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.

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 ρ_1 = 0.20 and 0.30, alpha is 0.05, and *N* = 20, 60, 100.

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

Alternative Hypothesis	**ρ1 ≠ ρ0 or ρs ≠ 0 or τ ≠ 0**
Simulations	**5000**
Random Seed	**2031884** (for Reproducibility)
Alpha	**0.05**
N (Sample Size)	**20 60 100**
ρ1 (Correlation|H1)	**0.2 0.3**
Distribution	**Bivariate Normal**

Output

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

Numeric Reports

Simulation Summary

	Simulation Distribution	
Variables	**Power (H1)**	**Alpha (H0)**
X and Y	Bivariate Normal	Bivariate Normal

Random Number Pool Size	10000
Number of Simulations	5000
Random Seed	2031884 (User-Entered)
Run Time	3.84 seconds

Power Comparison

Correlation Hypotheses: H1: ρ1 ≠ ρ0 or ρs ≠ 0 or τ ≠ 0

			Power		
N	Correlation | H1 ρ1	Target Alpha	Pearson Correlation	Spearman Rho	Kendall's Tau-b
20	0.2	0.05	0.138	0.127	0.119
60	0.2	0.05	0.306	0.299	0.291
100	0.2	0.05	0.554	0.517	0.515
20	0.3	0.05	0.223	0.201	0.191
60	0.3	0.05	0.684	0.636	0.629
100	0.3	0.05	0.848	0.799	0.796

Alpha Comparison

Correlation Hypotheses: H1: ρ1 ≠ ρ0 or ρs ≠ 0 or τ ≠ 0

		Alpha			
N	Correlation | H1 ρ1	Target	Pearson Correlation	Spearman Rho	Kendall's Tau-b
20	0.2	0.05	0.049	0.048	0.048
60	0.2	0.05	0.048	0.053	0.052
100	0.2	0.05	0.059	0.056	0.054
20	0.3	0.05	0.047	0.046	0.040
60	0.3	0.05	0.048	0.043	0.044
100	0.3	0.05	0.045	0.047	0.046

N	The size of the sample drawn from the population. It is the number of X-Y data points in a sample.
ρ0	The Pearson correlation coefficient assuming the null hypothesis, H0. It is set to zero which results in a test of non-correlation between X and Y.
ρ1	The Pearson correlation coefficient assuming the alternative hypothesis, H1. This is the value at which the power is computed.
Target Alpha	The desired probability of rejecting a true null hypothesis at which the tests were run.
Power	The probability of rejecting H0 when it is false. This is the actual value calculated by the power simulation.
Alpha	The alpha achieved by the test as calculated by the alpha simulation.
Pearson Corr Power and Alpha	The power (alpha) of the test whether Pearson product-moment correlation coefficient is zero.
Spearman Rho Power and Alpha	The power (alpha) of the test whether Spearman's Rho is zero.
Kendall's Tau-b Power and Alpha	The power (alpha) of the test whether Kendall's Tau-b is zero.

Dropout-Inflated Sample Size

Dropout Rate	Sample Size N	Dropout-Inflated Enrollment Sample Size N'	Expected Number of Dropouts D
20%	20	25	5
20%	60	75	15
20%	100	125	25

Dropout Rate The percentage of subjects (or items) that are expected to be lost at random during the course of the study and for whom no response data will be collected (i.e., will be treated as "missing"). Abbreviated as DR.
N The evaluable sample size at which power is computed (as entered by the user). If N subjects are evaluated out of the N' subjects that are enrolled in the study, the design will achieve the stated power.
N' The total number of subjects that should be enrolled in the study in order to obtain N evaluable subjects, based on the assumed dropout rate. N' is calculated by inflating N using the formula N' = N / (1 - DR), with N' always rounded up. (See Julious, S.A. (2010) pages 52-53, or Chow, S.C., Shao, J., Wang, H., and Lokhnygina, Y. (2018) pages 32-33.)
D The expected number of dropouts. D = N' - N.

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

Plots

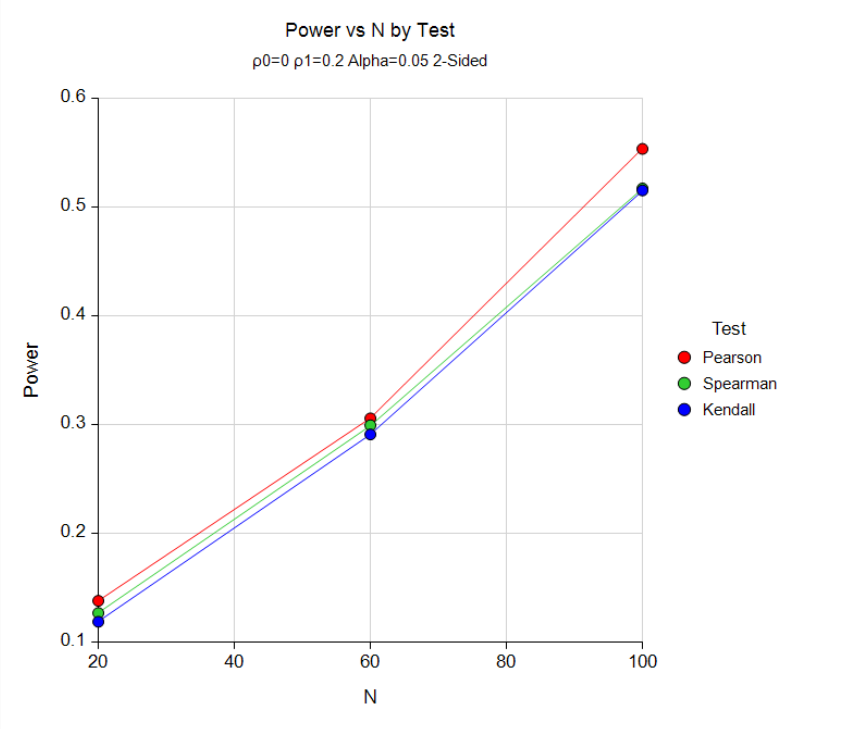

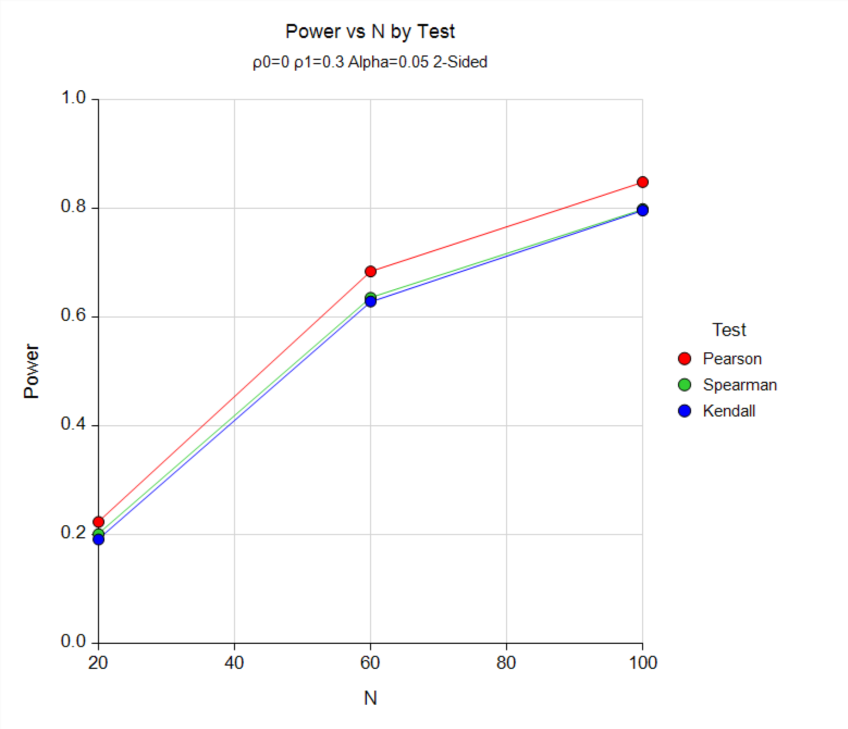

This plot gives a visual presentation of the results in the Numeric Report. We can see the impact on the power of increasing the sample size.

Example 2 – Validation using Zar (1984)

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, ρ_0 = 0, and ρ_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.

For reproducibility, we'll use a random seed of 5857325.

Setup

If the procedure window is not already open, use the PASS Home window to open it. The parameters for this example are listed below and are stored in the **Example 2** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Design Tab

Alternative Hypothesis .... **ρ1 ≠ ρ0 or ρs ≠ 0 or τ ≠ 0**
Simulations **50000**
Random Seed .... **5857325** (for Reproducibility)
Alpha.. **0.05**
N (Sample Size)... **12**
ρ1 (Correlation|H1) **0.866**
Distribution.. **Bivariate Normal**

Output

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

Simulation Summary

	Simulation Distribution	
Variables	**Power (H1)**	**Alpha (H0)**
X and Y	Bivariate Normal	Bivariate Normal

Random Number Pool Size	100000
Number of Simulations	50000
Random Seed	5857325 (User-Entered)
Run Time	1.58 seconds

Power Comparison

Correlation Hypotheses: H1: ρ1 ≠ ρ0 or ρs ≠ 0 or τ ≠ 0

			Power		
N	**Correlation | H1 ρ1**	**Target Alpha**	**Pearson Correlation**	**Spearman Rho**	**Kendall's Tau-b**
12	0.866	0.05	0.984	0.952	0.943

Alpha Comparison

Correlation Hypotheses: H1: ρ1 ≠ ρ0 or ρs ≠ 0 or τ ≠ 0

		Alpha			
N	Correlation | H1 ρ1	Target	Pearson Correlation	Spearman Rho	Kendall's Tau-b
12	0.866	0.05	0.051	0.053	0.045

Note that this procedure finds the power of the Pearson test to be 0.984, which matches Zar. The other two tests are a little less than 0.98.