

Interval Estimation

Maura Mezzetti

Department of Economics and Finance
Università Tor Vergata

Outline

- 1 Sampling from the Normal Distribution
- 2 Confidence Intervals
 - The Pivotal Approach
 - Example

Properties of Sample Mean and Sample Variance

Theorem Let $X_1, X_2, \dots, X_n$ be a random sample from a population $X \sim N(\mu, \sigma^2)$, Consider the following statistics: $\bar{X} = \frac{\sum_i^n X_i}{n}$ and $S^2 = \frac{\sum_i^n (X_i - \bar{X})^2}{n-1}$. Then:

- 1 $\bar{X}$ and S^2 are independent
- 2 $\bar{X} \sim N\left(\mu, \frac{\sigma^2}{n}\right)$
- 3 $\frac{(n-1)S^2}{\sigma^2}$ has a chi-squared distribution with $n - 1$ degrees of freedom.

Chi Squared random Variable

The notation $X \sim \chi_p^2$ indicates a Chi Squared Random Variable with p degrees of freedom:

- If $Z \sim N(0, 1)$, then $Z^2 \sim \chi_2^1$, that is, the square of a standard normal random variable is a Chi Squared random variable with one degree of freedom.
- If $X_1, X_2, \dots, X_n$ are independent and $X_i \sim \chi_{p_i}^2$ then

$$X_1 + X_2 + \dots + X_n \sim \chi_{p_1+p_2+\dots+p_n}^2,$$

that is, independent chi squared variables add to a chi squared variable, and the degrees of freedom also adds.

The Derived distributions: Students's t and Snedecor's F

- Let $X_1, X_2, \dots, X_n$ be a random sample from $N(\mu, \sigma^2)$ distribution, then the quantity:

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

has *Student's t distribution* with n degrees of freedom and we write $T \sim t_{n-1}$.

- Let $X_1, X_2, \dots, X_n$ be a random sample from $N(\mu_x, \sigma_x^2)$ and let $Y_1, Y_2, \dots, Y_m$ be a random sample from $N(\mu_y, \sigma_y^2)$ distribution, then the quantity:

$$F = \frac{S_X^2/\sigma_X^2}{S_Y^2/\sigma_Y^2}$$

has *Snedecor's F distribution* with $n - 1$ and $m - 1$ degrees of freedom and we write $F \sim F_{n-1, m-1}$.

The Derived distributions: Students's t and Snedecor's F

Theorem:

- 1 If $U \sim \chi_p$ and $V \sim \chi_q$ then $\frac{U/p}{V/q} \sim F_{p,q}$
- 2 If $X \sim F_{p,q}$, then $1/X \sim F_{q,p}$
- 3 If $X \sim t_p$, then $X^2 \sim F_{1,p}$
- 4 Higher the degrees of freedom, closer is the *Student's t distribution* to a Normal standard

Interval Estimation

The purpose of using an interval estimator rather than a point estimator is to have some guarantee of capturing the parameter of interest. The certainty of this guarantee is quantified in the following definitions.

Standard Error and Precision

Let $X_1, \dots, X_n$ a random sample with pdf $f(x; \theta)$. Let $T = T(\mathbf{X})$ a point estimator for θ .

Definition

The standard deviation of an estimator is called the standard error of the estimator (SE). In symbol

$$SE(T) = \sqrt{\text{Var}(T)}$$

- The precision of an estimator is measured by the SE of the estimator
- A precise estimator has a small standard error, but exactly how are the precision and the standard error related? (The answer will be given by the confidence intervals)

General Concepts

- Estimators provide a point estimate of their target θ , but give no indication of how precise they are.
- To obtain a better picture, we look for a **confidence interval**.
- Idea: first choose a **confidence level** α , which is the probability to make a mistake.
- Then find two functions $L(X_1, \dots, X_n) < U(X_1, \dots, X_n)$, such that the probability that θ is outside $[L, U]$ is less than or equal to α .

From point to interval estimators

Example

Let $X_1, \dots, X_n$ i.i.d $X \sim N(\mu, 1)$.

- $\bar{X}$ is a point estimator for μ
- $[\bar{X} - 1; \bar{X} + 1]$ is an interval estimator for μ

From point to interval estimators

Let $X_1, \dots, X_9$ i.i.d $X \sim N(\mu, 1)$

$[\bar{X} - 1; \bar{X} + 1]$ is an interval estimator for μ

From point to interval estimators

Let $X_1, \dots, X_9$ i.i.d $X \sim N(\mu, 1)$

$[\bar{X} - 1; \bar{X} + 1]$ is an interval estimator for μ

What is the coverage probability of $[\bar{X} - 1; \bar{X} + 1]$?

0.9973

$$P(\bar{X} - 1 \leq \mu \leq \bar{X} + 1) = 0.9973$$

From point to interval estimators

Let $X_1, \dots, X_9$ i.i.d $X \sim N(\mu, 1)$

$[\bar{X} - 1; \bar{X} + 1]$ is an interval estimator for μ

What is the coverage probability of $[\bar{X} - 1; \bar{X} + 1]$?

0.9973

$$P(\bar{X} - 1 \leq \mu \leq \bar{X} + 1) = 0.9973$$

Necessary Distribution of $\bar{X}$

$$\frac{\bar{X} - \mu}{\frac{1}{3}} \sim N(0, 1)$$

From point to interval estimators

Let $X_1, \dots, X_9$ i.i.d $X \sim N(\mu, 1)$

$[\bar{X} - 1; \bar{X} + 1]$ is an interval estimator for μ

What is the coverage probability of $[\bar{X} - 1; \bar{X} + 1]$?

0.9973

$$P(\bar{X} - 1 \leq \mu \leq \bar{X} + 1) = 0.9973$$

Necessary Distribution of $\bar{X}$

$$\frac{\bar{X} - \mu}{\frac{1}{3}} \sim N(0, 1)$$

From point to interval estimators

Example

Let $X_1, \dots, X_n$ i.i.d $X \sim N(\mu, 1)$.

- $\bar{X}$ is a point estimator for μ
- $[\bar{X} - 1; \bar{X} + 1]$ is an interval estimator for μ

- Estimating μ with $\bar{X}$ we get, for every n ,

$$P(\bar{X} = \mu) = 0$$

- Estimating μ with $[\bar{X} - 1; \bar{X} + 1]$ we get for $n = 9$, for example,

$$P(\bar{X} - 1 \leq \mu \leq \bar{X} + 1) = 0.9973$$

We loose in precision, but we gain some **confidence**, or assurance, that our assertion about μ is correct

Example: Uniform distribution

Example

Let $(X_1, X_2, \dots, X_n)$ i.i.d. Uniform random variables on $[0, \theta]$. The MLE of θ is $Y = \max(X_1, X_2, \dots, X_n)$, the sample maximum. Find the coverage probability for the interval estimators:

- 1 $[aY, bY]$, $1 \leq a < b$
- 2 $[Y + c, Y + d]$, $0 \leq c < d$

Distribution of Y ?

$$\begin{aligned}P(Y \leq y) &= \left(\frac{y}{\theta}\right)^n \\f_Y(y) &= n \left(\frac{y}{\theta}\right)^{n-1} \frac{1}{\theta}\end{aligned}$$

Example: Uniform distribution

Example

Let $(X_1, X_2, \dots, X_n)$ i.i.d. Uniform random variables on $[0, \theta]$. The MLE of θ is $Y = \max(X_1, X_2, \dots, X_n)$, the sample maximum. Find the coverage probability for the interval estimators:

- ① $[aY, bY]$, $1 \leq a < b$
- ② $[Y + c, Y + d]$, $0 \leq c < d$

Distribution of Y ?

$$P(Y \leq y) = \left(\frac{y}{\theta}\right)^n$$

$$f_Y(y) = n \left(\frac{y}{\theta}\right)^{n-1} \frac{1}{\theta}$$

Example: Uniform distribution

Example

Let $(X_1, X_2, \dots, X_n)$ i.i.d. Uniform random variables on $[0, \theta]$. The MLE of θ is $Y = \max(X_1, X_2, \dots, X_n)$, the sample maximum. Find the coverage probability for the interval estimators:

- 1 $[aY, bY]$, $1 \leq a < b$
- 2 $[Y + c, Y + d]$, $0 \leq c < d$

Prove these results

- 1 $P_\theta(\theta \in [aY, bY]) = P_\theta\left(\frac{1}{b} \leq \frac{Y}{\theta} \leq \frac{1}{a}\right) = \left(\frac{1}{a}\right)^n - \left(\frac{1}{b}\right)^n$
- 2 $P_\theta(\theta \in [Y + c, Y + d]) = \left(1 - \frac{c}{\theta}\right)^n - \left(1 - \frac{d}{\theta}\right)^n$

Example: Uniform distribution

Example

Let $(X_1, X_2, \dots, X_n)$ i.i.d. Uniform random variables on $[0, \theta]$. The MLE of θ is $Y = \max(X_1, X_2, \dots, X_n)$, the sample maximum. Find the coverage probability for the interval estimators:

- ① $[aY, bY]$, $1 \leq a < b$
- ② $[Y + c, Y + d]$, $0 \leq c < d$

Prove these results

- ① $P_\theta(\theta \in [aY, bY]) = P_\theta\left(\frac{1}{b} \leq \frac{Y}{\theta} \leq \frac{1}{a}\right) = \left(\frac{1}{a}\right)^n - \left(\frac{1}{b}\right)^n$
- ② $P_\theta(\theta \in [Y + c, Y + d]) = \left(1 - \frac{c}{\theta}\right)^n - \left(1 - \frac{d}{\theta}\right)^n$

Example: Uniform distribution

Example

Let $(X_1, X_2, \dots, X_n)$ i.i.d. Uniform random variables on $[0, \theta]$. The MLE of θ is $Y = \max(X_1, X_2, \dots, X_n)$, the sample maximum. Find the coverage probability for the interval estimators:

- 1 $[aY, bY]$, $1 \leq a < b$
- 2 $[Y + c, Y + d]$, $0 \leq c < d$

In the first case the coverage probability does not depend on θ and the confident coefficient is $\left(\frac{1}{a}\right)^n - \left(\frac{1}{b}\right)^n$.

In the second case the coverage probability depends on θ and the confident coefficient is zero: $\lim_{\theta \rightarrow \infty} \left[\left(1 - \frac{c}{\theta}\right)^n - \left(1 - \frac{d}{\theta}\right)^n \right] = 0$.

Example: Uniform distribution

Example

Let $(X_1, X_2, \dots, X_n)$ i.i.d. Uniform random variables on $[0, \theta]$. The MLE of θ is $Y = \max(X_1, X_2, \dots, X_n)$, the sample maximum. Find the coverage probability for the interval estimators:

- ① $[aY, bY]$, $1 \leq a < b$
- ② $[Y + c, Y + d]$, $0 \leq c < d$

In the first case the coverage probability does not depend on θ and the confident coefficient is $\left(\frac{1}{a}\right)^n - \left(\frac{1}{b}\right)^n$.

In the second case the coverage probability depends on θ and the confident coefficient is zero: $\lim_{\theta \rightarrow \infty} \left[\left(1 - \frac{c}{\theta}\right)^n - \left(1 - \frac{d}{\theta}\right)^n \right] = 0$.

Confidence Intervals

- Since point estimates only offer a single value as a guess to our population parameter, a confidence interval estimate is more desirable.
- With confidence intervals we offer a range of values for which our population parameter might lie and we give an evaluation of how confident we are that it falls within that range.

Definition of Confidence Intervals

Let $X \sim f(x; \theta)$, and let $(X_1, X_2, \dots, X_n)$ an iid random sample, consider two statistics L and U such that:

- $L < U$
- $P(L < \theta < U) = 1 - \alpha, \quad \forall \theta \in \Theta$

$[L(X), U(X)]$ is a random $(1 - \alpha)$ -confidence interval $(l(x), u(x))$ is numeric $(1 - \alpha)$ -confidence interval.

Definitions

Definition For an interval estimator $[L(X), U(X)]$ of a parameter θ , the coverage probability of $[L(X), U(X)]$ is the probability that the random interval $[L(X), U(X)]$ covers the true parameter, θ . In symbols, it is denoted by $P(\theta \in [L(X), U(X)])$

Definition For an interval estimator $[L(X), U(X)]$ of a parameter θ , the confidence coefficient of $[L(X), U(X)]$ is the minimum of the coverage probabilities,.

Confidence Intervals

To methods to build confidence intervals:

Inversion of a Test Procedure See later

The Pivotal Approach

Pivotal Random Variables

Definition

A **Pivotal random variable** $Q(\mathbf{X}, \theta)$ is a random variable (but in general, not a statistic) whose distribution does not depend on the parameter θ . Here $\mathbf{X}$ is the random sample.

- If the pivotal quantity can be inverted as a function of θ for each value of X , then we may find a confidence region for θ .
- Inversion means that, for each x and each subset A of the range of $Q(x, \theta)$, there exists a subset $R(x, A)$ of θ such that

$$Q(x, \theta) \in A \iff \theta \in R(x, A)$$

This implies that

$$P(Q(x, \theta) \in A) = P(\theta \in R(x, A))$$

Example of Pivotal Random Variables

$X \sim N(\mu, \sigma)$ σ known

$$\frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{n}}}$$

is a pivotal random variable from μ

$X \sim U(0, \theta)$

$$\frac{\max(X_1, \dots, X_n)}{\theta}$$

is a pivotal random variable from θ

Example of Pivotal Random Variables

$X \sim N(\mu, \sigma)$ σ known

$$\frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{n}}}$$

is a pivotal random variable from μ

$X \sim U(0, \theta)$

$$\frac{\max(X_1, \dots, X_n)}{\theta}$$

is a pivotal random variable from θ

$X \sim \text{Exp}(\lambda)$

$$2\lambda \sum_i X_i$$

Example of Pivotal Random Variables

$X \sim N(\mu, \sigma)$ σ known

$$\frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{n}}}$$

is a pivotal random variable from μ

$X \sim U(0, \theta)$

$$\frac{\max(X_1, \dots, X_n)}{\theta}$$

is a pivotal random variable from θ

$X \sim \text{Exp}(\lambda)$

$$2\lambda \sum_i X_i$$

Pivotal Random Variables

From a (nontrivial) pivotal random variable one can often obtain a confidence interval

In general there may be more than one pivotal random variable and confidence interval

Pivotal Random Variables

From a (nontrivial) pivotal random variable one can often obtain a confidence interval

In general there may be more than one pivotal random variable and confidence interval

Confidence interval at a fixed level $1 - \alpha$

Let T is a pivotal random variable for θ . If T is monotone as function of θ , then, fixed $0 < \alpha < 1$, we can find t_1 e t_2 such that:

$$P(t_1 < T(X_1, X_2, \dots, X_n, \theta) < t_2) = 1 - \alpha$$

for every θ

For example one possible solution is through the equation

$$P(T < t_2) = F(t_2) = 1 - \frac{\alpha}{2}$$


$$P(T < t_1) = F(t_1) = \frac{\alpha}{2}$$

where F is the cdf of T . The solution are the quantiles

$$t_2 = t_{1-\frac{\alpha}{2}}$$

$$t_1 = t_{\frac{\alpha}{2}}$$

Confidence interval



Interpretation of Confidence Intervals

- Remember the parameter is fix and the confidence interval is random.
- If different samples are drawn, the confidence intervals will be usually different
- In the long run, the proportion of time the interval will contain the actual parameter will be approximately equal to the confidence coefficient.

Confidence interval construction for the parameter θ

Let $t = T(X_1, X_2, \dots, X_n; \theta)$ be monotone in θ . Let $\theta = T^{-1}(X_1, X_2, \dots, X_n, t)$, the inverse. Then

$$P(t_1 < T(X_1, X_2, \dots, X_n, \theta) < t_2) = 1 - \alpha$$

is equivalent to

$$P(T^{-1}(X_1, X_2, \dots, X_n, t_1) < \theta < T^{-1}(X_1, X_2, \dots, X_n, t_2)) = 1 - \alpha$$

that is a confidence interval for θ .

- $L(X) = T^{-1}(X_1, X_2, \dots, X_n, t_1)$
- $U(X) = T^{-1}(X_1, X_2, \dots, X_n, t_2)$

Confidence Interval

The **casual interval** (L, U) is such that $P(L < \theta < U) = 1 - \alpha$.

Pivot for μ


- $X \sim N(\mu, \sigma^2)$, with σ known

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

- $X \sim N(\mu, \sigma^2)$, with σ unknown

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

Pivotal distribution for μ in $N(\mu, \sigma)$



Gaussian distribution

Let $X_1, \dots, X_n$ a random sample from a Gaussian distribution. Let $\bar{X}$ and S^2 the sample mean and the sample variance.

$X \sim N(\mu, \sigma^2)$, with σ known

The quantity $Z = \frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{n}}}$ has a $N(0, 1)$ distribution, it is a pivot for μ

$X \sim N(\mu, \sigma^2)$, with σ unknown

The quantity $T = \frac{\bar{X} - \mu}{\frac{S}{\sqrt{n}}}$ has a t -Student with $n - 1$ d.f. distribution, it is a pivot for μ

Gaussian distribution

Let $X_1, \dots, X_n$ a random sample from a Gaussian distribution. Let $\bar{X}$ and S^2 the sample mean and the sample variance.

$X \sim N(\mu, \sigma^2)$, with σ known

The quantity $Z = \frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{n}}}$ has a $N(0, 1)$ distribution, it is a pivot for μ

$X \sim N(\mu, \sigma^2)$, with σ unknown

The quantity $T = \frac{\bar{X} - \mu}{\frac{S}{\sqrt{n}}}$ has a t -Student with $n - 1$ d.f. distribution, it is a pivot for μ

$X \sim N(\mu, \sigma^2)$, with σ and μ unknown

The quantity $G = \frac{(n-1)S^2}{\sigma^2}$ has a χ^2 with $n - 1$ d.f. distribution, it is a pivot for σ

Gaussian distribution

Let $X_1, \dots, X_n$ a random sample from a Gaussian distribution. Let $\bar{X}$ and S^2 the sample mean and the sample variance.

$X \sim N(\mu, \sigma^2)$, with σ known

The quantity $Z = \frac{\bar{X} - \mu}{\frac{\sigma}{\sqrt{n}}}$ has a $N(0, 1)$ distribution, it is a pivot for μ

$X \sim N(\mu, \sigma^2)$, with σ unknown

The quantity $T = \frac{\bar{X} - \mu}{\frac{S}{\sqrt{n}}}$ has a t -Student with $n - 1$ d.f. distribution, it is a pivot for μ

$X \sim N(\mu, \sigma^2)$, with σ and μ unknown

The quantity $G = \frac{(n-1)S^2}{\sigma^2}$ has a χ^2 with $n - 1$ d.f. distribution, it is a pivot for σ

Confidence Intervals for population mean

(For the proof of the following equations see any Statistics Textbook), here it is just provided the analytic formula:

- $X \sim N(\mu, \sigma^2)$ Confidence Intervals for μ when σ is known:

$$(1 - \alpha)CI = \left[\bar{x} - z_{1-\alpha/2} \frac{\sigma}{\sqrt{n}}, \bar{x} + z_{1-\alpha/2} \frac{\sigma}{\sqrt{n}} \right]$$

- $X \sim N(\mu, \sigma^2)$ Confidence Intervals for μ when σ is unknown:

$$(1 - \alpha)CI = \left[\bar{x} - t_{1-\alpha/2}^{n-1} \frac{s_c}{\sqrt{n}}, \bar{x} + t_{1-\alpha/2}^{n-1} \frac{s_c}{\sqrt{n}} \right]$$

Confidence Intervals for population mean

- Let X have any distribution with $E(X) = \mu$ and $Var(X) = \sigma^2$, and let consider sample size large enough to apply central limit theorem: Confidence Intervals for μ :

$$(1 - \alpha)CI = \left[\bar{x} - z_{1-\alpha/2} \frac{s}{\sqrt{n}}, \bar{x} + z_{1-\alpha/2} \frac{s}{\sqrt{n}} \right]$$

- Where:
 - n is the sample size
 - $z_{1-\alpha/2}$ $(1 - \alpha/2)$ -th percentile of the standard gaussian distribution,
 - $t_{1-\alpha/2}^{n-1}$ $(1 - \alpha/2)$ -th percentile of the t-student distribution with $(n-1)$ degrees of freedom,

$$s^2 = \sum_i \frac{(x_i - \bar{x})^2}{n}, \quad s_c^2 = \sum_i \frac{(x_i - \bar{x})^2}{n-1}$$

The Rule for Sample Proportions

If numerous samples or repetitions of size n are taken, the frequency curve of the sample proportions $\hat{p}$ from various samples will be *approximately bell-shaped*. The mean of those sample proportions will be p (the population proportion). The variance will be (recall variance of a Binomial):

$$\frac{p(1 - p)}{n}$$

How can we estimate a confidence intervals for the parameter p ?

- Central limit theorem tells us that the sample proportion is normally distributed with mean p and variance

$$\frac{p(1-p)}{n}$$



$$\hat{p} \sim N\left(p, \frac{p(1-p)}{n}\right)$$

- Confidence Intervals for p :

$$(1-\alpha)CI = \left[\hat{p} - z_{1-\alpha/2} \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}, \hat{p} + z_{1-\alpha/2} \sqrt{\frac{\hat{p}(1-\hat{p})}{n}} \right]$$

Length of a confidence interval

- With α fixed, if n increase, confidence interval length decreases, conversely if n decreases, confidence interval length increases.
- With n fixed, if $(1 - \alpha)$ increases, $z_{1-\alpha}$ increases and confidence interval length increases, conversely, if $(1 - \alpha)$ decreases, $z_{1-\alpha}$ decreases and confidence interval length decreases.

The Interpretation of a Confidence Coefficient

The interpretation of the phrase $(1 - \alpha)$ **confidence** simply means this: Suppose hypothetically that we keep observing different data $X = (x_1, x_2, \dots, x_m, \dots)$ for a long time, and we keep constructing the corresponding observed confidence interval estimates $(L(x_1), U(x_1)), (L(x_2), U(x_2)), \dots, (L(x_m), U(x_m))$. In the long run, out of all these intervals constructed, approximately $100(1 - \alpha)\%$ would include the unknown value of the parameter θ .

Methods of evaluating Interval Estimators

Theorem Let $f(x)$ be a unimodal pdf. in the interval $[a, b]$ satisfies

① $\int_a^b f(x)dx = 1 - \alpha$

② $f(a) = f(b) > 0$

③ $a \leq x^* \leq b$ where x^* is a mode for $f(x)$

Then $[a, b]$ is the shortest among all intervals that satisfies condition (1).

Example: Uniform distribution

Suppose that $(X_1, X_2, \dots, X_n)$ are i.i.d. Uniform random variables on $[0, \theta]$. The MLE of θ is $X_{(n)} = \max(X_1, X_2, \dots, X_n)$, the sample maximum. The distribution function of $X_{(n)}/\theta$ is

$$G(x) = x^n \text{ for } 0 \leq x \leq 1$$

Thus $X_{(n)}/\theta$ is a pivotal random variable for θ .


Example: Uniform distribution

To find a $100p\%$ confidence interval for θ , we need to find a and b such that

$$P \left[a \leq \frac{X_{(n)}}{\theta} < b \right] = p$$

There are obviously infinitely many choices for a and b ; however, it can be shown that setting $b = 1$ and $a = (1 - p)^{1/n}$ results in the shortest possible confidence interval using the pivot $X_{(n)}/\theta$, namely $[X_{(n)}, X_{(n)}/(1 - p)^{1/n}]$.

Pivotal distribution for θ in $U(0, \theta)$



Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Sufficient Statistics for λ ?

$$\sum_{i=1}^n X_i$$

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Sufficient Statistics for λ ?

$$\sum_{i=1}^n X_i$$

Distribution of $\sum_{i=1}^n X_i$?

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Sufficient Statistics for λ ?

$$\sum_{i=1}^n X_i$$

Distribution of $\sum_{i=1}^n X_i$?

Sampling from an Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$Y \sim \text{Gamma}(n, \lambda)$$

$$aY \sim \text{Gamma}(n, \lambda/a)$$

Sampling from an Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$Y \sim \text{Gamma}(n, \lambda)$$

$$aY \sim \text{Gamma}(n, \lambda/a)$$

$$Y \sim \text{Gamma}(n, 1/2)$$

$$Y \sim \chi_{2n}^2$$

Sampling from an Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$Y \sim \text{Gamma}(n, \lambda)$$

$$aY \sim \text{Gamma}(n, \lambda/a)$$

$$Y \sim \text{Gamma}(n, 1/2)$$

$$Y \sim \chi^2_{2n}$$

Sampling from an Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$2\lambda \sum_{i=1}^n X_i \sim \text{Gamma}(n, 1/2)$$

Sampling from an Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$2\lambda \sum_{i=1}^n X_i \sim \text{Gamma}(n, 1/2)$$

$$2\lambda \sum_{i=1}^n X_i \sim \text{Gamma}(n, 1/2)$$

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}^2$$

Sampling from an Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$\sum_{i=1}^n X_i \sim \text{Gamma}(n, \lambda)$$

$$2\lambda \sum_{i=1}^n X_i \sim \text{Gamma}(n, 1/2)$$

$$2\lambda \sum_{i=1}^n X_i \sim \text{Gamma}(n, 1/2)$$

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}^2$$

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Distribution of $\sum_{i=1}^n X_i$?

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}^2$$

Distribution of $\sum_{i=1}^n X_i$?

By central limit theorem

$$\frac{\sum_{i=1}^n X_i}{n} \approx N(1/\lambda, 1/n\lambda^2)$$

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Distribution of $\sum_{i=1}^n X_i$?

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}$$

Distribution of $\sum_{i=1}^n X_i$?

By central limit theorem

$$\frac{\sum_{i=1}^n X_i}{n} \approx N(1/\lambda, 1/n\lambda^2)$$

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Distribution of $\sum_{i=1}^n X_i$?

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}$$

Distribution of $\sum_{i=1}^n X_i$?

By central limit theorem

$$\frac{\sum_{i=1}^n X_i}{n} \approx N(1/\lambda, 1/n\lambda^2)$$

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Distribution of $\sum_{i=1}^n X_i$?

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}$$

Distribution of $\sum_{i=1}^n X_i$?

By MLE properties

$$\hat{\lambda}_{MLE} \approx N\left(\lambda, \frac{\lambda^2}{n}\right)$$

$$\frac{1}{\bar{X}} \approx N\left(\lambda, \frac{\lambda^2}{n}\right)$$

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Distribution of $\sum_{i=1}^n X_i$?

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}$$

Distribution of $\sum_{i=1}^n X_i$?

By MLE properties

$$\hat{\lambda}_{MLE} \approx N\left(\lambda, \frac{\lambda^2}{n}\right)$$

$$\frac{1}{\bar{X}} \approx N\left(\lambda, \frac{\lambda^2}{n}\right)$$

Example: Exponential distribution

$$X \sim \text{Exp}(\lambda) \quad f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Distribution of $\sum_{i=1}^n X_i$?

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}$$

Distribution of $\sum_{i=1}^n X_i$?

By MLE properties

$$\hat{\lambda}_{MLE} \approx N\left(\lambda, \frac{\lambda^2}{n}\right)$$

$$\frac{1}{\bar{X}} \approx N\left(\lambda, \frac{\lambda^2}{n}\right)$$


Example: Exponential distribution

$$f(x|\lambda) = \lambda \exp(-\lambda x) \quad x > 0$$

Pivotal quantity $h(X_1, \dots, X_n; \lambda) = 2\lambda \sum_{i=1}^n X_i$

$$2\lambda \sum_{i=1}^n X_i \sim \chi_{2n}^2$$

Pivotal distribution for λ in $Exp(\lambda)$



Example: Exponential distribution

- Exact Intervals:**

$$P\left(\frac{\chi_{2n}^2(\alpha/2)}{2\sum_{i=1}^n X_i} \leq \lambda \leq \frac{\chi_{2n}^2(1-\alpha/2)}{2\sum_{i=1}^n X_i}\right) = 1 - \alpha$$

- Approximated Intervals:**

$$P\left(\bar{X} - z_{1-\alpha/2}\sqrt{\frac{1}{n\hat{\lambda}^2}} \leq 1/\lambda \leq \bar{X} + z_{1-\alpha/2}\sqrt{\frac{1}{n\hat{\lambda}^2}}\right) = 1 - \alpha$$

$$P\left(\frac{1}{\bar{X}} - z_{1-\alpha/2}\sqrt{\frac{\hat{\lambda}^2}{n}} \leq \lambda \leq \frac{1}{\bar{X}} + z_{1-\alpha/2}\sqrt{\frac{\hat{\lambda}^2}{n}}\right) = 1 - \alpha$$

Example: Exponential distribution

A theoretical model suggests that the time to breakdown of an insulating fluid between electrodes at a particular voltage has an exponential distribution with parameter λ . A random sample of $n = 10$ breakdown times yields the following sample data (in minutes):

41.53, 18.73, 2.99, 30.34, 12.33, 117.52, 73.02, 223.63, 4.00, 26.78.

We want to obtain a 95% confidence interval for λ and the average breakdown time $1/\lambda$,

In this problem, $n = 20$, $\alpha = 0.05$, look up the table, we have $\chi^2_{20}(0.975) = 34.17$ and $\chi^2_{20}(0.025) = 9.59$, and $\sum_{i=1}^{10} x_i = 550.87$. Inserting these numbers and we have the 95% confidence interval for λ is $(0.00871, 0.03101)$.

Conclusions for Confidence Intervals

- Confidence Interval (reflects sampling error) should always be reported along with the Point Estimate
- An interpretation of the Confidence Interval estimate should also be provided
- The sample size should be reported
- The confidence level should always be reported