Topic 5: Confidence Intervals

Ethan P. Marzban University of California, Santa Barbara PSTAT 120B

Outline

1. Asymptotic Confidence Intervals using the MLE

2. Non-Normal Confidence Intervals

Asymptotic Confidence Intervals using the MLE

UC SANTA BARBARA
Department of Statistics
and Applied Probability

Result

Theorem (Asymptotic MLE Result)

Given an i.i.d. sample $\vec{\boldsymbol{Y}}$ from a sample with unknown parameter θ and $\widehat{\theta}_{\text{MLE}}$, the maximum likelihood estimator for θ, we have, under certain "regularity conditions," that

$$\frac{\tau(\widehat{\theta}_{\text{MLE}}) - \tau(\theta)}{\sqrt{\frac{[\tau'(\theta)]^2}{\mathcal{I}_n(\theta)}}} \rightsquigarrow \mathcal{N}(0, 1)$$

UC SANTA BARBARA
Department of Statistics and Applied Probability

Example

Example

Let $Y_1, \cdots, Y_n \overset{\text{i.i.d.}}{\sim} \text{Exp}(\theta)$ for some unknown $\theta > 0$. Find the MLE for the population variance, and use this to construct a large-sample 95% confidence interval.

UC SANTA BARBARA
Department of Statistics and Applied Probability

Non-Normal Confidence Intervals

UC SANTA BARBARA
Department of Statistics
and Applied Probability

Leadup

- Last time, we discussed how to construct confidence intervals under the assumption of a normally-distributed population. What do we do if our population is *not* normally distributed?
- For example, if $Y_1, \cdots, Y_n \overset{\text{i.i.d.}}{\sim} \text{Exp}(\theta)$, how might we construct a $(1 - \alpha) \times 100\%$ confidence interval for θ?
- Well, there are quite a few options available to us!
- One of the most popular ways of constructing CIs is called the **pivotal method** (or **method of pivots**).

UC SANTA BARBARA
Department of Statistics and Applied Probability

Pivots

Definition (Pivot)

Given a sample $Y_1, \cdots, Y_n$ from a distribution with unknown parameter θ, we define a **pivot** (or **pivotal quantity**) for θ to be a function $U := g(\vec{\mathbf{Y}}, \theta)$ whose distribution doesn't depend on θ.

- So, for example, if $Y_1, \cdots, Y_n \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, 1)$, then $U := \sqrt{n}(\overline{Y}_n - \mu)$ is a pivot for μ since its distribution doesn't depend on μ.

UC SANTA BARBARA
Department of Statistics and Applied Probability

Constructing CIs using the Pivotal Method

- Consider a sample $Y_1, \cdots, Y_n$ from a distribution with unknown parameter θ, and assume $U(\vec{\mathbf{Y}}, \theta)$ is a pivotal quantity. Here's how we use the notion of a pivotal quantity to construct a $(1 - \alpha) \times 100\%$ confidence interval for θ:

(1) Set up the interval $\left\{a \leq U(\vec{\mathbf{Y}}, \theta) \leq b\right\}$, and use the distribution of $U(\vec{\mathbf{Y}}, \theta)$ to find a and b such that

$$\mathbb{P}(a \leq U(\vec{\mathbf{Y}}, \theta) \leq b) = 1 - \alpha$$

(2) Invert the interval $\left\{a \leq U(\vec{\mathbf{Y}}, \theta) \leq b\right\}$ to be of the form $\{\widehat{\theta}_L \leq \theta \leq \widehat{\theta}_U\}$, for random variables $\widehat{\theta}_L$ and $\widehat{\theta}_U$ that depend on $U(\vec{\mathbf{Y}}, \theta)$ and a and b.

UC SANTA BARBARA
Department of Statistics and Applied Probability

Example

Example

Suppose we have a single observation $Y \sim \text{Exp}(\theta)$. Construct a pivotal quantity for θ, and use this to construct a 95% confidence interval for θ.

Solution

- To find a pivotal quantity for θ, we want to find a function of Y and θ whose distribution doesn't depend on θ.
- As such, let's propose the following pivotal quantity: $U := (Y/\theta)$.
- To check this is a pivotal quantity, we need to find the distribution of U.
- But hey, we've done this a million times already (back in Topic 02 of this course)!

$$U \sim \mathrm{Exp}\left(\frac{1}{\theta} \cdot \theta\right) \sim \mathrm{Exp}(1)$$

- Because the distribution of U doesn't depend on θ, it is, by definition, a pivotal quantity for θ.

UC SANTA BARBARA
Department of Statistics and Applied Probability

Solution

- So, our 95% CI will take the form

$$\left\{a \leq \frac{Y}{\theta} \leq b\right\}$$

for constants a and b such that

$$\mathbb{P}\left(a \leq \frac{Y}{\theta} \leq b\right) = 0.95$$

- Let's pause, and sketch a picture.

UC SANTA BARBARA
Department of Statistics
and Applied Probability

Solution

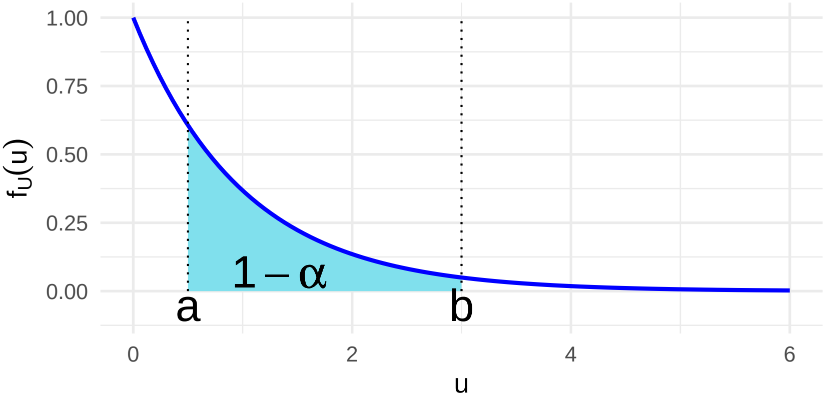

Solution

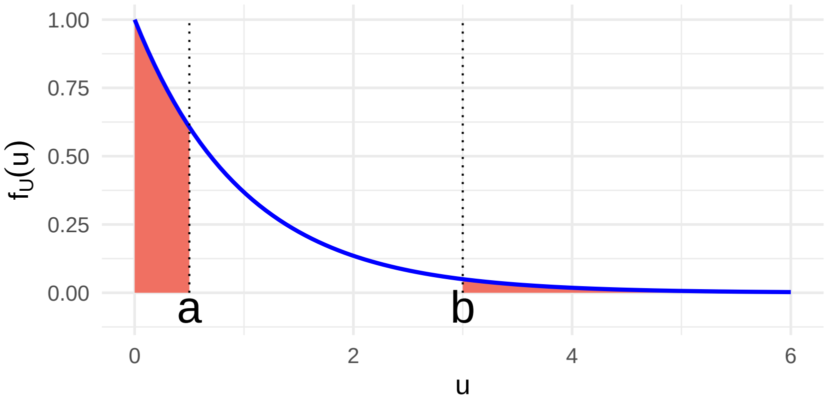

UC SANTA BARBARA
Department of Statistics and Applied Probability

Solution

- So, we want $\mathbb{P}(Y/\theta < a) = \alpha/2$ and $\mathbb{P}(Y/\theta > b) = \alpha/2$.
- Well, since we know the distribution of (Y/θ), we can compute these probabilities directly!

$$\mathbb{P}\left(\frac{Y}{\theta} < a\right) = \int_{-\infty}^{a} f_U(u)\,\mathrm{d}u = \int_{0}^{a} e^{-u}\,\mathrm{d}u = 1 - e^{-a}$$

$$\mathbb{P}\left(\frac{Y}{\theta} > b\right) = \int_{b}^{\infty} f_U(u)\,\mathrm{d}u = \int_{b}^{\infty} e^{-u}\,\mathrm{d}u = e^{-b}$$

- Hence, $1 - e^{-a} = \alpha/2$ and $e^{-b} = \alpha/2$. Equivalently: $a = \ln\left(\frac{2}{2-\alpha}\right)$ and $b = \ln\left(\frac{2}{\alpha}\right)$.

UC SANTA BARBARA
Department of Statistics and Applied Probability

Solution

- The final piece of the puzzle is to invert our interval.
- Let me quickly explain what I mean by this. The interval we started with is
$$\left\{ a \leq \frac{Y}{\theta} \leq b \right\}$$
where a and b are given by the quantities we just found on the previous slide
- But, this isn't the right form for a confidence interval - remember that a true confidence interval needs to have random *endpoints*. Right now, our randomness (i.e. the random variable Y) is in the *interior* of our interval.
- But, not to fret!

UC SANTA BARBARA
Department of Statistics and Applied Probability

Solution

- $\left\{a \leq \frac{Y}{\theta}\right\} \implies \{a\theta \leq Y\} \implies \left\{\theta \leq \frac{Y}{a}\right\}$
- $\left\{\frac{Y}{\theta} \leq b\right\} \implies \{Y \leq \theta b\} \implies \left\{\theta \geq \frac{Y}{b}\theta\right\}$
- So, combining everything:

$$\left\{a \leq \frac{Y}{\theta} \leq b\right\} \implies \left\{\frac{Y}{b} \leq \theta \leq \frac{Y}{a}\right\}$$

- This is now the right form for a CI - the randomness is in the endpoints!

Solution

- So, all in all, our $(1-\alpha)\times 100\%$ confidence interval for θ takes the form
$$\left[\frac{Y}{b}, \frac{Y}{a}\right]$$
or, substituting in our expressions for a and b:
$$\left[\frac{Y}{\ln\left(\frac{2}{\alpha}\right)}, \frac{Y}{\ln\left(\frac{2}{2-\alpha}\right)}\right]$$

UC SANTA BARBARA
Department of Statistics
and Applied Probability

Example

Example

The wait time of a randomly-selected person at *The Arbor* is assumed to follow an exponential distribution with parameter θ. If a single person was selected at random, and it was observed that they spent 3 minutes waiting in line at *The Arbor*, construct a 95% confidence interval for θ, the true average wait times at *The Arbor*.

Solution

- All we need to do is plug into our CI from the previous example!

$$\left[\frac{3}{\ln\left(\frac{2}{0.05}\right)} , \frac{3}{\ln\left(\frac{2}{2-0.05}\right)} \right] = [0.813 , 118.494]$$

- If that looks like a really wide interval.. it's because it is! But, remember that we only had a single sample. So, intuitively, it perhaps makes sense that our CI should be pretty wide.

Example

Example

Suppose $Y_1, \cdots, Y_n \overset{\text{i.i.d.}}{\sim} \text{Exp}(\theta)$, where $\theta > 0$ is an unknown parameter.

(a) Propose a pivotal quantity for θ that is a function of $\sum_{i=1}^n Y_i$, and show that your proposed quantity is in fact a pivot for θ.

(b) Use your pivot from part (a) to construct a $(1 - \alpha) \times 100\%$ confidence interval for θ.

Solution

- We know that $(\sum_{i=1}^{n} Y_i) \sim \text{Gamma}(n, \theta)$.
- Therefore, $(1/\theta)(\sum_{i=1}^{n} Y_i) = [\sum_{i=1}^{n}(Y_i/\theta)] \sim \text{Gamma}(n, 1)$, which doesn't depend on θ.
- Hence, a pivotal quantity for θ involving $\sum_{i=1}^{n} Y_i$ is

$$U := \sum_{i=1}^{n} \left(\frac{Y_i}{\theta} \right)$$

Solution

- Is this the only pivotal quantity for θ that involves $\sum_{i=1}^{n} Y_i$? No! In fact, we can multiply U by any positive constant [that doesn't depend on θ] and get another pivotal quantity for θ.
- For example,

$$U_2 := 2\sum_{i=1}^{n}\left(\frac{Y_i}{\theta}\right) \sim \text{Gamma}\,(n, 2) \sim \chi^2_{2n}$$

thereby showing that U_2 is *also* a pivotal quantity for θ.

UC SANTA BARBARA
Department of Statistics and Applied Probability

Solution

- Let's start off by using U_2 to construct a $(1-\alpha)\times 100\%$ CI for θ.
- Our initial interval takes the form

$$\left\{ a \leq 2 \sum_{i=1}^{n} \left(\frac{Y_i}{\theta} \right) \leq b \right\}$$

- By the definnition of coverage probability, we seek a and b such that

$$\mathbb{P}\left(a \leq 2 \sum_{i=1}^{n} \left(\frac{Y_i}{\theta} \right) \leq b \right) = 1 - \alpha$$

Solution

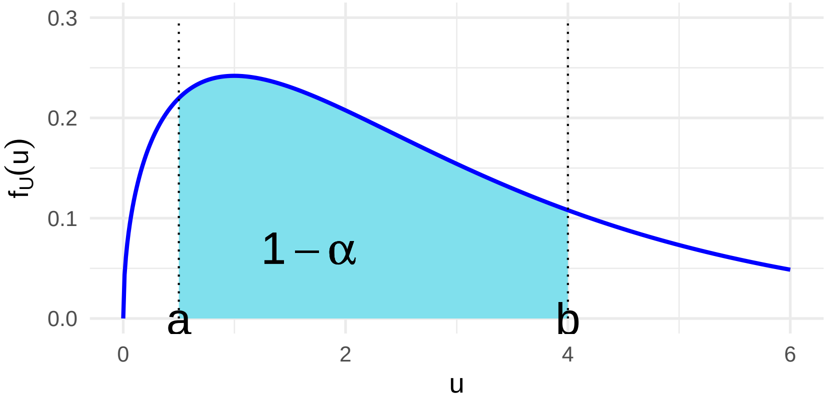

Solution

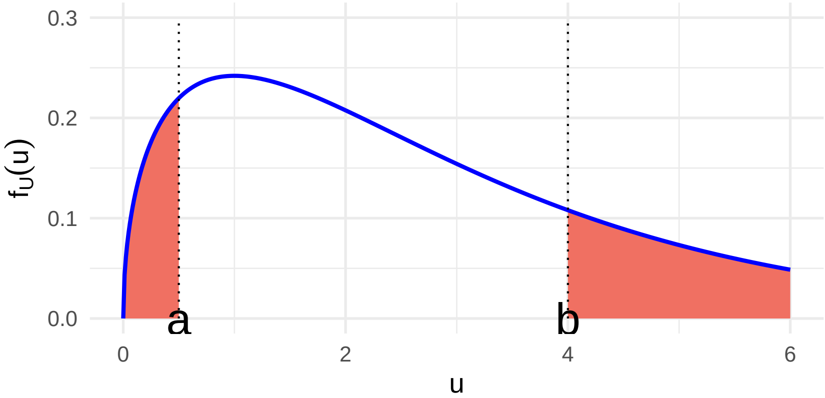

UC SANTA BARBARA
Department of Statistics
and Applied Probability

Solution

- So, we want

$$\mathbb{P}(U_2 < a) = \frac{\alpha}{2} \qquad \text{and} \qquad \mathbb{P}(U_2 > b) = \frac{\alpha}{2}$$

- Now, we don't have a simply closed-form expression for the CDF of the χ^2_{2n} distribution. But that's fine - we can still write

$$a = F^{-1}_{\chi^2_{2n}}\left(\frac{\alpha}{2}\right) \qquad \text{and} \qquad b = F^{-1}_{\chi^2_{2n}}\left(1 - \frac{\alpha}{2}\right)$$

Solution

- Finally, we invert the interval.
- $\left\{ a \leq \frac{2}{\theta} \sum_{i=1}^{n} Y_i \right\} \implies \left\{ \theta \leq \frac{2}{a} \sum_{i=1}^{n} Y_i \right\}$
- $\left\{ \frac{2}{\theta} \sum_{i=1}^{n} Y_i \leq b \right\} \implies \left\{ \theta \geq \frac{2}{b} \sum_{i=1}^{n} Y_i \right\}$
- So, our interval is

$$\left[\frac{2 \sum_{i=1}^{n} Y_i}{F_{\chi^2_{2n}}^{-1}\left(1 - \frac{\alpha}{2}\right)} , \frac{2 \sum_{i=1}^{n} Y_i}{F_{\chi^2_{2n}}^{-1}\left(\frac{\alpha}{2}\right)} \right]$$

UC SANTA BARBARA
Department of Statistics and Applied Probability

Solution

- Let's see what would have happened if we used our other pivotal quantity, $U_1 := (1/\theta)\sum_{i=1}^n Y_i$.
- We'd start with

$$\left\{ a \leq \frac{1}{\theta}\sum_{i=1}^{n} Y_i \right\}$$

- Ultimately, we'd find (and I encourage you to fill in these steps)

$$a = F^{-1}_{\mathrm{Gamma}(n,1)}\left(\frac{\alpha}{2}\right); \qquad b = F^{-1}_{\mathrm{Gamma}(n,1)}\left(1 - \frac{\alpha}{2}\right)$$

Solution

- Appropriately inverting the interval (and, again, you can try this on your own), we get

$$\left[\frac{\sum_{i=1}^{n} Y_i}{F^{-1}_{\text{Gamma}(n,1)}\left(1 - \frac{\alpha}{2}\right)} , \frac{\sum_{i=1}^{n} Y_i}{F^{-1}_{\text{Gamma}(n,1)}\left(\frac{\alpha}{2}\right)} \right]$$

- Do these give the same numerical values, given a particular observed dataset? Let's check!

UC SANTA BARBARA
Department of Statistics and Applied Probability

Solution

- Take $n = 7$ and $\overline{y}_7 = 4.77$ (so $\sum_{i=1}^{n} y_i = 33.39$) and adopt a 95% coverage probability.
- Our first CI, using the χ^2_{2n} distribution, becomes

$$\left[\frac{2(33.39)}{F^{-1}_{\chi^2_{14}}(0.975)} \ , \ \frac{2\sum_{i=1}^{n} Y_i}{F^{-1}_{\chi^2_{14}}(0.025)} \right] = [2.557 \ , \ 11.864]$$

- Our first CI, using the Gamma$(n, 1)$ distribution, becomes

$$\left[\frac{33.39}{F^{-1}_{\text{Gamma}(7,1)}(0.975)} \ , \ \frac{33.39}{F^{-1}_{\text{Gamma}(7,1)}(0.025)} \right] = [2.557 \ , \ 11.864]$$

Normal CI

- Finally, to tie together yesterday's lecture and today's lecture, let's see if we can actually use the pivotal method to recover the same sorts of normal-adjacent confidence intervals we constructed before.
- So, to start, let $Y_1, \cdots, Y_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mu, \sigma^2)$ where $\mu \in \mathbb{R}$ is unknown but $\sigma^2 > 0$ is known.
- A pivotal quantity is

$$U := \frac{\overline{Y}_n - \mu}{\sigma / \sqrt{n}} \sim \mathcal{N}(0, 1)$$

UC SANTA BARBARA
Department of Statistics and Applied Probability

Normal CI

- So our CI takes the form

$$\left\{ a \leq \frac{\overline{Y}_n - \mu}{\sigma/\sqrt{n}} \leq b \right\}$$

- Since we want a $(1-\alpha) \times 100\%$ coverage probability, we get

$$\mathbb{P}\left(a \leq \frac{\overline{Y}_n - \mu}{\sigma/\sqrt{n}} \leq b \right) = 1 - \alpha$$

UC SANTA BARBARA
Department of Statistics
and Applied Probability

Normal CI

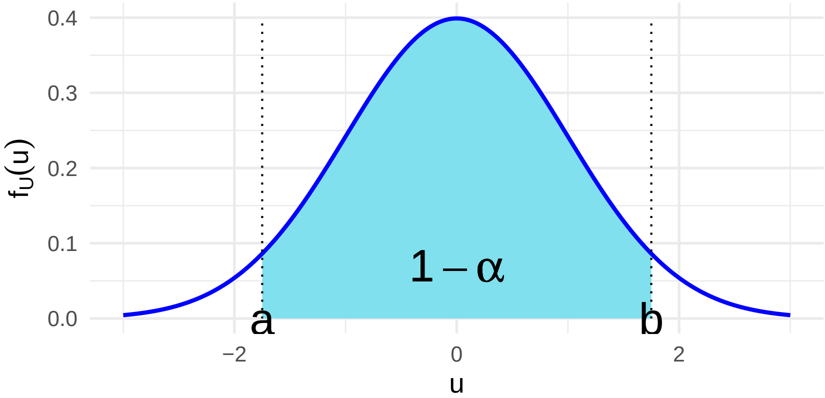

UC SANTA BARBARA
Department of Statistics and Applied Probability

Normal CI

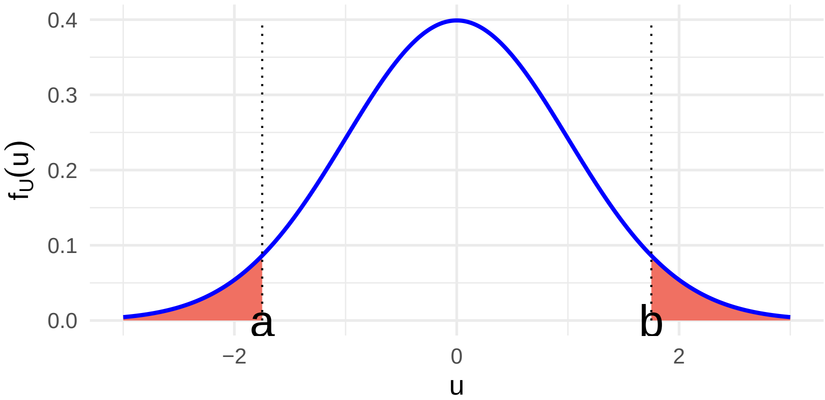

UC SANTA BARBARA
Department of Statistics
and Applied Probability

Normal CI

- $\mathbb{P}\left(\frac{\overline{Y}_n - \mu}{\sigma / \sqrt{n}} < a \right) = \Phi(a) \stackrel{!}{=} \frac{\alpha}{2} \implies a = \Phi^{-1}\left(\frac{\alpha}{2} \right)$
- $\mathbb{P}\left(\frac{\overline{Y}_n - \mu}{\sigma / \sqrt{n}} > b \right) = 1 - \Phi(b) \stackrel{!}{=} \frac{\alpha}{2} \implies b = \Phi^{-1}\left(1 - \frac{\alpha}{2} \right)$
- Now, we invert our interval:

UC SANTA BARBARA
Department of Statistics and Applied Probability

Normal CI

- $\left\{a \leq \frac{\overline{Y}_n - \mu}{\sigma/\sqrt{n}}\right\} \implies \left\{a \cdot \frac{\sigma}{\sqrt{n}} \leq \overline{Y}_n - \mu\right\} \implies \left\{\mu \leq \overline{Y}_n - a \cdot \frac{\sigma}{\sqrt{n}}\right\}$
- $\left\{\frac{\overline{Y}_n - \mu}{\sigma/\sqrt{n}} \leq b\right\} \left\{\overline{Y}_n - \mu \leq b \cdot \frac{\sigma}{\sqrt{n}}\right\} \implies \left\{\mu \geq \overline{Y}_n - b \cdot \frac{\sigma}{\sqrt{n}}\right\}$
- So our interval is equivalent to

$$\left\{\overline{Y}_n - b \cdot \frac{\sigma}{\sqrt{n}} \leq \mu \leq \overline{Y}_n - a \cdot \frac{\sigma}{\sqrt{n}}\right\}$$

Normal CI

- Therefore, plugging in the values for *a* and *b* we derived earlier on, our $(1-\alpha)\times 100\%$ CI, as derived using the method of pivotal quantities, is

$$\left[\overline{Y}_n - \Phi^{-1}\left(1-\frac{\alpha}{2}\right)\cdot\frac{\sigma}{\sqrt{n}}\,,\ \overline{Y}_n - \Phi^{-1}\left(\frac{\alpha}{2}\right)\cdot\frac{\sigma}{\sqrt{n}}\right]$$

- Finally, note that $-\Phi^{-1}\left(\frac{\alpha}{2}\right) = \Phi^{-1}\left(1-\frac{\alpha}{2}\right)$ to see that our interval is

$$\left[\overline{Y}_n - \Phi^{-1}\left(1-\frac{\alpha}{2}\right)\cdot\frac{\sigma}{\sqrt{n}}\,,\ \overline{Y}_n + \Phi^{-1}\left(1-\frac{\alpha}{2}\right)\cdot\frac{\sigma}{\sqrt{n}}\right] = \overline{Y}_n \pm \Phi^{-1}\left(1-\frac{\alpha}{2}\right)\cdot\frac{\sigma}{\sqrt{n}}$$

UC SANTA BARBARA
Department of Statistics and Applied Probability