STAT 511

Lecture 13: Additional Confidence Intervals' Related Topics

Devore: Section 7.3-7.4

Prof. Michael Levine

October 24, 2018

t-confidence intervals

- Large-sample confidence intervals are based on the fact that, for n large enough,

$$Z = \frac{\bar{X} - \mu}{S/\sqrt{n}}$$

  is approximately normally distributed
- But what if $n < 40$?
- For small n, this test statistic is denoted

$$T = \frac{\bar{X} - \mu}{S/\sqrt{n}}$$

  to stress the fact it is no longer normally distributed

t Distribution

- A t distribution is governed by one parameter ν which is called the number of degrees of freedom (df)
- Properties:
  1. t_ν curve is bell-shaped and centered at 0
  2. It has heavier tails than normal distribution (more spread out)
  3. As $\nu \to \infty$, the t_ν density curve approaches the normal curve

- Let $t_{\alpha,\nu}$ be the number on the horizontal axis such that the area to the left of it under t_ν curve is α; $t_{\alpha,\nu}$ is a *t critical value*.
- For fixed ν, $t_{\alpha,\nu}$ increases as α decreases
- For fixed α, as ν increases, the value $t_{\alpha,\nu}$ decreases. The process slows down as ν increases; that is why the table values are shown in increments of 2 between 30 df and 40 df, but then jump to $\nu = 50$, $\nu = 60$ etc.
- z_α is the last row of the table since t_∞ is the standard normal distribution

Figure :

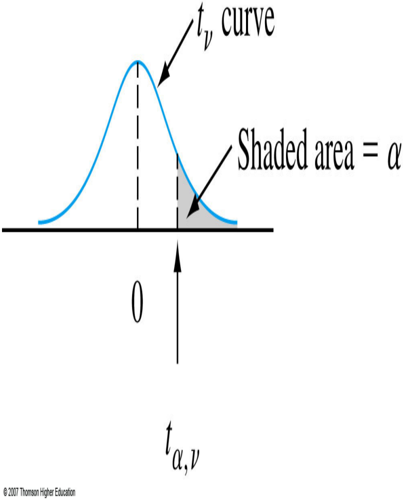
t_ν curve
Shaded area = α
0
$t_{\alpha,\nu}$
© 2007 Thomson Higher Education

One-sample t confidence interval

- The number of df for T is $n-1$ since S is based on deviations $X_1 - \bar{X}, \ldots, X_n - \bar{X}$ that add up to zero
- By definition of t critical value, we have

$$P(-t_{\alpha/2,n-1} < T < t_{\alpha/2,n-1}) = 1 - \alpha$$

- It is easy to show that $100(1-\alpha)\%$ confidence interval for μ is

$$\left(\bar{x} - t_{\alpha/2,n-1} \cdot \frac{s}{\sqrt{n}}, \bar{x} + t_{\alpha/2,n-1} \cdot \frac{s}{\sqrt{n}}\right)$$

- The alternative, more compact notation is

$$\bar{x} \pm t_{\alpha/2,n-1} \cdot \frac{s}{\sqrt{n}}$$

Example

- Sweetgum lumber is quite valuable but there's a general shortage of high-quality sweetgum today. Because of this, composite beams that are designed to add value to low-grade sweetgum lumber are commonly used.
- The sample consists of 30 observations on the modulus of rapture in psi

Figure :

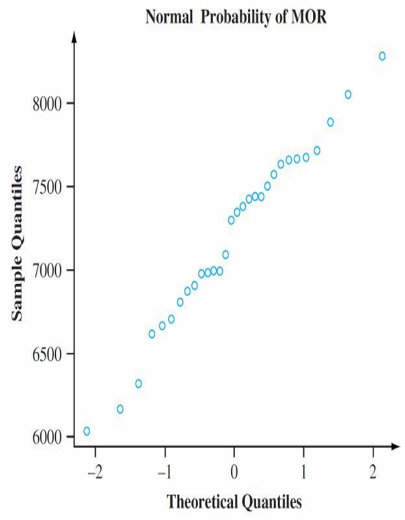
Normal Probability of MOR
Sample Quantiles
8000
7500
7000
6500
6000
-2
-1
0
1
2
Theoretical Quantiles

Calculations

- Take a confidence level of 95%.
- The CI is based on n 1 = 29 degrees of freedom, so the necessary t critical value is 2.045
-

$$\bar{x} \pm t_{0.25,29} \frac{s}{\sqrt{n}} = 7203.191 \pm (2.045) \frac{543.5400}{\sqrt{30}} = (7000.253, 7406.129)$$

Prediction interval

- Consider a random sample $X_1, \ldots, X_n$ from a normal population distribution. Suppose you want to predict X_{n+1}.
- A point predictor is $\bar{X}$; clearly, $E(\bar{X} - X_{n+1}) = \mu - \mu = 0$ and

$$V(\bar{X} - X_{n+1}) = V(\bar{X}) + V(X_{n+1}) = \sigma^2 + \frac{\sigma^2}{n} = \sigma^2 \left(1 + \frac{1}{n}\right)$$

- The prediction error is normally distributed and, therefore,

$$Z = \frac{\bar{X} - X_{n+1}}{\sqrt{\sigma^2 \left(1 + \frac{1}{n}\right)}}$$

has a standard normal distribution

- It is possible to show that

$$T = \frac{\bar{X} - X_{n+1}}{S\sqrt{1 + \frac{1}{n}}}$$

has t distribution with $n-1$ df
- Consequently, the prediction interval for X_{n+1} is

$$\bar{x} \pm t_{\alpha/2,n-1} \cdot s\sqrt{1 + \frac{1}{n}}$$

- Note the obvious difference with the t confidence interval for the mean μ...Why is the prediction interval wider?

- Note that the estimation error $\bar{X} - \mu$ is the deviation from the fixed value while the prediction error $\bar{X} - X_{n+1}$ is a difference between two random variables. The second has much more variability in it than the first...
- Even when $n \to \infty$, the PI approaches $\mu \pm z_{\alpha/2} \cdot \sigma$. This means that there is uncertainty about the true value X even when the infinite amount of information is available.

Example

- A meat inspector has randomly measured 30 packs of 95% lean beef. The sample resulted in the mean 96.2% with the sample standard deviation of 0.8%. Find a 99% prediction interval for a new pack. Assume normality
- For $\nu = 29$ df, we have the critical value $t_{0.005} = 2.756$. Hence a 99% prediction interval for a new observation x_0 is

$$96.2-(2.756)(0.8)\sqrt{1+\frac{1}{30}} < x_0 < 96.2-(2.756)(0.8)\sqrt{1+\frac{1}{30}}$$

which reduces to $(93.96, 98.44)$.

Bootstrap:The Introduction

- Suppose you have some distribution with the density $f(x;\theta)$ where θ is an unknown parameter
- Given a sample $x_1, \ldots, x_n$ from this distribution, you can obtain a point estimate $\hat{\theta}$; as an example, if you have normal distribution with mean μ, you can always estimate it by $\bar{x}$.
- If θ is the only unknown parameter, you can say that the (unknown) pdf $f(x;\theta)$ can be estimated by $f(x;\hat{\theta})$. Now you can generate multiple samples from $f(x;\hat{\theta})$ distribution to get

$$x_1^*, x_2^*, \ldots, x_n^* \tag{1}$$

- With B bootstrap samples at our disposal, we can have the bootstrap estimate of θ $\hat{\theta}^*$. For example, if the parameter in question is the mean μ, we have $\hat{\mu}^* = B^{-1} \sum x_i^*$.
- Why do we need bootstrap? An important issue is estimating the precision of the estimator $\hat{\theta}$; if $\theta = \sigma^2$, it is difficult to estimate the variance $\sigma_{\hat{\theta}}$.
- Using the bootstrap samples, we can estimate it as

$$S_{\hat{\theta}} = \sqrt{\frac{1}{B-1} \sum (\hat{\theta}_i^* - \bar{\theta}^*)^2}$$

Example

- Let X be the time to breakdown of an insulating fluid between electrodes at some voltage and assume it is exponentially distributed $f(x) = \lambda e^{-\lambda x}$
- A random sample of $n = 1$ breakdown times (min) is 41.53, 18,73, 2.99, 30.34, 12.33, 117.52, 73.02, 223.63, 4.00, 26.78
- A reasonable estimate of the distribution parameter is $\lambda = \frac{1}{\bar{x}} = 1/55.087 = 0.018153$
- Generate $B = 100$ samples, each of size 10, from $f(x; 0.018153)$
- Determine the value of $\hat{\lambda}_i^*$ for each $i = 1, \ldots, B$ and find $\bar{\lambda}^* = 0.02153$ and $s_{\hat{\lambda}} = 0.0091$; that last value can be used to construct a confidence interval for λ

Bootstrap Confidence Intervals

- Consider estimating the mean μ of a normal distribution with $\sigma = 1$.
- If $\bar{X}$ is used to estimate μ, $1.96/\sqrt{n}$ is the 97.5% percentile of the distribution of $\hat{\mu} - \mu$ since $P\left(\bar{X} - \mu < 1.96/\sqrt{n}\right) = P(Z < 1.96) = .9750$. Similarly, $-1.96/\sqrt{n}$ is the 2.5% percentile. Hence, we have $P\left(\hat{\mu} - 2.5\% \text{ percentile} > \mu > \hat{\mu} - 97.5\% \text{ percentile}\right)$.
- The above percentiles can be estimated using the bootstrap. If we have $B = 1000$ samples, we have 1000 differences $\hat{\theta}_1^* - \bar{\theta}^*, \ldots, \hat{\theta}_{1000}^* - \bar{\theta}^*$ from which we can compute the necessary percentiles.

Confidence Intervals for the Variance and Standard Deviation of a Normal Population

- Let $X_1, \ldots, X_n$ be a random sample from a normal distribution with mean μ and variance σ^2. Then, the RV

$$\frac{(n-1)S^2}{\sigma^2} = \frac{\sum (X_i - \bar{X})^2}{\sigma^2}$$

  has the so-called χ^2_{n-1} distribution.
- The colloquial name of this distribution is chi-squared with $n-1$ df.

Chi-squared distribution

- χ^2_ν is a continuous distribution with one parameter ν that is called the number of df. ν can be any positive integer.
- Here are the graphs χ^2_ν pdf with various df

Figure :

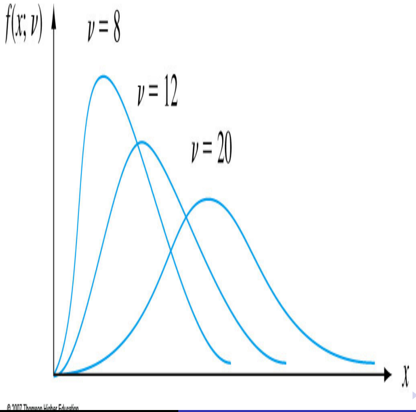

Some additional facts about χ^2_ν

- χ^2_ν is a special case of gamma distribution with $\alpha = \nu/2$ and $\beta = 2$
- A very useful representation of χ^2_ν distribution is as follows:
  1. Suppose we are given n iid standard normal random variables $Z_1, \ldots, Z_n$.
  2. Then, $\sum_{i=1}^n Z_i^2$ has a chi-squared distributions with $\nu = n$ degrees of freedom

Chi-squared confidence interval

- *Chi-squared critical value* $\chi^2_{\alpha,\nu}$ is a number on the measurement axis such that the area under the curve to the right of $\chi^2_{\alpha,\nu}$ is equal to α
- Note that χ^2_ν is not a symmetric distribution so extensive tabulation is needed for α from 0 to 1.
- Clearly,

$$P\left(\chi^2_{1-\alpha/2,n-1} < \frac{(n-1)S^2}{\sigma^2} < \chi^2_{\alpha/2,n-1}\right) = 1-\alpha$$

- Thus, we have

$$\frac{(n-1)S^2}{\chi^2_{\alpha/2,n-1}} < \sigma^2 < \frac{(n-1)S^2}{\chi^2_{1-\alpha/2,n-1}}$$

- We can conclude that $100(1-\alpha)\%$ confidence interval for the variance σ^2 of a normal population has a lower limit

$$\frac{(n-1)s^2}{\chi^2_{\alpha/2,n-1}}$$

and the upper limit

$$\frac{(n-1)s^2}{\chi^2_{1-\alpha/2,n-1}}$$

χ^2 critical values illustrated

Figure :

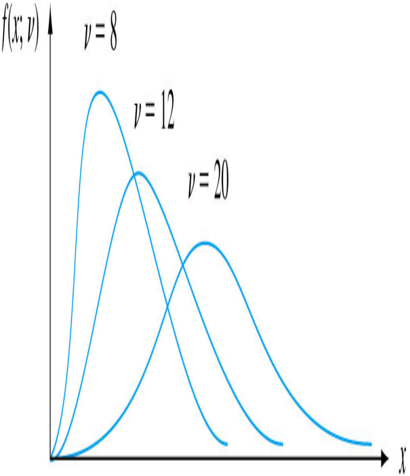

Example

- For a sample of 17 breakdown voltage measurements of electrically stressed circuits, the normal probability plot gives support to the normality assumption
- We can compute the sample variance $s^2 = 137,324.3$ which we use to estimate the true variance σ^2. For 15 df, the respective percentiles are $\chi^2_{.975,16} = 6.908$ and $\chi^2_{0.025,16} = 28.845$.
- The interval is

$$\left(\frac{16(137,324.3)}{28.845}, \frac{16(137,324.3)}{6.908}\right) = (76,172.3, 318,064.4)$$