

power onevariance — Power analysis for a one-sample variance test

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

Description

`power onevariance` computes sample size, power, or target variance for a one-sample variance test. By default, it computes sample size for given power and the values of the variance parameters under the null and alternative hypotheses. Alternatively, it can compute power for given sample size and values of the null and alternative variances or the target variance for given sample size, power, and the null variance. The results can also be obtained for an equivalent standard deviation test, in which case standard deviations are used instead of variances. Also see [PSS-2] [power](#) for a general introduction to the `power` command using hypothesis tests.

For precision and sample-size analysis for a CI for a population variance, see [PSS-3] [ciwidth onevariance](#).

Quick start

Sample size for test of $H_0: \sigma^2 = 4$ versus $H_a: \sigma^2 \neq 4$ with null variance $v_0 = 4$ and alternative variance $v_a = 9$ with default power of 0.8 and significance level $\alpha = 0.05$

```
power onevariance 4 9
```

As above, but for $H_0: \sigma = 2$ versus $H_a: \sigma \neq 2$ with null standard deviation $s_0 = 2$ and alternative standard deviation $s_a = 3$

```
power onevariance 2 3, sd
```

As above, but for $\alpha = 0.1$

```
power onevariance 2 3, sd alpha(0.1)
```

Sample sizes for v_a equal to 7, 8, 9, 10, and 11

```
power onevariance 4 (7(1)11)
```

As above, but display results in a graph showing sample size versus alternative variance

```
power onevariance 4 (7(1)11), graph
```

Specify v_0 and the ratio of variances

```
power onevariance 4, ratio(2.25)
```

Power for a sample size of 30

```
power onevariance 4 9, n(30)
```

As above, but specify standard deviations rather than variances

```
power onevariance 2 3, n(30) sd
```

As above, but specify a one-sided test

```
power onevariance 2 3, sd n(30) onesided
```

Effect size and target variance for $v_0 = 4$ with a sample size of 20 and power of 0.8
power onevariance 4, n(20) power(.8)

Menu

Statistics > Power, precision, and sample size

Syntax

Compute sample size

Variance scale

```
power onevariance  $v_0$   $v_a$  [ , power(numlist) options ]
```

Standard deviation scale

```
power onevariance  $s_0$   $s_a$  , sd [ power(numlist) options ]
```

Compute power

Variance scale

```
power onevariance  $v_0$   $v_a$  , n(numlist) [ options ]
```

Standard deviation scale

```
power onevariance  $s_0$   $s_a$  , sd n(numlist) [ options ]
```

Compute effect size and target parameter

Target variance

```
power onevariance  $v_0$  , n(numlist) power(numlist) [ options ]
```

Target standard deviation

```
power onevariance  $s_0$  , sd n(numlist) power(numlist) [ options ]
```

where v_0 and s_0 are the null (hypothesized) variance and standard deviation or the value of the variance and standard deviation under the null hypothesis, and v_a and s_a are the alternative (target) variance and standard deviation or the value of the variance and standard deviation under the alternative hypothesis. Each argument may be specified either as one number or as a list of values in parentheses (see [U] 11.1.8 [numlist](#)).

<i>options</i>	Description
<code>sd</code>	request computation using the standard deviation scale; default is the variance scale
Main	
* <code>alpha(numlist)</code>	significance level; default is <code>alpha(0.05)</code>
* <code>power(numlist)</code>	power; default is <code>power(0.8)</code>
* <code>beta(numlist)</code>	probability of type II error; default is <code>beta(0.2)</code>
* <code>n(numlist)</code>	sample size; required to compute power or effect size
<code>nfractional</code>	allow fractional sample size
* <code>ratio(numlist)</code>	ratio of variances, v_a/v_0 (or ratio of standard deviations, s_a/s_0 , if option <code>sd</code> is specified); specify instead of the alternative variance v_a (or standard deviation s_a)
<code>direction(upper lower)</code>	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
<code>onesided</code>	one-sided test; default is two sided
<code>parallel</code>	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	
<code>[no]table[(tablespec)]</code>	suppress table or display results as a table; see [PSS-2] power, table
<code>saving(filename [, replace])</code>	save the table data to <i>filename</i> ; use <code>replace</code> to overwrite existing <i>filename</i>
Graph	
<code>graph[(graphopts)]</code>	graph results; see [PSS-2] power, graph
Iteration	
<code>init(#)</code>	initial value for sample size or variance
<code>iterate(#)</code>	maximum number of iterations; default is <code>iterate(500)</code>
<code>tolerance(#)</code>	parameter tolerance; default is <code>tolerance(1e-12)</code>
<code>ftolerance(#)</code>	function tolerance; default is <code>ftolerance(1e-12)</code>
<code>[no]log</code>	suppress or display iteration log
<code>[no]dots</code>	suppress or display iterations as dots
<code>notitle</code>	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.

`collect` is allowed; see [U] 11.1.10 **Prefix commands**.

`sd` does not appear in the dialog box; specification of `sd` is done automatically by the dialog box selected.

`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>	number of subjects	N
<code>delta</code>	effect size	δ
<code>v0</code>	null variance	σ_0^2
<code>va</code>	alternative variance	σ_a^2
<code>s0</code>	null standard deviation	σ_0
<code>sa</code>	alternative standard deviation	σ_a
<code>ratio</code>	ratio of the alternative variance to the null variance or ratio of the alternative standard deviation to the null standard deviation (if <code>sd</code> is specified)	σ_a^2/σ_0^2 σ_a/σ_0
<code>target</code>	target parameter; synonym for <code>va</code>	
<code>_all</code>	display all supported columns	

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

Columns `s0` and `sa` are displayed in the default table in place of the `v0` and `va` columns when the `sd` option is specified.

Column `ratio` is shown in the default table if specified. If the `sd` option is specified, this column contains the ratio of standard deviations. Otherwise, this column contains the ratio of variances.

Options

`sd` specifies that the computation be performed using the standard deviation scale. The default is to use the variance scale.

Main

`alpha()`, `power()`, `beta()`, `n()`, `nfractional`; see [\[PSS-2\] power](#). The `nfractional` option is allowed only for sample-size determination.

`ratio(numlist)` specifies the ratio of the alternative variance to the null variance, v_a/v_0 , or the ratio of standard deviations, s_a/s_0 , if the `sd` option is specified. You can specify either the alternative variance v_a as a command argument or the ratio of the variances in `ratio()`. If you specify `ratio(#)`, the alternative variance is computed as $v_a = v_0 \times \#$. This option is not allowed with the effect-size determination.

`direction()`, `onesided`, `parallel`; see [\[PSS-2\] power](#).

Table

`table`, `table()`, `notable`; see [\[PSS-2\] power, table](#).

`saving()`; see [\[PSS-2\] power](#).

Graph

`graph`, `graph()`; see [PSS-2] **power**, **graph**. Also see the *column* table for a list of symbols used by the graphs.

Iteration

`init(#)` specifies an initial value for the iteration procedure. Iteration is used to compute variance for a two-sided test and to compute sample size. The default initial value for the sample size is obtained from a closed-form normal approximation. The default initial value for the variance is obtained from a closed-form solution for a one-sided test with the significance level of $\alpha/2$.

`iterate()`, `tolerance()`, `ftolerance()`, `log`, `nolog`, `dots`, `nodots`; see [PSS-2] **power**.

The following option is available with `power onevariance` but is not shown in the dialog box:

`notitle`; see [PSS-2] **power**.

Remarks and examples

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

Remarks are presented under the following headings:

- Introduction*
- Using power onevariance*
- Computing sample size*
- Computing power*
- Computing effect size and target variance*
- Performing a hypothesis test on variance*

This entry describes the `power onevariance` command and the methodology for power and sample-size analysis for a one-sample variance test. See [PSS-2] **Intro (power)** for a general introduction to power and sample-size analysis and [PSS-2] **power** for a general introduction to the `power` command using hypothesis tests.

Introduction

The study of variance arises in cases where investigators are interested in making an inference on the variability of a process. For example, the precision of a thermometer in taking accurate measurements, the variation in the weights of potato chips from one bag to another, the variation in mileage across automobiles of the same model. Before undertaking the actual study, we may want to find the optimal sample size to detect variations that exceed the tolerable limits or industry-wide standards.

This entry describes power and sample-size analysis for the inference about the population variance performed using hypothesis testing. Specifically, we consider the null hypothesis $H_0: \sigma^2 = \sigma_0^2$ versus the two-sided alternative hypothesis $H_a: \sigma^2 \neq \sigma_0^2$, the upper one-sided alternative $H_a: \sigma^2 > \sigma_0^2$, or the lower one-sided alternative $H_a: \sigma^2 < \sigma_0^2$.

Hypothesis testing of variances relies on the assumption of normality of the data. For a random sample of size n from a normal distribution, the distribution of the sample variance s^2 is scaled χ^2 . The χ^2 test statistic $(n-1)s^2/\sigma_0^2$, which has a χ^2 distribution χ_{n-1}^2 with $n-1$ degrees of freedom, is used to test hypotheses on variance, and the corresponding test is known as a χ^2 test.

The test of a variance is equivalent to the test of a standard deviation with the null hypothesis $H_0: \sigma = \sigma_0$. The standard deviation test uses the same χ^2 test statistic. The only difference between the two tests is the scale or metric of the variability parameter: variance for the variance test and standard deviation for the standard deviation test. In some cases, standard deviations may provide a more meaningful interpretation than variances. For example, standard deviations of test scores or IQ have the same scale as the mean and provide information about the spread of the observations around the mean.

The `power onevariance` command provides power and sample-size analysis for the χ^2 test of a one-sample variance or a one-sample standard deviation.

Using power onevariance

`power onevariance` computes sample size, power, or target variance for a one-sample variance test. If the `sd` option is specified, `power onevariance` computes sample size, power, or target standard deviation for an equivalent one-sample standard-deviation 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.

In what follows, we describe the use of `power onevariance` in a variance metric. The corresponding use in a standard deviation metric, when the `sd` option is specified, is the same except variances v_0 and v_a should be replaced with the respective standard deviations s_0 and s_a . Note that computations using the variance and standard deviation scales yield the same results; the difference is only in the specification of the parameters.

To compute sample size, you must specify the variances under the null and alternative hypotheses, v_0 and v_a , 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.

To compute power, you must specify the sample size in the `n()` option and the variances under the null and alternative hypotheses as arguments v_0 and v_a , respectively.

Instead of the null and alternative variances v_0 and v_a , you can specify the null variance v_0 and the ratio of the alternative variance to the null variance in the `ratio()` option.

To compute effect size, the ratio of the alternative to the null variances, and target variance, you must specify the sample size in the `n()` option, the power in the `power()` option, the null variance v_0 , and, optionally, the direction of the effect. The direction is upper by default, `direction(upper)`, which means that the target variance is assumed to be larger than the specified null value. You can change the direction to lower, which means that the target variance is assumed to be smaller than the specified null value, by specifying the `direction(lower)` option.

By default, the computed sample size is rounded up. You can specify the `nfractional` option to see the corresponding fractional sample size; see [Fractional sample sizes](#) in [PSS-4] **Unbalanced designs** for an example. The `nfractional` option is allowed only for sample-size determination.

The test statistic for a one-sample variance test follows a χ^2 distribution. Its degrees of freedom depends on the sample size; therefore, sample-size computations require iteration. The effect-size determination for a two-sided test also requires iteration. The default initial value of the sample size is obtained using a closed-form normal approximation. The default initial value of the variance for the effect-size determination is obtained by using the corresponding computation for a one-sided test with the significance level $\alpha/2$. The default initial values may be changed by specifying the `init()` option. See [PSS-2] **power** for the descriptions of other options that control the iteration procedure.

In the following sections, we describe the use of `power onevariance` accompanied by examples for computing sample size, power, and target variance.

Computing sample size

To compute sample size, you must specify the variances under the null and alternative hypotheses, v_0 and v_a , 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 one-sample variance test

Consider a study where interest lies in testing whether the variability in mileage (measured in miles per gallon) of automobiles of a certain car manufacturer equals a specified value. Industry-wide standards maintain that a variation of at most two miles per gallon (mpg) from an average value is acceptable for commercial production.

The process engineer suspects that a faulty assembly line has been producing the variation higher than the acceptable standard. He or she wishes to test the null hypothesis of $H_0: \sigma = 2$ versus a two-sided alternative $H_a: \sigma \neq 2$ or, equivalently, $H_0: \sigma^2 = 4$ versus $H_a: \sigma^2 \neq 4$. The engineer wants to find the minimum number of cars so that the 5%-level two-sided test achieves the power of 80% to detect the alternative variance of 9 (or standard deviation of 3 mpg) given the null variance of 4 (or standard deviation of 2 mpg). To obtain the sample size, we specify the null and alternative values of the variance in v_0 and v_a after the command name:

```
. power onevariance 4 9
Performing iteration ...
Estimated sample size for a one-sample variance test
Chi-squared test
H0: v = v0 versus Ha: v != v0
Study parameters:
      alpha =    0.0500
      power =    0.8000
      delta =    2.2500
      v0 =     4.0000
      va =     9.0000
Estimated sample size:
      N =          24
```

We find that a sample of 24 subjects is required for this study.

As we mentioned in the previous section, sample-size computation requires iteration. By default, `power onevariance` suppresses the iteration log, which may be displayed by specifying the `log` option.

▷ Example 2: Specifying ratio of variances

Instead of the alternative variance as in [example 1](#), we can specify the ratio of the alternative and null variances of $9/4 = 2.25$ in the `ratio()` option:

```
. power onevariance 4, ratio(2.25)
Performing iteration ...
Estimated sample size for a one-sample variance test
Chi-squared test
H0: v = v0 versus Ha: v != v0
Study parameters:
    alpha =    0.0500
    power =    0.8000
    delta =    2.2500
    v0 =      4.0000
    va =      9.0000
    ratio =    2.2500
Estimated sample size:
    N =          24
```

We obtain the same results as in [example 1](#). The ratio of the variances is now also displayed in the output.

◀

▷ Example 3: Standard deviation test

We can use the `sd` option to perform calculations in the standard deviation metric. We reproduce results from [example 1](#) using the corresponding null and standard deviations of 2 and 3.

```
. power onevariance 2 3, sd
Performing iteration ...
Estimated sample size for a one-sample standard-deviation test
Chi-squared test
H0: s = s0 versus Ha: s != s0
Study parameters:
    alpha =    0.0500
    power =    0.8000
    delta =    1.5000
    s0 =      2.0000
    sa =      3.0000
Estimated sample size:
    N =          24
```

The results are the same, except the output reports standard deviations instead of variances.

◀

Computing power

To compute power, you must specify the sample size in the `n()` option and the variances under the null and alternative hypotheses, v_0 and v_a , respectively.

▷ Example 4: Power of a one-sample variance test

Continuing with [example 1](#), we will suppose that we are designing a new study and anticipate obtaining a sample of 30 cars. To compute the power corresponding to this sample size given the study parameters from [example 1](#), we specify the sample size of 30 in the `n()` option:

```
. power onevariance 4 9, n(30)
Estimated power for a one-sample variance test
Chi-squared test
H0: v = v0 versus Ha: v != v0
Study parameters:
    alpha =    0.0500
      N =     30
    delta =    2.2500
     v0 =    4.0000
     va =    9.0000
Estimated power:
    power =    0.8827
```

With a larger sample size, the power of the test increases to about 88%.



▷ Example 5: Multiple values of study parameters

Suppose we would like to assess the effect of increasing the alternative variance on the power of the test. We do this by specifying a range of values in parentheses in the argument for the alternative variance:

```
. power onevariance 4 (4.5(0.5)10), n(30)
Estimated power for a one-sample variance test
Chi-squared test
H0: v = v0 versus Ha: v != v0
```

alpha	power	N	delta	v0	va
.05	.08402	30	1.125	4	4.5
.05	.1615	30	1.25	4	5
.05	.2694	30	1.375	4	5.5
.05	.391	30	1.5	4	6
.05	.511	30	1.625	4	6.5
.05	.6189	30	1.75	4	7
.05	.7098	30	1.875	4	7.5
.05	.7829	30	2	4	8
.05	.8397	30	2.125	4	8.5
.05	.8827	30	2.25	4	9
.05	.9147	30	2.375	4	9.5
.05	.9382	30	2.5	4	10

The power is an increasing function of the effect size, which is measured by the ratio of the alternative variance to the null variance.

For multiple values of parameters, the results are automatically displayed in a table, as we see above. For more examples of tables, see [\[PSS-2\] power, table](#). If you wish to produce a power plot, see [\[PSS-2\] power, graph](#).



Computing effect size and target variance

Effect size δ for a one-sample variance test is defined as the ratio of the alternative variance to the null variance $\delta = v_a/v_0$ or the ratio of the alternative standard deviation to the null standard deviation $\delta = s_a/s_0$ when the `sd` option is specified.

Sometimes, we may be interested in determining the smallest effect that yields a statistically significant result for prespecified sample size and power. In this case, power, sample size, and the null variance or the null standard deviation must be specified. In addition, you must also decide on the direction of the effect: upper, meaning $v_a > v_0$ ($s_a > s_0$), or lower, meaning $v_a < v_0$ ($s_a < s_0$). The direction may be specified in the `direction()` option; `direction(upper)` is the default.

► Example 6: Minimum detectable value of the variance

Continuing with [example 4](#), we may also be interested to find the minimum effect size that can be detected with a power of 80% given a sample of 30 subjects. To compute the smallest effect size and the corresponding target variance, after the command name, we specify the null variance of 4, sample size `n(30)`, and power `power(0.8)`:

```
. power onevariance 4, n(30) power(0.8)
Performing iteration ...
Estimated target variance for a one-sample variance test
Chi-squared test
H0: v = v0 versus Ha: v != v0; va > v0
Study parameters:
      alpha =    0.0500
      power =    0.8000
      N      =     30
      v0     =    4.0000
Estimated effect size and target variance:
      delta =    2.0343
      va    =    8.1371
```

The smallest detectable value of the effect size, the ratio of the variances, is 2.03, which corresponds to the alternative variance of 8.14. Compared with [example 1](#), for the same power of 80%, this example shows a smaller variance with a larger sample of 30 subjects.

Above we assumed the effect to be in the upper direction. The effect size and target variance in the lower direction can be obtained by specifying `direction(lower)`.

```

. power onevariance 4, n(30) power(0.8) direction(lower)
Performing iteration ...
Estimated target variance for a one-sample variance test
Chi-squared test
H0: v = v0 versus Ha: v != v0; va < v0
Study parameters:
      alpha =    0.0500
      power =    0.8000
         N =     30
        v0 =    4.0000
Estimated effect size and target variance:
      delta =    0.4567
        va =    1.8267

```

The smallest detectable value of the effect size is 0.46, which corresponds to the alternative variance of 1.83.

◀

Performing a hypothesis test on variance

In this section, we demonstrate the use of the `sdtest` command for testing hypotheses about variances; see [R] [sdtest](#) for details. Suppose we wish to test the hypothesis that the variance or standard deviation is different from a specific null value on the collected data. We can use the `sdtest` command to do this.

► Example 7: Testing for variance

We use `auto.dta` to demonstrate the use of `sdtest`. We have data on mileage ratings of 74 automobiles and wish to test whether the overall standard deviation is different from 3 miles per gallon (mpg).

```

. use https://www.stata-press.com/data/r17/auto
(1978 automobile data)
. sdtest mpg == 3
One-sample test of variance

```

Variable	Obs	Mean	Std. err.	Std. dev.	[95% conf. interval]	
mpg	74	21.2973	.6725511	5.785503	19.9569	22.63769

```

      sd = sd(mpg)                                c = chi2 = 271.4955
H0: sd = 3                                       Degrees of freedom = 73
      Ha: sd < 3                                Ha: sd != 3                                Ha: sd > 3
Pr(C < c) = 1.0000                            2*Pr(C > c) = 0.0000                            Pr(C > c) = 0.0000

```

We find statistical evidence to reject the null hypothesis of $H_0: \sigma_{\text{mpg}} = 3$ versus a two-sided alternative $H_a: \sigma_{\text{mpg}} \neq 3$ at the 5% significance level; the p -value < 0.0000 .

We use the estimates of this study to perform a sample-size analysis we would have conducted before the study.

```

. power onevar 3 5.78, sd
Performing iteration ...
Estimated sample size for a one-sample standard-deviation test
Chi-squared test
H0: s = s0 versus Ha: s != s0
Study parameters:
      alpha =    0.0500
      power =    0.8000
      delta =    1.9267
      s0     =    3.0000
      sa     =    5.7800
Estimated sample size:
      N =      10

```

We find that the sample size required to detect a standard deviation of 5.78 mpg given the null value of 3 mpg with 80% power using a 5%-level two-sided test is only 10.

◀

Stored results

`power onevariance` 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>	sample size
<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(v0)</code>	variance under the null hypothesis (for variance scale, default)
<code>r(va)</code>	variance under the alternative hypothesis (for variance scale, default)
<code>r(s0)</code>	standard deviation under the null hypothesis (if option <code>sd</code> is specified)
<code>r(sa)</code>	standard deviation under the alternative hypothesis (if option <code>sd</code> is specified)
<code>r(ratio)</code>	ratio of the alternative variance to the null variance (or the ratio of standard deviations if option <code>sd</code> is specified)
<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 size or variance
<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

<code>r(type)</code>	test
<code>r(method)</code>	onevariance
<code>r(direction)</code>	upper or lower
<code>r(columns)</code>	displayed table columns
<code>r(labels)</code>	table column labels
<code>r(widths)</code>	table column widths
<code>r(formats)</code>	table column formats
<code>r(scale)</code>	variance or standard deviation

Matrices

<code>r(pss_table)</code>	table of results
---------------------------	------------------

Methods and formulas

Consider a random sample of size n from a normal population with mean μ and variance σ^2 . Let σ_0^2 and σ_a^2 denote the null and alternative values of the variance parameter, respectively.

A one-sample variance test involves testing the null hypothesis $H_0: \sigma^2 = \sigma_0^2$ versus the two-sided alternative hypothesis $H_a: \sigma^2 \neq \sigma_0^2$, the upper one-sided alternative $H_a: \sigma^2 > \sigma_0^2$, or the lower one-sided alternative $H_a: \sigma^2 < \sigma_0^2$.

The sampling distribution of the test statistic $\chi^2 = (n-1)s^2/\sigma^2$ under the null hypothesis follows a χ^2 distribution with $n-1$ degrees of freedom, where s^2 is the sample variance. The corresponding test is known as a χ^2 test.

The following formulas are based on [Dixon and Massey \(1983, 110–112\)](#).

Let α be the significance level, β be the probability of a type II error, and $\chi_{n-1,1-\alpha}^2$ and $\chi_{n-1,\beta}^2$ be the $(1-\alpha)$ th and the β th quantiles of the χ^2 distribution with $n-1$ degrees of freedom.

The following equality holds at the critical value of the accept/reject boundary for H_0 :

$$\frac{\chi_{n-1,1-\alpha}^2}{n-1} \sigma_0^2 = \frac{\chi_{n-1,\beta}^2}{n-1} \sigma_a^2$$

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

$$\pi = \begin{cases} 1 - \chi_{n-1}^2 \left(\frac{\sigma_0^2}{\sigma_a^2} \chi_{n-1,1-\alpha}^2 \right) & \text{for an upper one-sided test} \\ \chi_{n-1}^2 \left(\frac{\sigma_0^2}{\sigma_a^2} \chi_{n-1,\alpha}^2 \right) & \text{for a lower one-sided test} \\ 1 - \chi_{n-1}^2 \left(\frac{\sigma_0^2}{\sigma_a^2} \chi_{n-1,1-\alpha/2}^2 \right) + \chi_{n-1}^2 \left(\frac{\sigma_0^2}{\sigma_a^2} \chi_{n-1,\alpha/2}^2 \right) & \text{for a two-sided test} \end{cases} \quad (1)$$

where $\chi_{n-1}^2(\cdot)$ is the cdf of a χ^2 distribution with $n-1$ degrees of freedom.

Sample size n is obtained by iteratively solving the corresponding power equation from (1) for n . The default initial value for the sample size is obtained by using a large-sample normal approximation.

For a large n , the log-transformed sample variance is approximately normal with mean $2 \ln(\sigma)$ and standard deviation $\sqrt{2/n}$. The approximate sample size is then given by

$$n = \frac{1}{2} \left\{ \frac{z_{1-\alpha/k} - z_\beta}{\ln \left(\frac{\sigma_a}{\sigma_0} \right)} \right\}^2$$

where $k = 1$ for a one-sided test and $k = 2$ for a two-sided test.

For a one-sided test, the minimum detectable value of the variance is computed as follows:

$$\sigma_a^2 = \begin{cases} \sigma_0^2 \frac{\chi_{n-1,1-\alpha}^2}{\chi_{n-1,\beta}^2} & \text{for an upper one-sided test} \\ \sigma_0^2 \frac{\chi_{n-1,\alpha}^2}{\chi_{n-1,1-\beta}^2} & \text{for a lower one-sided test} \end{cases} \quad (2)$$

For a two-sided test, the minimum detectable value of the variance is computed by iteratively solving the two-sided power equation from (1) for σ_a^2 . The default initial value is obtained from (2) with α replaced by $\alpha/2$.

If the `nfractional` option is not specified, the computed sample size is rounded up.

Reference

Dixon, W. J., and F. J. Massey, Jr. 1983. *Introduction to Statistical Analysis*. 4th ed. New York: McGraw-Hill.

Also see

[PSS-2] **power** — Power and sample-size analysis for hypothesis tests

[PSS-2] **power, graph** — Graph results from the power command

[PSS-2] **power, table** — Produce table of results from the power command

[PSS-3] **ciwidth onevariance** — Precision analysis for a one-variance CI

[PSS-5] **Glossary**

[R] **sdtest** — Variance-comparison tests