

1-sample Confidence Intervals (CI)

Module 8

Statistics 251: Statistical Methods

Updated 2021

Introduction

Statistical Inference: using information obtained from a proper sample to make an educated judgment about a population

Three types of inference

- (1) point estimation
- (2) interval estimation (aka confidence intervals (CIs))
- (3) statistical tests (aka hypothesis tests)

Terms I

Point Estimation: Point estimators are functions that are used to obtain an approximate value of a population parameter from random samples of the population. They use the sample data of a population to calculate a point estimate or a statistic that serves as the best estimate of an unknown parameter of a population. The selected statistic is called the **point estimate** of the parameter. Examples are: \bar{X} estimates μ , $\hat{\pi}$ estimates π , s^2 estimates σ^2 , and so on.

A point estimate is just a single number and by itself provides no information about the precision and reliability of estimation; it gives no feedback on how close our estimate was to the parameter

Properties of point estimators

- (1) *$E(\hat{\theta}) = \theta$ or Bias:* the bias of a point estimator is defined as the difference between the expected value of the estimator and the value of the parameter being estimated. When the estimated value of the parameter and the value of the parameter being estimated are equal, the estimator is considered unbiased
- (2) *Consistency:* tells us how close the point estimator stays to the value of the parameter as it increases in size. The point estimator requires a large sample size for it to be more consistent and accurate
- (3) *Efficiency:* the most efficient point estimator is the one with the smallest variance of all the unbiased and consistent estimators. The variance measures the level of dispersion from the estimate, and the smallest variance should vary the least from one sample to the other.

Generally, the efficiency of the estimator depends on the distribution of the population. For example, in a normal distribution, the mean is considered more efficient than the median, but the same does not apply in asymmetrical distributions

Terms II

Interval estimation: an alternative to reporting a sensible value for the parameter being estimated is to calculate an entire interval of plausible values, called **interval estimation**, specifically we call them **confidence intervals (CIs)**.

A CI is an estimate of an interval in statistics that may contain a population parameter. The unknown population parameter is found through a sample parameter calculated from the sampled data. For example, the population mean μ is found using the sample mean \bar{X} .

The interval is generally defined by its lower and upper bounds. The CI is expressed as a percentage (the most frequently quoted percentages are 90%, 95%, and 99%). The percentage reflects the confidence level (CL)

Assumptions

Assumptions: conditions that should be true in order for the data to properly fit the model used for estimations (1) Independence: observations are independent from one another (2) Randomization: proper randomization was used (takes care of independence issue if there is one) (3) Means need an *approximate* normal distribution (if $n \geq 30$, then it is approximately normal), and proportions need to meet the S/F condition: $np \geq 5$ and $nq \geq 5$ (if $n \geq 60$, S/F condition is met)

If assumptions are violated, the results from the analyses are not as valid nor reliable

General Form

All CIs (even more complex ones) have the same form:

$$\text{point estimate} \pm \text{bound}$$

Where the bound on the error of estimation is $z^*(se)$ or $t^*(se)$ and $se_{mean} = \frac{\sigma}{\sqrt{n}}$, $se_{\pi} = \sqrt{\hat{\pi}(1-\pi)/n}$, or $se_{mean} = \frac{s}{\sqrt{n}}$ (the one you use depends on the situation; explanations to come)

Form for CI on μ when σ is known

$$\bar{X} \pm z^*(se_{mean})$$

Page 443 and 445 in the textbook show examples on how to do these calculations on TI calculators.

To solve for a sample size given a CL and bound:


$$n = \left(\frac{z^* \sigma}{\text{bound}} \right)^2$$

Unlike normal rounding conventions, we *always round up*

Finding z^*

z^* is a critical value of z based on the confidence level. $1 - CL = \alpha$ where α is the “left-over” area

Standard Normal



How to find z^*


Using statdistributions.com:

- (1) There are 4 distributions; the standard normal (z) is default
- (2) Subtract the confidence level (CL) from 1: $1 - CL$
- (3) $1 - CL = \alpha$, where α is the left over area, called pvalue on website
- (4) Input pvalue = α
- (5) Select two-tail (2T)

z^* for 90% CI

90%: $z_{90\%}$ $\alpha = 1 - 0.9 = 0.1$, 2T (which is in between two values); the z-score for 90% CI is: ± 1.645
pvalue=.1,2T


90% CI



z^* for 95% CI

95%: $z_{95\%}$ $\alpha = 1 - 0.95 = 0.05$, 2T; $z = \pm 1.96$ pvalue=.05,2T


95% CI



z^* for 99% CI

99%: $z_{99\%} = 1 - 0.99 = 0.01$, 2T (which is in between two values); the z -score for 99% CI is: ± 2.575
pvalue=.01,2T

99% CI



Generic Interpretation of CI

To interpret the CI, the statement we use discusses the level of confidence, the parameter we are trying to estimate, and the values of the interval.

“We are (CL)% confident the true (mean, proportion, etc.) of (insert context) is between (lower number) and (upper number) (units of measurement).”

Estimating μ when σ is known

A consumer testing agency wants to evaluate the claim made by a manufacturer of discount tires; the claim is that its tires can be driven at least 35,000 miles before wearing out. Assume that these tires have a normal distribution and its standard deviation is 5,000 miles.

To determine the average number of miles that can be obtained from the tires, the agency randomly selects 60 tires from the manufacturer’s warehouse and places the tires on 15 similar cars driven by test drivers on a 2-mile oval track, with sample mean of 31.47. (The mean was scaled down from 31,470 miles for ease of

calculations but will NOT change the answer). $\bar{X} \sim N\left(31.47, \frac{5}{\sqrt{60}}\right) = \bar{X} \sim N(31.47, 0.646)$ (we could do it with the values in thousands, $\bar{X} \sim N(31470, 645.5)$, and other than rounding errors, there would be no difference)

Examples for μ with known σ

- (1) Estimate μ , the true average lifetime miles of these tires with 90% confidence, interpret
- (2) Estimate μ , the true average lifetime miles of these tires with 95% confidence, interpret
- (3) Estimate μ , the true average lifetime miles of these tires with 98% confidence, interpret
- (4) Calculate the necessary sample size for a new study that has 95% CL and a bound of 1.05 miles

90% CI for μ example (1)

$$31.47 \pm (1.645)(0.646) = 31.47 \pm 1.063 = (30.41, 32.53)$$

Interpretation: We are 90% confident the true average lifetime of these tires is between 30,410 and 32,530 miles.

95% CI for μ example (2)

$$31.47 \pm (1.96)(0.646) = 31.47 \pm 1.266 = (30.2, 32.74)$$

Interpretation: We are 95% confident the true average lifetime of these tires is between 30,200 and 32,740 miles.

99% CI for μ example (3)

$$31.47 \pm (2.575)(0.646) = 31.47 \pm 1.663 = (29.81, 33.13)$$

Interpretation: We are 99% confident the true average lifetime of these tires is between 29,810 and 33,130 miles.

Calculate the necessary sample size for a new study that has 95% CL and a bound of 1.05

$bound = 1.05$, $z^* = 1.96$, and $\sigma = 5$

$$n = \left(\frac{z^* \sigma}{bound}\right)^2 = \left(\frac{(1.96)(5)}{1.05}\right)^2 = 87.11 \rightarrow 88$$

CI for π , proportion

$$\hat{\pi} \pm bound = \hat{\pi} \pm z^*(se_{\hat{\pi}})$$

Where $bound = z^*(se)$ and $se_{\hat{\pi}} = \sqrt{\hat{\pi}(1-\hat{\pi})/n}$, z^* comes from the table the same way as previous example with μ CI

To solve for a sample size given a CL and bound:

$$n = [\hat{\pi}(1-\hat{\pi})] \left(\frac{z^*}{bound}\right)^2$$

When no previous information about the proportion should be, use $\hat{\pi} = 0.5$. Again, unlike normal rounding conventions, we *always round up*

Estimating π

In a given town, a random sample of 892 voters contained 500 people who favored a particular bond issue.

- (1) Find the true proportion of voters who favor the particular bond with 90% confidence
- (2) If another sample is to be taken later, how many people should be sampled to be within 0.02 of the estimate?
- (3) Suppose for another sample, there is no previous information about the proportion of people who would favor the bond issue. How many people should be sampled to be within 0.02 of the estimate?

90% CI for π example

$\hat{\pi} = \frac{X}{n}$ where X is the count of successes and n is the sample size. $\frac{500}{892} = 0.561$

$$se_{\hat{\pi}} = \sqrt{\hat{\pi}(1 - \hat{\pi})/n} = \sqrt{(0.561)(1 - 0.561)/892} = 0.0166$$

$$0.561 \pm (1.645)(0.0166) = 0.561 \pm 0.0273 = (0.5337, 0.5883)$$

Interpretation: We are 90% confident the true proportion of voters who favor the particular bond issue is between 53.37% and 58.83%.

Calculate new sample size with 90% CL and bound of 0.02

$bound = 0.02$, $z^* = 1.645$, and $\hat{\pi} = 0.561$

$$\begin{aligned} n &= \hat{\pi}(1 - \hat{\pi}) \left(\frac{z^*}{bound} \right)^2 = (0.561)(1 - 0.561) \left(\frac{(1.645)}{0.02} \right)^2 \\ &= 1666.093 \rightarrow 1667 \end{aligned}$$

Calculate new sample size with 90% CL and bound of 0.02, no information about π

$bound = 0.02$, $z^* = 1.645$, and with no known information about the proportion, use $\hat{\pi} = 0.5$

$$\begin{aligned} n &= \hat{\pi}(1 - \hat{\pi}) \left(\frac{z^*}{bound} \right)^2 = (0.5)(1 - 0.5) \left(\frac{(1.645)}{0.02} \right)^2 \\ &= 1691.266 \rightarrow 1692 \end{aligned}$$

Student's t Distribution

What happens when we do not know the standard deviation (or variance), or the sample size is very small? Sometimes the z distribution may not be conservative enough. There is a distribution that is similar to the z distribution but made more for distributions with heavier tails (more area in the tails than z).

Statistical beer?

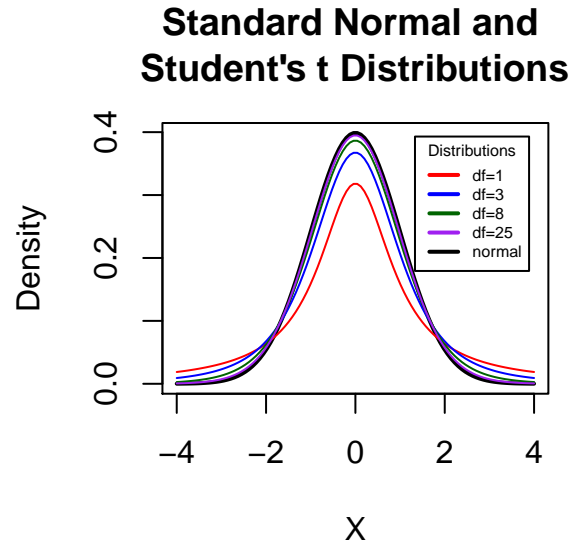
It is called Student's t distribution. It was created by a quality control manager at Guinness Beer named Gosset in 1908. He worked in quality control and worked with small samples. He couldn't publish the results of his study under his true name because of his work contract; no company secrets to be given out under penalty of law.

t logistics


t needs two things: (1) *df*, degrees of freedom ($df = n - 1$), and (2) $\alpha = 1 - CL$

Degrees of freedom arises from the fact that we are now using an estimate, \bar{X} , for the true mean, μ . We have lost a bit of information, and now have one less degree of freedom. For more detailed description, see Lecture link for this module on class website (links provided there)

Student's *t* graphs



z and *t* with varying *df*



How to find t^*


Using statdistributions.com:

- (1) Choose **Student's t**
- (2) Subtract the confidence level (CL) from 1: $1 - CL$
- (3) $1 - CL = \alpha$, where α is the left over area, called **pvalue** on website
- (4) Input **pvalue**= α

- (5) Input $df = df$
- (6) Select **two-tail** (2T)
- (7) t is symmetric so that the values can be negative for the other side of the curve (below the mean)


t^* for 90% CI

90%: $t_{90\%,df}$ $1 - 0.9 = 0.1$, $df = 19$, 2T. $t^* = t_{90\%,19} = 1.729$ $pvalue=.1,df=19,2T$




t^* for 95% CI

95%: $t_{95\%,df}$ $1 - 0.95 = 0.05$, $df = 19$, 2T. $t^* = t_{95\%,19} = 2.093$ $pvalue=.05,df=19,2T$



t^* for 99% CI

99%: $t_{99\%,df}$ $1 - 0.99 = 0.01$, $df = 19$, 2T. $t^* = t_{99\%,19} = 2.861$ $pvalue=.01,df=19,2T$



Form for CI on μ when σ is not known (we use s)

$$\bar{X} \pm t^*(se_{mean})$$

Where $t^* = t_{\alpha/2, df}$, $df = n - 1$, and $se_{mean} = \frac{s}{\sqrt{n}}$

Page 453+ in the textbook show examples on how to do these calculations on TI calculators.

Estimating μ when σ is not known, using s

A product comes in cans labeled “38 oz”, and a random sample of 10 cans had the following weights: {34.6, 39.65, 34.75, 40, 39.5, 38.9, 34.25, 36.8, 39, 37.2}

```
rbind(mean, sd)
```

```
      [,1]
mean 37.465000
sd   2.265324
```

Estimate μ , the true average weight of the product, with 98% confidence, interpret

95% CI for μ when σ is unknown

$t^* = t_{98\%, df}$ where $df = n - 1 = 10 - 1 = 9$ so $t_{98\%, 9} = t_{0.02/2, 9} = 2.821$, $se_{mean} = \frac{s}{\sqrt{n}} = \frac{2.265}{\sqrt{10}} = 0.716$

$$37.465 \pm (2.821)(0.716) = 37.465 \pm 2.0198 = (35.45, 39.48)$$

Interpretation: We are 98% confident the true average weight of product is between 35.45 and 39.48 ounces.

Reality

In practice, the use of Z seldom happens. You have to know a lot about the population and such ahead of time and that does not happen often. Statisticians primarily use t since it is good with data samples, small sample sizes, and if the sample size happens to be large enough, t will eventually converge to z so use of t is just better. It is a little more conservative so you may not get as many false positives that way (which is good).

Rather than do all the calculations by hand, most of what is done is interpreting output. You NOT will be responsible for running a software program and coding; rather you will read and interpret output from a program, R, that is provided for you.

Functions in R for CIs

Examples shown will use `t.test()` for all CIs for \bar{X} (regardless of whether σ is known or not) and $\hat{\pi}$

- (1) Tires (CI for mean with σ known; still use t)
- (2) Voting (CI for proportions, using t)
- (3) Product weights (CI for mean with σ unknown; use t)

Tires output 90% CI

```
t.test(tires,conf.level=.9)
```

```
One Sample t-test
```

```
data: tires
t = 47.274, df = 59, p-value < 2.2e-16
alternative hypothesis: true mean is not equal to 0
90 percent confidence interval:
 31328.30 33624.34
sample estimates:
mean of x
 32476.32
```

With 90% confidence, the true mean mileage of the tires is between 31,328.3 and 33,624.34 miles.

Tires output 95% CI

```
t.test(tires)
```

```
One Sample t-test
```

```
data: tires
t = 47.274, df = 59, p-value < 2.2e-16
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 31101.66 33850.98
sample estimates:
mean of x
 32476.32
```

With 95% confidence, the true mean mileage of the tires is between 31,101.66 and 33,850.98 miles.

Tires output 98% CI

```
t.test(tires,conf.level=.98)
```

```
One Sample t-test
```

```
data: tires
t = 47.274, df = 59, p-value < 2.2e-16
alternative hypothesis: true mean is not equal to 0
```

```
98 percent confidence interval:
 30833.58 34119.06
sample estimates:
mean of x
 32476.32
```

With 98% confidence, the true mean mileage of the tires is between 30,833.58 and 34,119.06 miles.

Voting output 90% CI

```
t.test(votes,conf.level=.9)
```

```
One Sample t-test
```

```
data: votes
t = 33.712, df = 891, p-value < 2.2e-16
alternative hypothesis: true mean is not equal to 0
90 percent confidence interval:
 0.5331600 0.5879162
sample estimates:
mean of x
0.5605381
```

With 90% confidence, the true proportion of voters who favor the bond is between 53.32% and 58.79%.

Voting output 95% CI

```
t.test(votes)
```

```
One Sample t-test
```

```
data: votes
t = 33.712, df = 891, p-value < 2.2e-16
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 0.5279047 0.5931715
sample estimates:
mean of x
0.5605381
```

With 95% confidence, the true proportion of voters who favor the bond is between 52.79% and 59.32%.

Voting output 98% CI

```
t.test(votes,conf.level=.98)
```

```
One Sample t-test
```

```
data: votes
t = 33.712, df = 891, p-value < 2.2e-16
alternative hypothesis: true mean is not equal to 0
98 percent confidence interval:
 0.5217873 0.5992889
sample estimates:
mean of x
```

0.5605381

With 98% confidence, the true proportion of voters who favor the bond is between 52.18% and 59.93%.

Product weight output 90% CI

```
t.test(product,conf.level=.9)
```

One Sample t-test

```
data: product
t = 52.299, df = 9, p-value = 1.716e-12
alternative hypothesis: true mean is not equal to 0
90 percent confidence interval:
 36.15183 38.77817
sample estimates:
mean of x
 37.465
```

With 90% confidence, the true average weight of product is between 36.2 and 38.8 ounces.

Product weight output 95% CI

```
t.test(product)
```

One Sample t-test

```
data: product
t = 52.299, df = 9, p-value = 1.716e-12
alternative hypothesis: true mean is not equal to 0
95 percent confidence interval:
 35.84448 39.08552
sample estimates:
mean of x
 37.465
```

With 95% confidence, the true average weight of product is between 35.8 and 39.1 ounces.

Product weight output 99% CI

```
t.test(product,conf.level=.99)
```

One Sample t-test

```
data: product
t = 52.299, df = 9, p-value = 1.716e-12
alternative hypothesis: true mean is not equal to 0
99 percent confidence interval:
 35.13695 39.79305
sample estimates:
mean of x
 37.465
```

With 99% confidence, the true average weight of product is between 35.1 and 39.8 ounces.