

power twocorrelations — Power analysis for a two-sample correlations test

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`power twocorrelations` computes sample size, power, or the experimental-group correlation for a two-sample correlations test. By default, it computes sample size for given power and the values of the control-group and experimental-group correlations. Alternatively, it can compute power for given sample size and values of the control-group and experimental-group correlations or the experimental-group correlation for given sample size, power, and the control-group correlation. Also see [\[PSS\] power](#) for a general introduction to the `power` command using hypothesis tests.

Quick start

Sample size for a test of $H_0: \rho_1 = \rho_2$ versus $H_a: \rho_1 \neq \rho_2$ when control-group correlation $r_1 = 0.4$ and experimental-group correlation $r_2 = 0.1$ with default power of 0.8 and significance level $\alpha = 0.05$

```
power twocorrelations .4 .1
```

As above, for r_2 equal to -0.05 , 0 , 0.05 , and 0.1

```
power twocorrelations .4 (-.05(.05).1)
```

As above, but for a one-sided test

```
power twocorrelations .4 (-.05(.05).1), onesided
```

Specify r_1 and difference $r_2 - r_1 = -0.3$

```
power twocorrelations .4, diff(-.3)
```

As above, but specify a sample size of 200 for group 1 and compute sample size for group 2

```
power twocorrelations 0.4, diff(-0.3) n1(200) compute(N2)
```

Power for $r_1 = 0.4$ and $r_2 = -0.15$ with a sample size of 100 and default $\alpha = 0.05$

```
power twocorrelations .4 -.15, n(100)
```

As above, but specify sample size of 50 and 65 for groups 1 and 2, respectively

```
power twocorrelations .4 -.15 , n1(50) n2(65)
```

Power for sample sizes of 60, 70, 80, 90, and 100

```
power twocorrelations .4 -.15, n(60(10)100)
```

As above, but display results in a graph of power versus sample size

```
power twocorrelations .4 -.15, n(60(10)100) graph
```

Effect size and experimental-group correlation for $r_1 = 0.4$ with sample size of 150 and power of 0.85

```
power twocorrelations .4, n(150) power(.85)
```

Menu

Statistics > Power and sample size

Syntax

Compute sample size

```
power twocorrelations  $r_1$   $r_2$  [ , power(numlist) options ]
```

Compute power

```
power twocorrelations  $r_1$   $r_2$  , n(numlist) [options ]
```

Compute effect size and experimental-group correlation

```
power twocorrelations  $r_1$  , n(numlist) power(numlist) [options ]
```

where r_1 is the correlation in the control (reference) group and r_2 is the correlation in the experimental (comparison) group. r_1 and r_2 may each be specified either as one number or as a list of values in parentheses (see [U] [11.1.8 numlist](#)).

<i>options</i>	Description
Main	
* <u>alpha</u> (<i>numlist</i>)	significance level; default is <code>alpha(0.05)</code>
* <u>power</u> (<i>numlist</i>)	power; default is <code>power(0.8)</code>
* <u>beta</u> (<i>numlist</i>)	probability of type II error; default is <code>beta(0.2)</code>
* <u>n</u> (<i>numlist</i>)	total sample size; required to compute power or effect size
* <u>n1</u> (<i>numlist</i>)	sample size of the control group
* <u>n2</u> (<i>numlist</i>)	sample size of the experimental group
* <u>nratio</u> (<i>numlist</i>)	ratio of sample sizes, $N2/N1$; default is <code>nratio(1)</code> , meaning equal group sizes
<u>compute</u> ($N1 N2$)	solve for $N1$ given $N2$ or for $N2$ given $N1$
<u>nfractional</u>	allow fractional sample sizes
* <u>diff</u> (<i>numlist</i>)	difference between the experimental-group and control-group correlations, $r_2 - r_1$; specify instead of the experimental-group correlation r_2
<u>direction</u> (<u>upper</u> <u>lower</u>)	direction of the effect for effect-size determination; default is <code>direction(upper)</code> , which means that the postulated value of the parameter is larger than the hypothesized value
<u>onesided</u>	one-sided test; default is two sided
<u>parallel</u>	treat number lists in starred options or in command arguments as parallel when multiple values per option or argument are specified (do not enumerate all possible combinations of values)
Table	
[<u>no</u>] <u>table</u> [(<i>tablespec</i>)]	suppress table or display results as a table; see [PSS] power, table
<u>saving</u> (<i>filename</i> [, <u>replace</u>])	save the table data to <i>filename</i> ; use <code>replace</code> to overwrite existing <i>filename</i>
Graph	
<u>graph</u> [(<i>graphopts</i>)]	graph results; see [PSS] power, graph
Iteration	
<u>init</u> (#)	initial value for sample sizes or experimental-group correlation
<u>iterate</u> (#)	maximum number of iterations; default is <code>iterate(500)</code>
<u>tolerance</u> (#)	parameter tolerance; default is <code>tolerance(1e-12)</code>
<u>ftolerance</u> (#)	function tolerance; default is <code>ftolerance(1e-12)</code>
[<u>no</u>] <u>log</u>	suppress or display iteration log
[<u>no</u>] <u>dots</u>	suppress or display iterations as dots
<u>notitle</u>	suppress the title

*Specifying a list of values in at least two starred options, or at least two command arguments, or at least one starred option and one argument results in computations for all possible combinations of the values; see [U] [11.1.8 numlist](#). Also see the `parallel` option.

`notitle` does not appear in the dialog box.

where *tablespec* is

```
column[:label] [column[:label] [...]] [, tableopts]
```

column is one of the columns defined [below](#), and *label* is a column label (may contain quotes and compound quotes).

<i>column</i>	Description	Symbol
<code>alpha</code>	significance level	α
<code>power</code>	power	$1 - \beta$
<code>beta</code>	type II error probability	β
<code>N</code>	total number of subjects	N
<code>N1</code>	number of subjects in the control group	N_1
<code>N2</code>	number of subjects in the experimental group	N_2
<code>nratio</code>	ratio of sample sizes, experimental to control	N_2/N_1
<code>delta</code>	effect size	δ
<code>r1</code>	control-group correlation	ρ_1
<code>r2</code>	experimental-group correlation	ρ_2
<code>diff</code>	difference between the experimental-group correlation and the control-group correlation	$\rho_2 - \rho_1$
<code>target</code>	target parameter; synonym for <code>r2</code>	
<code>_all</code>	display all supported columns	

Column `beta` is shown in the default table in place of column `power` if specified.

Columns `diff` and `nratio` are shown in the default table if specified.

Options

Main

`alpha()`, `power()`, `beta()`, `n()`, `n1()`, `n2()`, `nratio()`, `compute()`, `nfractional`; see [\[PSS\] power](#).

`diff(numlist)` specifies the difference between the experimental-group correlation and the control-group correlation, $r_2 - r_1$. You can specify either the experimental-group correlation r_2 as a command argument or the difference between the two correlations in `diff()`. If you specify `diff(#)`, the experimental-group correlation is computed as $r_2 = r_1 + \#$. This option is not allowed with the effect-size determination.

`direction()`, `onesided`, `parallel`; see [\[PSS\] power](#).

Table

`table`, `table()`, `notable`; see [\[PSS\] power, table](#).

`saving()`; see [\[PSS\] power](#).

Graph

`graph`, `graph()`; see [\[PSS\] power, graph](#). Also see the *column* table for a list of symbols used by the graphs.

Iteration

`init(#)` specifies the initial value for the estimated parameter. For sample-size determination, the estimated parameter is either the control-group size n_1 or, if `compute(N2)` is specified, the experimental-group size n_2 . For the effect-size determination, the estimated parameter is the experimental-group correlation r_2 . The default initial values for a two-sided test are obtained as a closed-form solution for the corresponding one-sided test with the significance level $\alpha/2$.

`iterate()`, `tolerance()`, `ftolerance()`, `log`, `nolog`, `dots`, `nodots`; see [PSS] [power](#).

The following option is available with `power twocorrelations` but is not shown in the dialog box: `notitle`; see [PSS] [power](#).

Remarks and examples

[stata.com](http://www.stata.com)

Remarks are presented under the following headings:

Introduction

Using power twocorrelations

Computing sample size

Computing power

Computing effect size and experimental-group correlation

Testing a hypothesis about two independent correlations

This entry describes the `power twocorrelations` command and the methodology for power and sample-size analysis for a two-sample correlations test. See [PSS] [intro](#) for a general introduction to power and sample-size analysis and [PSS] [power](#) for a general introduction to the `power` command using hypothesis tests.

Introduction

There are many examples of studies where a researcher may want to compare two correlations. A winemaker may want to test the null hypothesis that the correlation between fermentation time and alcohol level is the same for pinot noir grapes and merlot grapes. An education researcher may want to test the null hypothesis that the correlation of verbal and math SAT scores is the same for males and females. Or a genetics researcher may want to test the null hypothesis that the correlation of the cholesterol levels in identical twins raised together is equal to the correlation of the cholesterol levels in identical twins raised apart.

Correlation analysis emanates from studies quantifying dependence between random variables. The correlation coefficient ρ is commonly used to measure the strength of such dependence. We consider Pearson's correlation obtained by standardizing the covariance between two random variables. As a result, the correlation coefficient has no units and ranges between -1 and 1 .

This entry describes power and sample-size analysis for the inference about two population correlation coefficients performed using hypothesis testing. Specifically, we consider the null hypothesis $H_0: \rho_2 = \rho_1$ versus the two-sided alternative hypothesis $H_a: \rho_2 \neq \rho_1$, the upper one-sided alternative $H_a: \rho_2 > \rho_1$, or the lower one-sided alternative $H_a: \rho_2 < \rho_1$.

One common approach in testing hypotheses concerning the correlation parameter is to use an inverse hyperbolic tangent transformation, $\tanh^{-1}(x) = 0.5 \ln(1+x)/\ln(1-x)$, also known as Fisher's z transformation when applied to the correlation coefficient (Fisher 1915). Specifically, if $\hat{\rho}$ is the sample correlation coefficient and n is the sample size, Fisher (1915) showed that

$$\sqrt{n-3} \{ \tanh^{-1}(\hat{\rho}) - \tanh^{-1}(\rho) \} \sim N(0, 1)$$

for n as small as 10, although the approximation tends to perform better for $n > 25$. For a two-sample correlations test, the null hypothesis $H_0: \rho_1 = \rho_2$ is equivalent to $H_0: \tanh^{-1}(\rho_1) = \tanh^{-1}(\rho_2)$. The latter test is referred to as the two-sample Fisher's z test.

`power twocorrelations` performs computations based on the asymptotic two-sample Fisher's z test.

Using power twocorrelations

`power twocorrelations` computes sample size, power, or experimental-group correlation for a two-sample correlations test. All computations are performed for a two-sided hypothesis test where, by default, the significance level is set to 0.05. You may change the significance level by specifying the `alpha()` option. You can specify the `onesided` option to request a one-sided test. By default, all computations assume a balanced or equal-allocation design; see [PSS] **unbalanced designs** for a description of how to specify an unbalanced design.

To compute the total sample size, you must specify the control- and experimental-group correlations, r_1 and r_2 , respectively, and, optionally, the power of the test in the `power()` option. The default power is set to 0.8.

Instead of the total sample size, you can compute one of the group sizes given the other one. To compute the control-group sample size, you must specify the `compute(N1)` option and the sample size of the experimental group in the `n2()` option. Likewise, to compute the experimental-group sample size, you must specify the `compute(N2)` option and the sample size of the control group in the `n1()` option.

To compute power, you must specify the total sample size in the `n()` option and the control- and experimental-group correlations, r_1 and r_2 , respectively.

Instead of the experimental-group correlation r_2 , you may specify the difference $r_2 - r_1$ between the experimental-group correlation and the control-group correlation in the `diff()` option when computing sample size or power.

To compute effect size, the difference between the experimental-group and the control-group correlation, and the experimental-group correlation, you must specify the total sample size in the `n()` option, the power in the `power()` option, the control-group correlation r_1 , and optionally, the direction of the effect. The direction is upper by default, `direction(upper)`, which means that the experimental-group correlation is assumed to be larger than the specified control-group value. You can change the direction to be lower, which means that the experimental-group correlation is assumed to be smaller than the specified control-group value, by specifying the `direction(lower)` option.

Instead of the total sample size `n()`, you can specify individual group sizes in `n1()` and `n2()`, or specify one of the group sizes and `nratio()` when computing power or effect size. Also see *Two samples* in [PSS] **unbalanced designs** for more details.

In the following sections, we describe the use of `power twocorrelations` accompanied with examples for computing sample size, power, and experimental-group correlation.

Computing sample size

To compute sample size, you must specify the control- and experimental-group correlations, r_1 and r_2 , respectively, and, optionally, the power of the test in the `power()` option. A default power of 0.8 is assumed if `power()` is not specified.

▷ Example 1: Sample size for a two-sample correlations test

Consider a study in which investigators are interested in testing whether the correlation between height and weight differs for males and females. The null hypothesis $H_0: \rho_F = \rho_M$ is tested against the alternative hypothesis $H_a: \rho_F \neq \rho_M$, where ρ_F is the correlation between height and weight for females, and ρ_M is the correlation between height and weight for males.

Before conducting the study, investigators wish to determine the minimum sample size required to detect a difference between the correlation of 0.3 for females and a correlation of 0.5 for males with 80% power using a two-sided 5%-level test. We specify the values 0.3 and 0.5 as the control- and experimental-group correlations after the command name. We omit options `alpha(0.05)` and `power(0.8)` because the specified values are their defaults. To compute the total sample size assuming equal-group allocation, we type

```
. power twocorrelations 0.3 0.5
Performing iteration ...
Estimated sample sizes for a two-sample correlations test
Fisher's z test
Ho: r2 = r1 versus Ha: r2 != r1
Study parameters:
    alpha =    0.0500
    power =    0.8000
    delta =    0.2000
    r1 =     0.3000
    r2 =     0.5000
Estimated sample sizes:
    N =        554
    N per group =    277
```

A total sample of 554 individuals, 277 in each group, must be obtained to detect the difference between correlations of females and males when the correlations are 0.3 and 0.5, respectively, with 80% power using a two-sided 5%-level test.

◀

▷ Example 2: Computing one of the group sizes

Continuing with [example 1](#), we will suppose that for some reason, we can enroll only 250 male subjects in our study. We want to know how many female subjects we need to recruit to maintain the 80% power for detecting the difference as described in example 1. To do so, we use the combination of `compute(N1)` and `n2()`:

```

. power twocorrelations 0.3 0.5, n2(250) compute(N1)
Performing iteration ...
Estimated sample sizes for a two-sample correlations test
Fisher's z test
Ho: r2 = r1 versus Ha: r2 != r1
Study parameters:
      alpha =    0.0500
      power =    0.8000
      delta =    0.2000
      r1 =     0.3000
      r2 =     0.5000
      N2 =      250
Estimated sample sizes:
      N =      559
      N1 =     309

```

We need 309 females for a total sample size of 559 subjects, which is larger than the required total sample size for the corresponding balanced design from [example 1](#).



▷ Example 3: Unbalanced design

By default, `power twocorrelations` computes sample size for a balanced or equal-allocation design. If we know the allocation ratio of subjects between the groups, we can compute the required sample size for an unbalanced design by specifying the `nratio()` option.

Continuing with [example 1](#), we will suppose that we anticipate to recruit twice as many males as females; that is, $n_2/n_1 = 2$. We specify the `nratio(2)` option to compute the required sample size for the specified unbalanced design.

```

. power twocorrelations 0.3 0.5, nratio(2)
Performing iteration ...
Estimated sample sizes for a two-sample correlations test
Fisher's z test
Ho: r2 = r1 versus Ha: r2 != r1
Study parameters:
      alpha =    0.0500
      power =    0.8000
      delta =    0.2000
      r1 =     0.3000
      r2 =     0.5000
      N2/N1 =    2.0000
Estimated sample sizes:
      N =      624
      N1 =     208
      N2 =     416

```

Also see [Two samples in \[PSS\] unbalanced designs](#) for more examples of unbalanced designs for two-sample tests.



Computing power

To compute power, you must specify the sample size in the `n()` option and the control- and the experimental-group correlations, r_1 and r_2 , respectively.

▷ Example 4: Power of a two-sample correlations test

Continuing with [example 1](#), we will suppose that we are designing a study and anticipate a total sample of 500 subjects. To compute the power corresponding to this sample size given the study parameters from [example 1](#), we specify the sample size in `n()`:

```
. power twocorrelations 0.3 0.5, n(500)
Estimated power for a two-sample correlations test
Fisher's z test
Ho: r2 = r1 versus Ha: r2 != r1
Study parameters:
      alpha =    0.0500
         N =     500
  N per group =   250
      delta =    0.2000
         r1 =    0.3000
         r2 =    0.5000
Estimated power:
      power =    0.7595
```

With a smaller sample of 500 subjects compared with [example 1](#), we obtain a power of roughly 76%. ◀

▷ Example 5: Multiple values of study parameters

In this example, we assess the effect of varying the value of the correlation of the male group on the power of our study. We supply a list of correlations in parentheses for the second command argument:

```
. power twocorrelations 0.3 (0.4(0.1)0.9), n(500)
Estimated power for a two-sample correlations test
Fisher's z test
Ho: r2 = r1 versus Ha: r2 != r1
```

alpha	power	N	N1	N2	delta	r1	r2
.05	.2452	500	250	250	.1	.3	.4
.05	.7595	500	250	250	.2	.3	.5
.05	.9894	500	250	250	.3	.3	.6
.05	1	500	250	250	.4	.3	.7
.05	1	500	250	250	.5	.3	.8
.05	1	500	250	250	.6	.3	.9

From the table, we see that the power increases from 25% to 100% as the correlation increases from 0.4 to 0.9.

For multiple values of parameters, the results are automatically displayed in a table, as we see above. For more examples of tables, see [\[PSS\] power, table](#). If you wish to produce a power plot, see [\[PSS\] power, graph](#).

Computing effect size and experimental-group correlation

Effect size δ for a two-sample correlations test is defined as the difference between the experimental-group and control-group correlations, $\delta = \rho_2 - \rho_1$.

Sometimes, we may be interested in determining the smallest effect (for a given control-group correlation) that yields a statistically significant result for prespecified sample size and power. In this case, power, sample size, and control-group correlation must be specified. In addition, you must also decide on the direction of the effect: upper, meaning $\rho_2 > \rho_1$, or lower, meaning $\rho_2 < \rho_1$. The direction may be specified in the `direction()` option; `direction(upper)` is the default.

► Example 6: Minimum detectable value of the experimental-group correlation

Continuing with [example 4](#), we will compute the smallest positive correlation of the male group that can be detected given a total sample of 500 individuals and a power of 80%. To solve for the experimental-group correlation, after the command, we specify the control-group correlation of 0.3, total sample size `n(500)`, and power `power(0.8)`:

```
. power twocorrelations 0.3, n(500) power(0.8)
Performing iteration ...
Estimated experimental-group correlation for a two-sample correlations test
Fisher's z test
Ho: r2 = r1 versus Ha: r2 != r1; r2 > r1
Study parameters:
      alpha =    0.0500
      power =    0.8000
         N =     500
N per group =    250
         r1 =    0.3000
Estimated effect size and experimental-group correlation:
      delta =    0.2092
         r2 =    0.5092
```

We find that the minimum value of the experimental-group correlation that would yield a statistically significant result in this study is 0.5092, and the corresponding effect size is 0.2092.

In this example, we computed the correlation assuming an upper direction, $\rho_2 > \rho_1$, or a positive effect, $\delta > 0$. To request a lower direction, or a negative effect, we can specify the `direction(lower)` option.

◀

Testing a hypothesis about two independent correlations

After data are collected, we can use the `mvtest` command to test the equality of two independent correlations using an asymptotic likelihood-ratio test; see [\[MV\] mvtest](#) for details. We can also manually perform a two-sample Fisher's z test on which `power twocorrelations` bases its computations. We demonstrate our examples using the `genderpsych` dataset from the [\[MV\] mvtest correlations](#) entry.

► Example 7: Comparing two correlations using mvtest

Consider a sample of 64 individuals with equal numbers of males and females. We would like to know whether the correlation between the pictorial inconsistencies (variable `y1`) and vocabulary (variable `y4`) is the same between males and females.

```
. use http://www.stata-press.com/data/r15/genderpsych
(Four Psychological Test Scores, Rencher and Christensen (2012))
. mvtest correlations y1 y4, by(gender)
Test of equality of correlation matrices across samples
      Jennrich chi2(1) =      2.16
      Prob > chi2 =      0.1415
```

The reported p -value is 0.1415, so we do not have sufficient evidence to reject the null hypothesis about the equality of the two correlations.

◀

▶ Example 8: Two-sample Fisher's z test

To compute a two-sample Fisher's z test manually, we perform the following steps. We first compute the estimates of correlation coefficients for each group by using the `correlate` command; see [R] [correlate](#). We then compute Fisher's z test statistic and its corresponding p -value.

We compute and store correlation estimates for males in the `r1` scalar and the corresponding sample size in the `N1` scalar.

```
. /* compute and store correlation and sample size for males */
. correlate y1 y4 if gender==1
(obs=32)
```

	y1	y4
y1	1.0000	
y4	0.5647	1.0000

```
. scalar r1 = r(rho)
. scalar N1 = r(N)
```

We store the corresponding results for females in `r2` and `N2`.

```
. /* compute and store correlation and sample size for females */
. correlate y1 y4 if gender==2
(obs=32)
```

	y1	y4
y1	1.0000	
y4	0.2596	1.0000

```
. scalar r2 = r(rho)
. scalar N2 = r(N)
```

We now compute the z test statistic, stored in the `Z` scalar, by applying Fisher's z transformation to the obtained correlation estimates `r1` and `r2` and using the cumulative function of the standard normal distribution `normal()` to compute the p -value.

```
. /* compute Fisher's z statistic and p-value and display results */
. scalar mu_Z = atanh(r2) - atanh(r1)
. scalar sigma_Z = sqrt(1/(N1-3)+1/(N2-3))
. scalar Z = mu_Z/sigma_Z
. scalar pvalue = 2*normal(-abs(Z))
. display "Z statistic = " %8.4g Z _n "P-value = " %8.4g pvalue
Z statistic = -1.424
P-value = .1543
```

The p -value is 0.1543 and is close to the p -value reported by `mvtest` in [example 7](#).

The estimates of the correlations obtained from `correlate` are 0.5647 for males and 0.2596 for females. We may check how many subjects we need to detect the difference between these two correlation values.

```
. power twocorrelations 0.5647 0.2596
Performing iteration ...
Estimated sample sizes for a two-sample correlations test
Fisher's z test
Ho: r2 = r1 versus Ha: r2 != r1
Study parameters:
      alpha =    0.0500
      power =    0.8000
      delta =   -0.3051
         r1 =    0.5647
         r2 =    0.2596
Estimated sample sizes:
           N =      232
      N per group =    116
```

We need a total of 232 subjects, 116 per group in a balanced design, to have a power of 80% to detect the difference between the given values of correlations for males and females.

◀

Stored results

`power twocorrelations` stores the following in `r()`:

Scalars

<code>r(alpha)</code>	significance level
<code>r(power)</code>	power
<code>r(beta)</code>	probability of a type II error
<code>r(delta)</code>	effect size
<code>r(N)</code>	total sample size
<code>r(N_a)</code>	actual sample size
<code>r(N1)</code>	sample size of the control group
<code>r(N2)</code>	sample size of the experimental group
<code>r(nratio)</code>	ratio of sample sizes, $N2/N1$
<code>r(nratio_a)</code>	actual ratio of sample sizes
<code>r(nfractional)</code>	1 if <code>nfractional</code> is specified; 0 otherwise
<code>r(onesided)</code>	1 for a one-sided test; 0 otherwise
<code>r(r1)</code>	control-group correlation
<code>r(r2)</code>	experimental-group correlation
<code>r(diff)</code>	difference between the experimental- and control-group correlations
<code>r(separator)</code>	number of lines between separator lines in the table
<code>r(divider)</code>	1 if <code>divider</code> is requested in the table; 0 otherwise
<code>r(init)</code>	initial value for sample sizes or experimental-group correlation
<code>r(maxiter)</code>	maximum number of iterations
<code>r(iter)</code>	number of iterations performed
<code>r(tolerance)</code>	requested parameter tolerance
<code>r(deltax)</code>	final parameter tolerance achieved
<code>r(ftolerance)</code>	requested distance of the objective function from zero
<code>r(function)</code>	final distance of the objective function from zero
<code>r(converged)</code>	1 if iteration algorithm converged; 0 otherwise

Macros	
r(type)	test
r(method)	twocorrelations
r(direction)	upper or lower
r(columns)	displayed table columns
r(labels)	table column labels
r(widths)	table column widths
r(formats)	table column formats
Matrices	
r(pss_table)	table of results

Methods and formulas

Let ρ_1 and ρ_2 denote Pearson's correlation for the control and the experimental groups, respectively. Let $\hat{\rho}_1$ and $\hat{\rho}_2$ denote the corresponding estimators, the sample correlation coefficients.

A two-sample correlations test involves testing the null hypothesis $H_0: \rho_2 = \rho_1$ versus the two-sided alternative hypothesis $H_a: \rho_2 \neq \rho_1$, the upper one-sided alternative $H_a: \rho_2 > \rho_1$, or the lower one-sided alternative $H_a: \rho_2 < \rho_1$.

The exact distribution of the difference between the sample correlation coefficients $\hat{\rho}_1$ and $\hat{\rho}_2$ is complicated for testing the hypothesis $H_0: \rho_1 - \rho_2 \neq 0$. An alternative is to apply an inverse hyperbolic tangent, $\tanh^{-1}(\cdot)$, or Fisher's z transformation to the estimators to approximate their sampling distribution by a normal distribution (Fisher 1915). The hypothesis $H_0: \rho_1 = \rho_2$ is then equivalent to $H_0: \tanh^{-1}(\rho_1) = \tanh^{-1}(\rho_2)$.

Let $Z_1 = \tanh^{-1}(\hat{\rho}_1) = 0.5 \ln\{(1 + \hat{\rho}_1)/(1 - \hat{\rho}_1)\}$ denote Fisher's z transformation of $\hat{\rho}_1$, and let Z_2 denote the corresponding transformation of $\hat{\rho}_2$. Then the difference $\delta_z = Z_2 - Z_1$ follows a normal distribution with mean $\mu_z = \mu_2 - \mu_1 = \tanh^{-1}(\rho_2) - \tanh^{-1}(\rho_1)$ and standard deviation $\sigma_z = \sqrt{1/(n_1 - 3) + 1/(n_2 - 3)}$, where n_1 and n_2 are the sample sizes of the control and the experimental groups, respectively (for example, see Graybill [1961, 211]; Anderson [2003, 134]).

Let α be the significance level, β be the probability of a type II error, and $z_{1-\alpha}$ and z_β be the $(1 - \alpha)$ th and the β th quantiles of the standard normal distribution.

The power $\pi = 1 - \beta$ is computed using

$$\pi = \begin{cases} \Phi\left(\frac{\delta_z}{\sigma_z} - z_{1-\alpha}\right) & \text{for an upper one-sided test} \\ \Phi\left(-\frac{\delta_z}{\sigma_z} - z_{1-\alpha}\right) & \text{for a lower one-sided test} \\ \Phi\left(\frac{\delta_z}{\sigma_z} - z_{1-\alpha/2}\right) + \Phi\left(-\frac{\delta_z}{\sigma_z} - z_{1-\alpha/2}\right) & \text{for a two-sided test} \end{cases} \quad (1)$$

where $\Phi(\cdot)$ is the cdf of a standard normal distribution.

Let $R = n_2/n_1$ denote the allocation ratio. Then $n_2 = R \times n_1$ and power can be viewed as a function of n_1 . Therefore, for sample-size determination, the control-group sample size n_1 is computed first. The experimental-group size n_2 is then computed as $R \times n_1$, and the total sample size is computed as $n = n_1 + n_2$. By default, sample sizes are rounded to integer values; see *Fractional sample sizes* in [PSS] **unbalanced designs** for details.

For a one-sided test, the control-group sample size n_1 is computed as a positive root of the following quadratic equation:

$$\frac{(R + 1)n_1 - 6}{(n_1 - 3)(Rn_1 - 3)} = \left(\frac{\delta_z}{z_{1-\alpha} - z_\beta}\right)^2$$

For a one-sided test, if one of the group sizes is known, the other one is computed using the following formula. For example, to compute n_1 given n_2 we use

$$n_1 = 3 + \frac{1}{\left(\frac{\delta_z}{z_{1-\alpha} - z_\beta}\right)^2 - \frac{1}{n_2 - 3}}$$

The minimum detectable value of the experimental-group correlation is obtained by first finding the corresponding minimum value of μ_2 and then applying the inverse Fisher's z transformation to that μ_2 :

$$\rho_2 = \frac{e^{2\mu_2} - 1}{e^{2\mu_2} + 1}$$

For a one-sided test, $\mu_2 = \mu_1 + \sigma_z(z_{1-\alpha} - z_\beta)$ when $\mu_2 > \mu_1$, and $\mu_2 = \mu_1 - \sigma_z(z_{1-\alpha} - z_\beta)$ when $\mu_2 < \mu_1$.

For a two-sided test, the sample sizes and μ_2 are obtained by iteratively solving the two-sided power equation in (1) for the respective parameters. The initial values are obtained from the respective formulas for the one-sided computation with the significance level $\alpha/2$.

References

- Anderson, T. W. 2003. *An Introduction to Multivariate Statistical Analysis*. 3rd ed. New York: Wiley.
- Fisher, R. A. 1915. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika* 10: 507–521.
- Graybill, F. A. 1961. *An Introduction to Linear Statistical Models, Vol. 1*. New York: McGraw–Hill.

Also see

- [PSS] **power** — Power and sample-size analysis for hypothesis tests
- [PSS] **power, graph** — Graph results from the power command
- [PSS] **power, table** — Produce table of results from the power command
- [PSS] **Glossary**
- [MV] **mvttest** — Multivariate tests
- [R] **correlate** — Correlations (covariances) of variables or coefficients