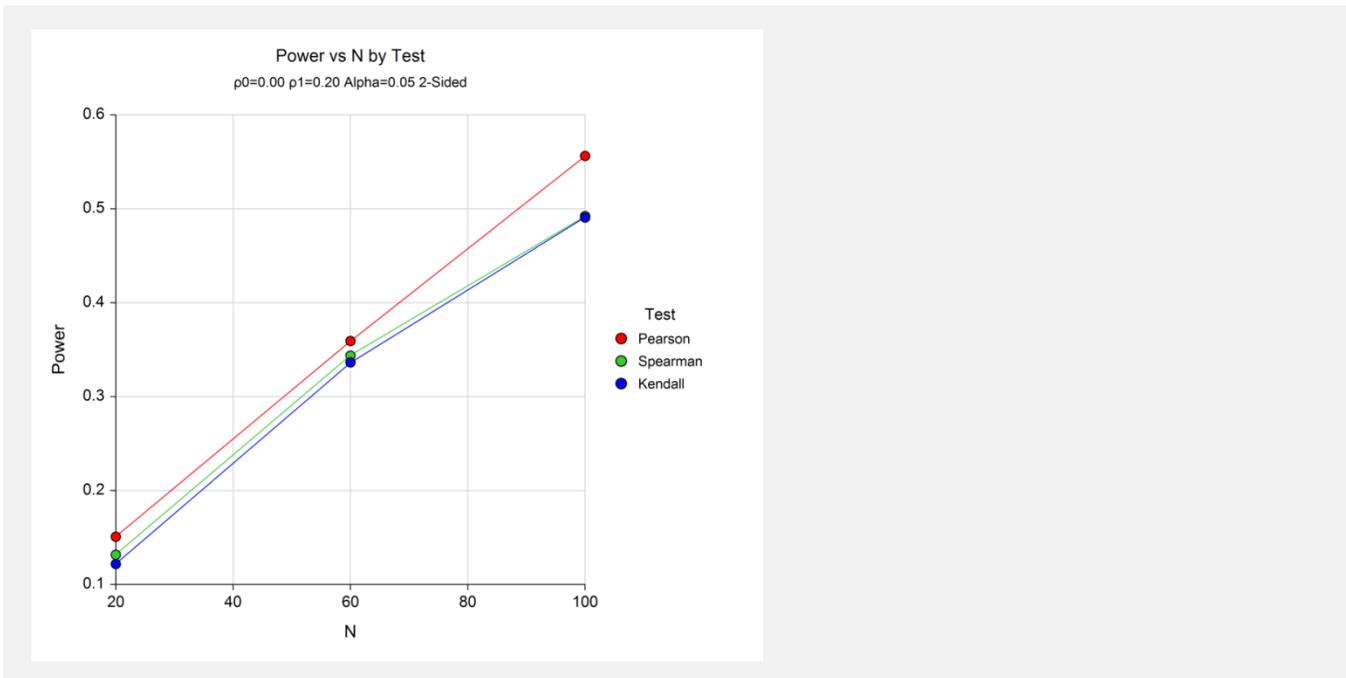
Kendall, M. and Gibbons, J.D. 1990. Rank Correlation Methods, 5th Edition. Oxford University Press. New York.

Zar, Jerrold H. 1984. Biostatistical Analysis. Second Edition. Prentice-Hall. Englewood Cliffs, New Jersey.

Devroye, Luc. 1986. Non-Uniform Random Variate Generation. Springer-Verlag. New York.

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

Plots Section



This plot gives a visual presentation of the results in the Numeric Report. We can see the impact on the power of increasing the sample size.

Power Comparison of Correlation Tests (Simulation)

Example 2 – Validation using Zar

Zar (1984) page 312 presents an example in which the power of a Pearson's correlation coefficient is calculated. If $N = 12$, $\alpha = 0.05$, $\rho_0 = 0$, and $\rho_1 = 0.866$, Zar calculates a power of 98% for a two-sided test. We know that the power of Spearman's rho and Kendall's tau will be a little less than Pearson's correlation.

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 Correlation Tests (Simulation)** procedure window by expanding **Correlation**, then **Correlation**, then clicking on **Test (Inequality)**, and then clicking on **Power Comparison of Correlation Tests (Simulation)**. You may then make the appropriate entries as listed below, or open **Example 2** by going to the **File** menu and choosing **Open Example Template**.

<u>Option</u>	<u>Value</u>
Design Tab	
Alternative Hypothesis	$\rho_1 \neq \rho_0$ or $\rho_s \neq 0$ or $\tau \neq 0$
Simulations	50000
Alpha	0.05
N (Sample Size)	12
ρ_1 (Correlation H0)	0.866
Distribution	Bivariate Normal

Output

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

Numeric Results

Power Comparison for Testing Correlation Hypotheses: H1: $\rho_0 \neq \rho_1$ or $\rho_s \neq 0$ or $\tau \neq 0$					
N	H1 Corr ρ_1	Target Alpha	Pearson Corr Power	Spearman Rho Power	Kendall Tau b Power
12	0.866	0.050	0.984	0.951	0.944
Run Time: 4.31 seconds.					
Alpha Comparison for Testing Correlation Hypotheses: H1: $\rho_0 \neq \rho_1$ or $\rho_s \neq 0$ or $\tau \neq 0$					
N	H1 Corr ρ_1	Target Alpha	Pearson Corr Alpha	Spearman Rho Alpha	Kendall Tau b Alpha
12	0.866	0.050	0.050	0.051	0.045

Note that this procedure finds the power of the Pearson test to be 0.98, which matches Zar. The other two tests are a little less than 0.98.