Description

ci computes confidence intervals for population means, proportions, variances, and standard deviations.

cii is the immediate form of ci; see [U] **19 Immediate commands** for a general discussion of immediate commands.

Quick start

Confidence intervals for means of normally distributed variables v1, v2, and v3

    ci means v1-v3

Confidence interval for mean of Poisson-distributed variable v4

    ci means v4, poisson

Confidence interval for rate of v4 with total exposure recorded in v5

    ci means v4, poisson exposure(v5)

Confidence interval for proportion of binary variable v6

    ci proportions v6

Confidence intervals for variances of v1, v2, and v3

    ci variances v1-v3

Same as above, but Bonett confidence intervals are produced

    ci variances v1-v3, bonett

90% Bonett confidence intervals for standard deviations of v1, v2, and v3

    ci variances v1-v3, sd bonett level(90)

Confidence interval for a mean based on a sample with 85 observations, a sample mean of 10, and a standard deviation of 3

    cii means 85 10 3

90% confidence interval for rate from a sample with 4,379 deaths over 11,394 person-years

    cii means 11394 4379, poisson level(90)

Agresti–Coull confidence interval for proportion based on a sample with 2,377 observations and 136 successes

    cii proportions 2377 136, agresti

Bonett confidence interval for variance based on a sample with 20 observations, sample variance of 9, and estimated kurtosis of 1.8

```
cii variances 20 9 1.8, bonett
```

Same as above, but with confidence interval for standard deviation

```
cii variances 20 3 1.8, sd bonett
```

Menu

ci

Statistics > Summaries, tables, and tests > Summary and descriptive statistics > Confidence intervals

cii for a normal mean

Statistics > Summaries, tables, and tests > Summary and descriptive statistics > Normal mean CI calculator

cii for a Poisson mean

Statistics > Summaries, tables, and tests > Summary and descriptive statistics > Poisson mean CI calculator

cii for a proportion

Statistics > Summaries, tables, and tests > Summary and descriptive statistics > Proportion CI calculator

cii for a variance

Statistics > Summaries, tables, and tests > Summary and descriptive statistics > Variance CI calculator

cii for a standard deviation

Statistics > Summaries, tables, and tests > Summary and descriptive statistics > Standard deviation CI calculator

Syntax

Confidence intervals for means, normal distribution

> ci <u>mean</u>s [*varlist*] [*if*] [*in*] [*weight*] [, *options*]

> cii <u>mean</u>s $\#_{obs}$ $\#_{mean}$ $\#_{sd}$ [, <u>l</u>evel(#)]

Confidence intervals for means, Poisson distribution

> ci <u>mean</u>s [*varlist*] [*if*] [*in*] [*weight*] , <u>poisson</u> [<u>expo</u>sure(*varname*) *options*]

> cii <u>mean</u>s $\#_{exposure}$ $\#_{events}$, <u>poisson</u> [<u>l</u>evel(#)]

Confidence intervals for proportions

> ci <u>prop</u>ortions [*varlist*] [*if*] [*in*] [*weight*] [, *prop_options options*]

> cii <u>prop</u>ortions $\#_{obs}$ $\#_{succ}$ [, *prop_options* <u>l</u>evel(#)]

Confidence intervals for variances

> ci <u>var</u>iances [*varlist*] [*if*] [*in*] [*weight*] [, <u>bon</u>ett *options*]

> cii <u>var</u>iances $\#_{obs}$ $\#_{variance}$ [, <u>l</u>evel(#)]

> cii <u>var</u>iances $\#_{obs}$ $\#_{variance}$ $\#_{kurtosis}$, <u>bon</u>ett [<u>l</u>evel(#)]

Confidence intervals for standard deviations

> ci <u>var</u>iances [*varlist*] [*if*] [*in*] [*weight*] , sd [<u>bon</u>ett *options*]

> cii <u>var</u>iances $\#_{obs}$ $\#_{sd}$, sd [<u>l</u>evel(#)]

> cii <u>var</u>iances $\#_{obs}$ $\#_{sd}$ $\#_{kurtosis}$, sd <u>bon</u>ett [<u>l</u>evel(#)]

$\#_{obs}$ must be a positive integer. $\#_{exposure}$, $\#_{sd}$, and $\#_{variance}$ must be a positive number. $\#_{succ}$ and $\#_{events}$ must be a nonnegative integer or between 0 and 1. If the number is between 0 and 1, Stata interprets it as the fraction of successes or events and converts it to an integer number representing the number of successes or events. The computation then proceeds as if two integers had been specified. If the bonett option is specified, you must additionally specify $\#_{kurtosis}$ with cii variances.

prop_options	Description
exact	calculate exact confidence intervals; the default
wald	calculate Wald confidence intervals
wilson	calculate Wilson confidence intervals
agresti	calculate Agresti–Coull confidence intervals
jeffreys	calculate Jeffreys confidence intervals

options	Description
level(*#*)	set confidence level; default is level(95)
separator(*#*)	draw separator line after every # variables; default is separator(5)
total	add output for all groups combined (for use with by only)

by, collect, and statsby are allowed with ci, and collect is allowed with cii; see [U] **11.1.10 Prefix commands**.

aweights are allowed with ci means for normal data, and fweights are allowed with all ci subcommands; see [U] **11.1.6 weight**.

Options

Options are presented under the following headings:

> Options for ci and cii means
> Options for ci and cii proportions
> Options for ci and cii variances

Options for ci and cii means

> [Main]

poisson specifies that the variables (or numbers for cii) are Poisson-distributed counts; exact Poisson confidence intervals will be calculated. By default, confidence intervals for means are calculated based on a normal distribution.

exposure(*varname*) is used only with poisson. You do not need to specify poisson if you specify exposure(); poisson is assumed. *varname* contains the total exposure (typically a time or an area) during which the number of events recorded in *varlist* was observed.

level(*#*) specifies the confidence level, as a percentage, for confidence intervals. The default is level(95) or as set by set level; see [R] **level**.

separator(*#*) specifies how often separation lines should be inserted into the output. The default is separator(5), meaning that a line is drawn after every five variables. separator(10) would draw the line after every 10 variables. separator(0) suppresses the separation line.

total is used with the by prefix. It requests that in addition to output for each by-group, output be added for all groups combined.

Options for ci and cii proportions

> **Main**

exact, wald, wilson, agresti, and jeffreys specify how binomial confidence intervals are to be calculated.

exact is the default and specifies exact (also known in the literature as Clopper–Pearson [1934]) binomial confidence intervals.

wald specifies calculation of Wald confidence intervals.

wilson specifies calculation of Wilson confidence intervals.

agresti specifies calculation of Agresti–Coull confidence intervals.

jeffreys specifies calculation of Jeffreys confidence intervals.

See Brown, Cai, and DasGupta (2001) for a discussion and comparison of the different binomial confidence intervals.

level(#) specifies the confidence level, as a percentage, for confidence intervals. The default is level(95) or as set by set level; see [R] **level**.

separator(#) specifies how often separation lines should be inserted into the output. The default is separator(5), meaning that a line is drawn after every five variables. separator(10) would draw the line after every 10 variables. separator(0) suppresses the separation line.

total is used with the by prefix. It requests that in addition to output for each by-group, output be added for all groups combined.

Options for ci and cii variances

> **Main**

sd specifies that confidence intervals for standard deviations be calculated. The default is to compute confidence intervals for variances.

bonett specifies that Bonett confidence intervals be calculated. The default is to compute normal-based confidence intervals, which assume normality for the data.

level(#) specifies the confidence level, as a percentage, for confidence intervals. The default is level(95) or as set by set level; see [R] **level**.

separator(#) specifies how often separation lines should be inserted into the output. The default is separator(5), meaning that a line is drawn after every five variables. separator(10) would draw the line after every 10 variables. separator(0) suppresses the separation line.

total is used with the by prefix. It requests that in addition to output for each by-group, output be added for all groups combined.

Remarks and examples

Remarks are presented under the following headings:

Confidence intervals for means

ci means computes a confidence interval for the population mean for each of the variables in *varlist*.

Normal-based confidence intervals

▷ Example 1: Normal-based confidence intervals

Without the poisson option, ci means produces normal-based confidence intervals that are correct if the variable is normally distributed and asymptotically correct for all other distributions satisfying the conditions of the central limit theorem.

```
. use https://www.stata-press.com/data/r19/auto
(1978 automobile data)
. ci means mpg price
```

Variable	Obs	Mean	Std. err.	[95% conf. interval]	
mpg	74	21.2973	.6725511	19.9569	22.63769
price	74	6165.257	342.8719	5481.914	6848.6

The standard error of the mean of mpg is 0.67, and the 95% confidence interval is $[\,19.96, 22.64\,]$. We can obtain wider confidence intervals, 99%, by typing

```
. ci means mpg price, level(99)
```

Variable	Obs	Mean	Std. err.	[99% conf. interval]	
mpg	74	21.2973	.6725511	19.51849	23.07611
price	74	6165.257	342.8719	5258.405	7072.108

◁

▷ Example 2: The by prefix

The by prefix breaks out the confidence intervals according to by-group; total adds an overall summary. For instance,

```
. by foreign: ci means mpg, total
```

-> foreign = Domestic

Variable	Obs	Mean	Std. err.	[95% conf. interval]
mpg	52	19.82692	.657777	18.50638 21.14747

-> foreign = Foreign

Variable	Obs	Mean	Std. err.	[95% conf. interval]
mpg	22	24.77273	1.40951	21.84149 27.70396

-> Total

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

◁

▷ Example 3: Controlling the format

You can control the formatting of the numbers in the output by specifying a display format for the variable; see [U] **12.5 Formats: Controlling how data are displayed**. For instance,

```
. format mpg %9.2f
. ci means mpg
```

Variable	Obs	Mean	Std. err.	[95% conf. interval]
mpg	74	21.30	0.67	19.96 22.64

◁

Poisson confidence intervals

If you specify the poisson option, ci means assumes count data and computes exact Poisson confidence intervals.

▷ Example 4: Poisson confidence intervals

We have data on the number of bacterial colonies on a Petri dish. The dish has been divided into 36 small squares, and the number of colonies in each square has been counted. Each observation in our dataset represents a square on the dish. The variable count records the number of colonies in each square counted, which varies from 0 to 5.

```
. use https://www.stata-press.com/data/r19/petri, clear
. ci means count, poisson
```

Variable	Exposure	Mean	Std. err.	Poisson exact [95% conf. interval]
count	36	2.333333	.2545875	1.861158 2.888825

ci reports that the average number of colonies per square is 2.33. If the expected number of colonies per square were as low as 1.86, the probability of observing 2.33 or more colonies per square would be 2.5%. If the expected number were as large as 2.89, the probability of observing 2.33 or fewer colonies per square would be 2.5%.

◁

▷ Example 5: Option exposure()

The number of "observations"—how finely the Petri dish is divided—makes no difference. The Poisson distribution is a function only of the count. In example 4, we observed a total of $2.33 \times 36 = 84$ colonies and a confidence interval of $[1.86 \times 36, 2.89 \times 36] = [67, 104]$. We would obtain the same $[67, 104]$ confidence interval if our dish were divided into, say, 49 squares rather than 36.

For the counts, it is not even important that all the squares be of the same size. For rates, however, such differences do matter but in an easy-to-calculate way. Rates are obtained from counts by dividing by exposure, which is typically a number multiplied by either time or an area. For our Petri dishes, we divide by an area to obtain a rate, but if our example were cast in terms of being infected by a disease, we might divide by person-years to obtain the rate. Rates are convenient because they are easier to compare: we might have 2.3 colonies per square inch or 0.0005 infections per person-year.

So let's assume that we wish to obtain the number of colonies per square inch and, moreover, that not all the "squares" on our dish are of equal size. We have a variable called `area` that records the area of each square:

```
. ci means count, exposure(area)
```

| | | | | Poisson exact | |
Variable	Exposure	Mean	Std. err.	[95% conf.	interval]
count	3	28	3.055051	22.3339	34.66591

The rates are now in more familiar terms. In our sample, there are 28 colonies per square inch, and the 95% confidence interval is $[22.3, 34.7]$. When we did not specify `exposure()`, `ci means` with the `poisson` option assumed that each observation contributed 1 to exposure.

◁

❑ Technical note

If there were no colonies on our dish, `ci means` with the `poisson` option would calculate a one-sided confidence interval:

```
. use https://www.stata-press.com/data/r19/petrinone
. ci means count, poisson
```

| | | | | Poisson exact | |
Variable	Exposure	Mean	Std. err.	[95% conf.	interval]
count	36	0	0	0	.1024689*

(*) one-sided, 97.5% confidence interval

❑

Confidence intervals for proportions

The `ci proportions` command assumes binary (0/1) data and computes binomial confidence intervals.

▷ Example 6: Exact binomial (Clopper–Pearson) confidence interval

We have data on employees, including a variable marking whether the employee was promoted last year.

```
. use https://www.stata-press.com/data/r19/promo
. ci proportions promoted
```

Variable	Obs	Proportion	Std. err.	Binomial exact [95% conf. interval]	
promoted	20	.1	.067082	.0123485	.3169827

The exact binomial, also known as the Clopper–Pearson (1934) interval, is computed by default.

Nominally, the interpretation of a 95% confidence interval is that under repeated samples or experiments, 95% of the resultant intervals would contain the unknown parameter in question. However, for binomial data, the actual coverage probability, regardless of method, usually differs from that interpretation. This result occurs because of the discreteness of the binomial distribution, which produces only a finite set of outcomes, meaning that coverage probabilities are subject to discrete jumps and that the exact nominal level cannot always be achieved. Therefore, the term "exact confidence interval" refers to its being derived from the binomial distribution, the distribution exactly generating the data, rather than resulting in exactly the nominal coverage.

For the Clopper–Pearson interval, the actual coverage probability is guaranteed to be greater than or equal to the nominal confidence level, here 95%. Because of the way it is calculated—see *Methods and formulas*—it may also be interpreted as follows: If the true probability of being promoted were 0.012, the chances of observing a result as extreme or more extreme than the result observed ($20 \times 0.1 = 2$ or more promotions) would be 2.5%. If the true probability of being promoted were 0.317, the chances of observing a result as extreme or more extreme than the result observed (two or fewer promotions) would be 2.5%.

◁

▷ Example 7: Other confidence intervals

The Clopper–Pearson interval is desirable because it guarantees nominal coverage; however, by dropping this restriction, you may obtain accurate intervals that are not as conservative. In this vein, you might opt for the Wilson (1927) interval,

```
. ci proportions promoted, wilson
```

Variable	Obs	Proportion	Std. err.	Wilson [95% conf. interval]	
promoted	20	.1	.067082	.0278665	.3010336

the Agresti–Coull (1998) interval,

```
. ci proportions promoted, agresti
```

Variable	Obs	Proportion	Std. err.	Agresti–Coull [95% conf. interval]	
promoted	20	.1	.067082	.0156562	.3132439

or the Bayesian-derived Jeffreys interval (Brown, Cai, and DasGupta 2001),

```
. ci proportions promoted, jeffreys
```

| | | | | Jeffreys | |
Variable	Obs	Proportion	Std. err.	[95% conf. interval]
promoted	20	.1	.067082	.0213725	.2838533

Picking the best interval is a matter of balancing accuracy (coverage) against precision (average interval length) and depends on sample size and success probability. Brown, Cai, and DasGupta (2001) recommend the Wilson or Jeffreys interval for small sample sizes (≤ 40) yet favor the Agresti–Coull interval for its simplicity, decent performance for sample sizes less than or equal to 40, and performance comparable to Wilson or Jeffreys for sample sizes greater than 40. They also deem the Clopper–Pearson interval to be "wastefully conservative and [...] not a good choice for practical use", unless of course one requires, at a minimum, the nominal coverage level.

Finally, the binomial Wald confidence interval is obtained by specifying the `wald` option. The Wald interval is the one taught in most introductory statistics courses and, for the above, is simply, for level $1 - \alpha$, Proportion$\pm z_{\alpha/2}$(Std. err.), where $z_{\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard normal. Because its overall poor performance makes it impractical, the Wald interval is available mainly for pedagogical purposes. The binomial Wald interval is also similar to the interval produced by treating binary data as normal data and using `ci means`, with two exceptions. First, the calculation of the standard error in `ci proportions` uses denominator n rather than $n - 1$, used for normal data in `ci means`. Second, confidence intervals for normal data are based on the t distribution rather than the standard normal. Of course, both discrepancies vanish as sample size increases.

◁

❑ Technical note

Let's repeat example 6, but this time with data in which there are no promotions over the observed period:

```
. use https://www.stata-press.com/data/r19/promonone
. ci proportions promoted
```

| | | | | Binomial exact | |
Variable	Obs	Proportion	Std. err.	[95% conf. interval]
promoted	20	0	0	0	.1684335*

(*) one-sided, 97.5% confidence interval

The confidence interval is $[\,0, 0.168\,]$, and this is the confidence interval that most books publish. It is not, however, a true 95% confidence interval because the lower tail has vanished. As Stata notes, it is a one-sided, 97.5% confidence interval. If you wanted to put 5% in the right tail, you could type `ci proportions promoted, level(90)`.

❑

❑ Technical note

ci proportions ignores any variables that do not take on the values 0 and 1 exclusively. For instance, with our automobile dataset,

```
. use https://www.stata-press.com/data/r19/auto
(1978 automobile data)
. ci proportions mpg foreign
```

Variable	Obs	Proportion	Std. err.	Binomial exact [95% conf. interval]	
foreign	74	.2972973	.0531331	.196584	.4148353

Note: The results are produced only for binary (0/1) variables.

We also requested the confidence interval for mpg, but Stata ignored us. It does that so you can type ci proportions and obtain correct confidence intervals for all the variables that are 0/1 in your data.

❑

Confidence intervals for variances

The ci variances command computes confidence intervals for the variances or, if the sd option is specified, for the standard deviations. The default is a normal-based confidence interval that assumes the data are normal and uses a χ^2 distribution to construct the confidence intervals. If normality is suspect, you may consider using the bonett option to compute Bonett (2006) confidence intervals, which are more robust to nonnormality.

▷ Example 8: Normal-based confidence intervals

So far, we have restricted our attention to confidence intervals for means and proportions. Typically, when people think of statistical inference, they usually have in mind inferences concerning population means. However, the population parameter of interest will vary from one situation to another. In many scenarios, the population variance is as important as the population mean. For example, in a quality control study, a machine that fills 16-ounce canned peas is investigated at regular time intervals. A random sample of $n = 8$ containers is selected every hour. Ideally, the amount of peas in a can should vary only slightly about the 16-ounce value. If the variance was large, then a large proportion of cans would be either underfilled, thus cheating the customer, or overfilled, thus resulting in economic loss to the manufacturing company. Suppose that the weights of 16-ounce cans filled by the machine are normally distributed. The acceptable variability in the weights is expected to be 0.09 with the respective standard deviation of 0.3 ounces. To monitor the machine's performance, we can compute confidence intervals for the variance of the weights of cans:

```
. use https://www.stata-press.com/data/r19/peas_normdist
(Weights of canned peas, normal distribution)
. ci variances weight
```

Variable	Obs	Variance	[95% conf. interval]	
weight	8	.3888409	.1699823	1.610708

The command reports the sample estimate of the variance of 0.39 with the 95% confidence interval of [0.17, 1.61].

Instead of the variance, we may be interested in confidence intervals for the standard deviation. We can specify the sd option to compute such confidence intervals.

```
. ci variances weight, sd
    Variable |        Obs      Std. dev.      [95% conf. interval]
-------------+-----------------------------------------------------
      weight |          8      .6235711       .4122891    1.269137
```

The 95% confidence interval for the standard deviation of the weights is [0.41, 1.27]. Because the desired value for the standard deviation, 0.3 ounces, falls outside the interval, the machine may require some tuning.

◁

Confidence intervals in example 8 are based on the assumption that the random sample is selected from a population having a normal distribution. Nonnormality of the population distribution, in the form of skewness or heavy tails, can have a drastic impact on the asymptotic coverage probability of the normal-based confidence intervals. This is the case even for distributions that are similar to normal. Scheffé (1959, 336) showed that the normal-based interval has an asymptotic coverage probability of about 0.76, 0.63, 0.60, and 0.51 for the logistic, t with seven degrees of freedom, Laplace, and t with five degrees of freedom distributions, respectively. Miller (1997, 264) describes this situation as "catastrophic" because these distributions are symmetric and not easily distinguishable from a normal distribution unless the sample size is large. Hence, it is judicious to evaluate the normality of the data prior to constructing the normal-based confidence intervals for variances or standard deviations.

Bonett (2006) proposed a confidence interval that performs well in small samples under moderate departures from normality. His interval performs only slightly worse than the exact normal-based confidence interval when sampling from a normal distribution. A larger sample size provides Bonett confidence intervals with greater protection against nonnormality.

▷ Example 9: Bonett confidence interval for normal data

We will repeat example 8 and construct a Bonett confidence interval for the standard deviation by specifying the bonett option. The results are similar, and both examples lead to the same inferential conclusion.

```
. ci variances weight, sd bonett
                                                    Bonett
    Variable |        Obs      Std. dev.      [95% conf. interval]
-------------+-----------------------------------------------------
      weight |          8      .6235711       .3997041    1.288498
```

The Bonett confidence interval is wider than the normal-based confidence interval in example 8. For normal data, Bonett (2006) suggested that if Bonett confidence interval is used for a sample of size $n+3$, then its average width will be about the same as the average width of the normal-based confidence interval from a sample size of n. Sampling three more observations may be a small price to pay because Bonett confidence intervals perform substantially better than the normal-based confidence intervals for nonnormal data.

◁

> ## Example 10: Bonett confidence interval for nonnormal data

The following data have been generated from a t distribution with five degrees of freedom to illustrate the effect of wrongfully using the normal-based confidence interval when the data-generating process is not normal.

```
. use https://www.stata-press.com/data/r19/peas_tdist
(Weights of canned peas, t distribution)
. ci variances weight, sd
```

Variable	Obs	Std. dev.	[95% conf. interval]	
weight	8	2.226558	1.472143	4.531652

The standard deviation of a t distribution with five degrees of freedom is $\sqrt{5/3} \approx 1.29$ and falls outside the confidence interval limits. If we suspect that data may not be normal, the Bonett confidence interval is typically a better choice:

```
. ci variances weight, sd bonett
```

Variable	Obs	Std. dev.	Bonett [95% conf. interval]	
weight	8	2.226558	1.137505	5.772519

The value 1.29 is within the limits of the Bonett confidence interval $[\,1.14, 5.77\,]$

◁

Immediate form

So far, we computed confidence intervals for various parameters using data in memory. We can also compute confidence intervals using only data summaries, without any data in memory. Each of the considered `ci` commands has an immediate `cii` version that computes the respective confidence intervals using data summaries.

> ## Example 11: Confidence interval for a normal mean

We are reading a soon-to-be-published paper by a colleague. In it is a table showing the number of observations, mean, and standard deviation of the 1980 median family income for the Northeast and West. We correctly think that the paper would be much improved if it included the confidence intervals. The paper claims that for 166 cities in the Northeast, the average of median family income is \$19,509 with a standard deviation of \$4,379:

For the Northeast:

```
. cii means 166 19509 4379
```

Variable	Obs	Mean	Std. err.	[95% conf. interval]	
	166	19509	339.8763	18837.93	20180.07

For the West:

```
. cii means 256 22557 5003
```

Variable	Obs	Mean	Std. err.	[95% conf. interval]	
	256	22557	312.6875	21941.22	23172.78

◁

▷ Example 12: Confidence interval for a Poisson mean

The number of reported traffic accidents in Santa Monica over a 24-hour period is 27. We need know nothing else to compute a confidence interval for the mean number of accidents for a day:

```
. cii means 1 27, poisson

                                               Poisson exact
    Variable |  Exposure       Mean    Std. err.    [95% conf. interval]
             |
             |         1         27    5.196152      17.79317    39.28358
```

◁

▷ Example 13: Confidence interval for a proportion

We flip a coin 10 times, and it comes up heads only once. We are shocked and decide to obtain a 99% confidence interval for this coin:

```
. cii proportions 10 1, level(99)

                                               Binomial exact
    Variable |    Obs  Proportion   Std. err.    [99% conf. interval]
             |
             |     10         .1    .0948683      .0005011    .5442871
```

◁

▷ Example 14: Confidence interval for a variance

A company fills 32-ounce tomato juice jars with a quantity of juice having a normal distribution with a claimed variance not exceeding 0.2. A random sample of 15 jars is collected to evaluate this claim. The sample variance is 0.5:

```
. cii variances 15 0.5
    Variable |    Obs     Variance       [95% conf. interval]
             |
             |     15           .5       .2680047    1.243621
```

Because the advertised value of 0.2 does not fall inside the confidence interval, the company is allowing too much variation in the amount of tomato juice per jar.

◁

▷ Example 15: Confidence interval for a standard deviation

Suppose the director of statistical development at a statistical software company is a big soccer fan and requires all developers to play on the company team in the city's local soccer league. Ten developers are randomly selected to participate in the game. To ensure an advantage over other teams, the director requires each of the 10 developers to cover 6 miles on average each game. Being merciful, she will tolerate a standard deviation of 0.3 miles across different players, arguing that this will keep the team's performance consistent. The distance covered by each player is measured using a pedometer. At the end of the game, the sample standard deviation of the distances covered by the 10 players was 0.56 miles:

```
. cii variances 10 0.56, sd
    Variable |    Obs     Std. dev.       [95% conf. interval]
             |
             |     10          .56       .3851877    1.022342
```

Because the confidence interval does not include the designated value for the standard deviation, 0.3 miles, it is clear the team is not meeting standards, and an unpleasant meeting is planned.

◁

▷ Example 16: Confidence interval for a standard deviation of nonnormal data

Continuing with example 15, a clever statistician points out that distances covered by company players in a soccer match do not follow the normal distribution because some players, mostly econometricians, walk on the field, while others, mostly statisticians, do all the running. Therefore, the normal-based confidence interval (which assumes normality) is not valid. Instead, we should use the Bonett confidence interval, which additionally requires an estimate of kurtosis; see *Methods and formulas*. If kurtosis is estimated to be 5, we would obtain the following:

```
. cii variances 10 0.56 5, sd bonett
```

			Bonett	
Variable	Obs	Std. dev.	[95% conf. interval]	
	10	.56	.2689449	1.45029

The Bonett confidence interval now contains the specified value for the standard deviation, 0.3 miles. The director of statistics concludes that overall team performance is acceptable. An uncomfortable meeting is still planned but for a smaller group.

◁

Stored results

ci means and cii means store the following in r():

Scalars
r(N)	number of observations or, if poisson is specified, exposure
r(mean)	mean
r(se)	estimate of standard error
r(lb)	lower bound of confidence interval
r(ub)	upper bound of confidence interval
r(level)	confidence level of confidence interval

Macros
r(citype)	normal or poisson; type of confidence interval
r(exposure)	name of exposure variable with poisson

ci proportions and cii proportions store the following in r():

Scalars
r(N)	number of observations
r(proportion)	proportion
r(se)	estimate of standard error
r(lb)	lower bound of confidence interval
r(ub)	upper bound of confidence interval
r(level)	confidence level of confidence interval

Macros
r(citype)	exact, wald, wilson, agresti, or jeffreys; type of confidence interval

ci variances and cii variances store the following in r():

Scalars
r(N)	number of observations
r(Var)	variance

`r(sd)`	standard deviation, if `sd` is specified
`r(kurtosis)`	kurtosis, only if `bonett` is specified
`r(lb)`	lower bound of confidence interval
`r(ub)`	upper bound of confidence interval
`r(level)`	confidence level of confidence interval

Macros

`r(citype)`	`normal` or `bonett`, type of confidence interval

Methods and formulas

Methods and formulas are presented under the following headings:

> Normal mean
> Poisson mean
> Binomial proportion
> Variance and standard deviation

Normal mean

Define n, $\overline{x}$, and s^2 as, respectively, the number of observations, (weighted) average, and (unbiased) estimated variance of the variable in question; see [R] **summarize**.

The standard error of the mean, s_μ, is defined as $\sqrt{s^2/n}$.

Let α be $1 - l/100$, where l is the confidence level specified by the user in the `level()` option. Define $t_{\alpha/2}$ as the two-sided t statistic corresponding to a significance level of α with $n-1$ degrees of freedom; $t_{\alpha/2}$ is obtained from Stata as `invttail(`n`-1,0.5*`α`)`. The lower and upper confidence bounds are, respectively, $\overline{x} - s_\mu t_{\alpha/2}$ and $\overline{x} + s_\mu t_{\alpha/2}$.

Poisson mean

Given the total cases, k, the estimate of the expected count λ is k, and its standard error is $\sqrt{k}$. `ci means` with the `poisson` option calculates the exact confidence interval $[\lambda_1, \lambda_2]$ such that

$$\Pr(K \geq k | \lambda = \lambda_1) = \alpha/2$$

and

$$\Pr(K \leq k | \lambda = \lambda_2) = \alpha/2$$

where K is Poisson with mean λ. Solution is obtained by Newton's method. If $k = 0$, the calculation of λ_1 is skipped. All values are then reported as rates, which are the above numbers divided by the total exposure.

Binomial proportion

Given k successes of n trials, the estimated probability of a success is $\hat{p} = k/n$ with standard error $\sqrt{\hat{p}(1-\hat{p})/n}$. `ci` calculates the exact (Clopper–Pearson) confidence interval $[p_1, p_2]$ such that

$$\Pr(K \geq k | p = p_1) = \alpha/2$$

and

$$\Pr(K \leq k | p = p_2) = \alpha/2$$

where K is distributed as binomial(n, p). The endpoints may be obtained directly by using Stata's `invbinomial()` function. If $k = 0$ or $k = n$, the calculation of the appropriate tail is skipped.

The Wald interval is $\hat{p} \pm z_{\alpha/2}\sqrt{\hat{p}(1 - \hat{p})/n}$, where $z_{\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard normal. The interval is obtained by inverting the acceptance region of the large-sample Wald test of $H_0: p = p_0$ versus the two-sided alternative. That is, the confidence interval is the set of all p_0 such that

$$\left| \frac{\hat{p} - p_0}{\sqrt{n^{-1}\hat{p}(1 - \hat{p})}} \right| \leq z_{\alpha/2}$$

The Wilson interval is a variation on the Wald interval, using the null standard error $\sqrt{n^{-1}p_0(1 - p_0)}$ in place of the estimated standard error $\sqrt{n^{-1}\hat{p}(1 - \hat{p})}$ in the above expression. Inverting this acceptance region is more complicated yet results in the closed form

$$\frac{k + z_{\alpha/2}^2/2}{n + z_{\alpha/2}^2} \pm \frac{z_{\alpha/2}n^{1/2}}{n + z_{\alpha/2}^2} \left\{ \hat{p}(1 - \hat{p}) + \frac{z_{\alpha/2}^2}{4n} \right\}^{1/2}$$

The Agresti–Coull interval is basically a Wald interval that borrows its center from the Wilson interval. Defining $\tilde{k} = k + z_{\alpha/2}^2/2$, $\tilde{n} = n + z_{\alpha/2}^2$, and (hence) $\tilde{p} = \tilde{k}/\tilde{n}$, the Agresti–Coull interval is

$$\tilde{p} \pm z_{\alpha/2}\sqrt{\tilde{p}(1 - \tilde{p})/\tilde{n}}$$

When $\alpha = 0.05$, $z_{\alpha/2}$ is near enough to 2 that $\tilde{p}$ can be thought of as a typical estimate of proportion where two successes and two failures have been added to the sample (Agresti and Coull 1998). This typical estimate of proportion makes the Agresti–Coull interval an easy-to-present alternative for introductory statistics students.

The Jeffreys interval is a Bayesian credible interval and is based on the Jeffreys prior, which is the Beta$(1/2, 1/2)$ distribution. Assigning this prior to p results in a posterior distribution for p that is Beta with parameters $k + 1/2$ and $n - k + 1/2$. The Jeffreys interval is then taken to be the $1 - \alpha$ central posterior probability interval, namely, the $\alpha/2$ and $1 - \alpha/2$ quantiles of the Beta$(k + 1/2, n - k + 1/2)$ distribution. These quantiles may be obtained directly by using Stata's `invibeta()` function. See [BAYES] **bayesstats summary** for more details about credible intervals.

Variance and standard deviation

Let $X_1, \ldots, X_n$ be a random sample and assume that $X_i \sim N(\mu, \sigma^2)$. Because $(n-1)s^2/\sigma^2 \sim \chi^2_{n-1}$, we have $\Pr\{\chi^2_{n-1,\alpha/2} \leq (n-1)s^2/\sigma^2 \leq \chi^2_{n-1,1-\alpha/2}\} = 1 - \alpha$, where $\chi^2_{n-1,\alpha/2}$ and $\chi^2_{n-1,1-\alpha/2}$ are the $\alpha/2$ and $1 - \alpha/2$ quantiles of the χ^2_{n-1} distribution. Thus, the normal-based confidence interval for the population variance σ^2 with $100(1-\alpha)\%$ confidence level is given by

$$I_{\text{normal}} = \left[\frac{(n-1)s^2}{\chi^2_{n-1,1-\alpha/2}}, \frac{(n-1)s^2}{\chi^2_{n-1,\alpha/2}} \right]$$

$\chi^2_{n-1,1-\alpha/2}$ and $\chi^2_{n-1,\alpha/2}$ are obtained from Stata as `invchi2tail(n-1,0.5*α)` and `invchi2(n-1,0.5*α)`, respectively.

The normal-based confidence interval is very sensitive to minor departures from the normality assumption, and its performance does not improve with increasing sample size. For scenarios in which the population distribution is not normal, the actual coverage probability of the normal-based confidence interval can be drastically lower than the nominal confidence level α.

Bonett (2006) proposed an alternative to the normal-based confidence interval that is nearly exact under normality and has coverage probability close to $1 - \alpha$ under moderate nonnormality. It also has $1 - \alpha$ asymptotic coverage probability for nonnormal distributions with finite fourth moment. Instead of assuming that $X_i \sim N(\mu, \sigma^2)$, Bonett's approach requires continuous i.i.d. random variables with finite fourth moment. The variance of s^2 may be expressed as $\sigma^4\{\gamma_4 - (n-3)/(n-1)\}/n$ (see Casella and Berger [2002, ex. 5.8, 257]), where $\gamma_4 = \mu^4/\sigma^4$ is the kurtosis and $\mu^4 = E(X_i - \mu)^4$ is the population fourth central moment. The variance-stabilizing transformation $\ln(s^2)$ and the delta method can be used to construct an asymptotic $100(1-\alpha)\%$ confidence interval for σ^2,

$$\left[\exp\left\{ \ln(s^2) - z_{\alpha/2}se \right\}, \; \exp\left\{ \ln(s^2) + z_{\alpha/2}se \right\} \right]$$

where $se = \{\hat{\gamma}_4 - (n-3)/(n-1)\}/n \approx \text{Var}\{\ln(s^2)\}$ and $\hat{\gamma}_4$ is an estimate of the kurtosis. Bonett introduced three adjustments to improve the small-sample properties of the above confidence interval. First, he swapped the inner and outer denominator in the expression for se and changed it to $\{\hat{\gamma}_4 - (n-3)/n\}/(n-1)$. This was suggested by Shoemaker (2003) who used it to improve the small-sample performance of his variance test. Second, with regard to the estimation of kurtosis, Bonett proposed $\hat{\gamma}_4 = n\sum(X_i - m)^4/\left\{\sum(X_i - \overline{X})^2\right\}^2$, where m is a trimmed mean with a trim-proportion equal to $1/\{2(n-4)^{1/2}\}$. This kurtosis estimator reduces the negative bias in symmetric and skewed heavy-tailed distributions. Last, he empirically derived a small-sample correction factor $c = n/(n - z_{\alpha/2})$ that helps equalize the tail probabilities. These modifications yield

$$I_{\text{Bonett}} = \left[\exp\left\{ \ln(cs^2) - z_{\alpha/2}se \right\}, \; \exp\left\{ \ln(cs^2) + z_{\alpha/2}se \right\} \right]$$

where $z_{\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard normal and $se = c\left[\{\hat{\gamma}_4 - (n-3)/n\}/(n-1)\right]$.

Taking the square root of the endpoints of both intervals gives confidence intervals for the standard deviation σ.

Edwin Bidwell (E. B.) Wilson (1879–1964) majored in mathematics at Harvard and studied and taught at Yale and MIT before returning to Harvard in 1922. He worked in mathematics, physics, and statistics. His method for binomial intervals can be considered a precursor, for a particular problem, of Neyman's concept of confidence intervals.

Jerzy Neyman (1894–1981) was born in Bendery, Russia, now Moldavia. He studied and then taught at Kharkov University, moving from physics to mathematics. In 1921, Neyman moved to Poland, where he worked in statistics at Bydgoszcz and then Warsaw. Neyman received a Rockefeller Fellowship to work with Karl Pearson at University College London. There he collaborated with Egon Pearson, Karl's son, on the theory of hypothesis testing. Life in Poland became progressively more difficult, and Neyman returned to UCL to work there from 1934 to 1938. At this time, he published on the theory of confidence intervals. He then was offered a post in California at Berkeley, where he settled. Neyman established an outstanding statistics department and remained highly active in research, including applications in astronomy, meteorology, and medicine. He was one of the great statisticians of the 20th century.

Acknowledgment

References

Agresti, A., and B. A. Coull. 1998. Approximate is better than "exact" for interval estimation of binomial proportions. *American Statistician* 52: 119–126. https://doi.org/10.1080/00031305.1998.10480550.

Bonett, D. G. 2006. Approximate confidence interval for standard deviation of nonnormal distributions. *Computational Statistics and Data Analysis* 50: 775–782. https://doi.org/10.1016/j.csda.2004.10.003.

Brown, L. D., T. T. Cai, and A. DasGupta. 2001. Interval estimation for a binomial proportion. *Statistical Science* 16: 101–133. https://doi.org/10.1214/ss/1009213286.

Campbell, M. J., D. Machin, and S. J. Walters. 2007. *Medical Statistics: A Textbook for the Health Sciences*. 4th ed. Chichester, UK: Wiley.

Casella, G., and R. L. Berger. 2002. *Statistical Inference*. 2nd ed. Pacific Grove, CA: Duxbury.

Clopper, C. J., and E. S. Pearson. 1934. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika* 26: 404–413. https://doi.org/10.1093/biomet/26.4.404.

Cook, A. 1990. Sir Harold Jeffreys, 2 April 1891–18 March 1989. *Biographical Memoirs of Fellows of the Royal Society* 36: 303–333. https://doi.org/10.1098/rsbm.1990.0034.

Earnest, A. 2017. *Essentials of a Successful Biostatistical Collaboration*. Boca Raton, FL: CRC Press.

Jeffreys, H. 1946. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London,* A ser., 186: 453–461. https://doi.org/10.1098/rspa.1946.0056.

Lash, T. L., T. J. VanderWeele, S. Haneuse, and K. J. Rothman. 2021. *Modern Epidemiology*. 4th ed. Philadelphia: Wolters Kluwer.

Lindley, D. V. 2001. "Harold Jeffreys". In *Statisticians of the Centuries*, edited by C. C. Heyde and E. Seneta, 402–405. New York: Springer. https://doi.org/10.1007/978-1-4613-0179-0_86.

Miller, R. G., Jr. 1997. *Beyond ANOVA: Basics of Applied Statistics*. London: Chapman and Hall.

Reid, C. 1982. *Neyman—from Life*. New York: Springer. https://doi.org/10.1007/978-1-4612-5754-7_1.

Scheffé, H. 1959. *The Analysis of Variance*. New York: Wiley.

Shoemaker, L. H. 2003. Fixing the F test for equal variances. *American Statistician* 57: 105–114. https://doi.org/10.1198/0003130031441.

Stigler, S. M. 1997. "Wilson, Edwin Bidwell". In *Leading Personalities in Statistical Sciences: From the Seventeenth Century to the Present*, edited by N. L. Johnson and S. Kotz, 344–346. New York: Wiley. https://doi.org/10.1002/9781118150719.ch107.

Ward, B. W. 2019. kg_nchs: A command for Korn–Graubard confidence intervals and National Center for Health Statistics' Data Presentation Standards for Proportions. *Stata Journal* 19: 510–522.

Wilson, E. B. 1927. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association* 22: 209–212. https://doi.org/10.2307/2276774.

Also see

[R] **ameans** — Arithmetic, geometric, and harmonic means

[R] **bitest** — Binomial probability test

[R] **centile** — Report centile and confidence interval

[R] **prtest** — Tests of proportions

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

[R] **summarize** — Summary statistics

[R] **ttest** — t tests (mean-comparison tests)

[D] **pctile** — Create variable containing percentiles