San José State University

Math 161A: Applied Probability & Statistics I

Interval Estimation

Prof. Guangliang Chen

Sec 7.1: Basic properties of confidence intervals

Sec 7.3: Intervals based on a normal population distribution

Sec 7.4 Confidence intervals for the variance of a normal population

Introduction

Last time we started considering the new setting in which **we only know the distribution type**, but not the values of its parameters.

The new goal is to use a random sample to infer about the unknown population parameter. This is called **statistical inference**.

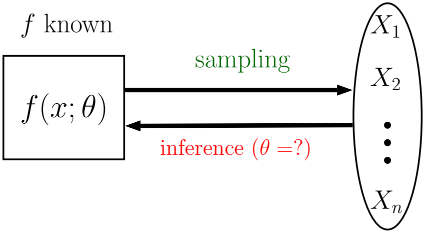

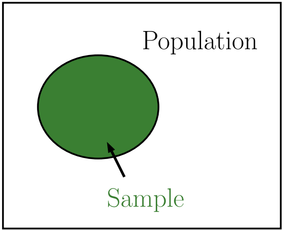

We also mentioned three different kinds of inference tasks

✓ **Point estimation**: What is a single (best) guess of the value of θ?

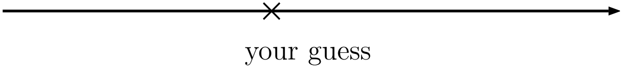

- **Interval estimation**: Can you find an interval to capture the value of θ?

- **Hypothesis testing**: It is claimed that $\theta = \theta_0$ (θ_0 represents a specific number). How do you test the hypothesis based on a random sample from the population?

Recall that mathematically, a **point estimator** $\hat{\theta}$ of θ is a (reasonable) statistic used to estimate θ.

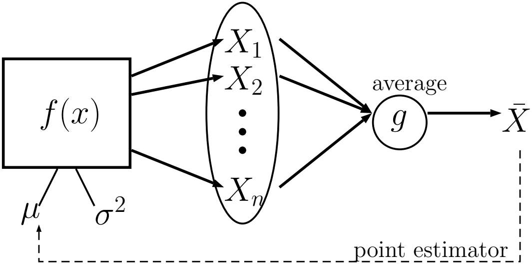

For any specific realization of the random sample, the corresponding value of $\hat{\theta}$ is called a point estimate of θ.

Limitations with point estimation:

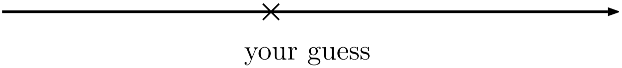

- Point estimates are rarely exactly correct (even when point estimators that are unbiased and have least variance are used).

  For example, for a random sample from the $N(\mu, \sigma^2)$ population, the point estimator $\bar{X}$ of μ is a MVUE. For any small $\epsilon > 0$, the probability that $\bar{X}$ is within a distance of ϵ from μ is

$$P(\mu - \epsilon < \bar{X} < \mu + \epsilon) \approx 2\epsilon f(\mu).$$

- Point estimates provide no error information.

Confidence intervals

Question: Can we make the "coverage probability" much higher than 0?

The answer is yes (by using an interval around $\bar{X}$). One extreme case is

$$P(\mu \in (\bar{X} - \infty, \bar{X} + \infty) = 1)$$

but it is useless.

A more favorable solution is to find a "short" interval with "high" coverage probability:

$$P(\mu \in (\bar{X} - m, \bar{X} + m)) = 1 - \alpha \qquad \text{(for some small } \alpha).$$

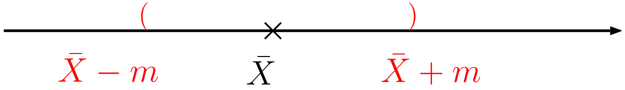

Rewrite as

$$P(\bar{X} - m < \mu < \bar{X} + m) = 1 - \alpha.$$

In the equation,

- μ: population mean (unknown parameter to be estimated)
- $\bar{X}$: sample mean (statistic)
- m: half width (fixed scalar, to be found)
- $1 - \alpha$: coverage probability (specified by user)
- $(\bar{X} - m, \bar{X} + m)$: interval estimator (random)

Task: Given α, find m.

Theorem 0.1. Assume $X_1, \ldots, X_n \overset{iid}{\sim} N(\mu, \sigma^2)$ where μ is unknown, but σ^2 is known. For any given $0 < \alpha < 1$, we have

$$m = z_{\alpha/2} \frac{\sigma}{\sqrt{n}}.$$

Proof. The equation on the preceding slide is equivalent to

$$P(-m < \bar{X} - \mu < m) = 1 - \alpha, \quad \text{or} \quad P\left(-\frac{m}{\sigma/\sqrt{n}} < Z < \frac{m}{\sigma/\sqrt{n}}\right) = 1 - \alpha.$$

This implies that

$$\frac{m}{\sigma/\sqrt{n}} = z_{\alpha/2}, \qquad \text{and accordingly,} \quad m = z_{\alpha/2} \frac{\sigma}{\sqrt{n}}.$$

Interval estimator

We have just obtained that

$$P\left(\mu \in \left(\bar{X} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}}, \bar{X} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\right)\right) = 1 - \alpha.$$

Def 0.1. We call the interval estimator

$$\left(\bar{X} - z_{\alpha/2}\frac{\sigma}{\sqrt{n}}, \bar{X} + z_{\alpha/2}\frac{\sigma}{\sqrt{n}}\right) \equiv \bar{X} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

a $1 - \alpha$ **random interval** for μ. The quantity $m = z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$ is called the **margin of error** of the point estimator $\bar{X}$.

Remark. If $\alpha = 0.05$ (i.e., $1 - \alpha = 0.95$), then $m = 1.96\frac{\sigma}{\sqrt{n}}$.

Confidence interval

Def 0.2. For any specific sample $X_1 = x_1, \ldots, X_n = x_n$ (along with the observed value $\bar{x}$ of $\bar{X}$), the <u>interval estimate</u>

$$\bar{x} \pm z_{\alpha/2} \frac{\sigma}{\sqrt{n}}$$

is called a $1-\alpha$ **confidence interval** for μ. In this setting, $1-\alpha$ is called the **confidence level**.

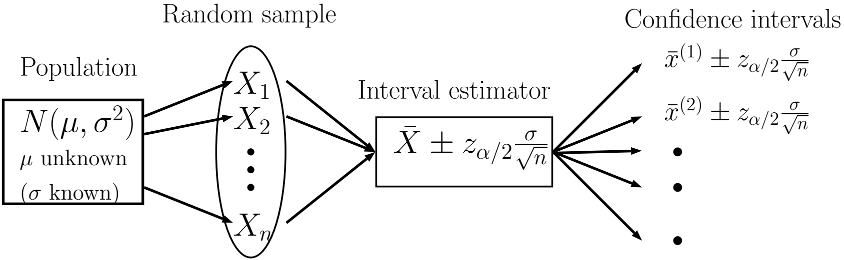

Example 0.1. Recall the brown egg example where $n = 12, \bar{x} = 65.5$ and $\sigma = 2$, a 95% confidence interval is

$$\bar{x} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}} = 65.5 \pm 1.96 \cdot \frac{2}{\sqrt{12}} = 65.5 \pm 1.1 = (64.4, 66.6).$$

Confidence intervals

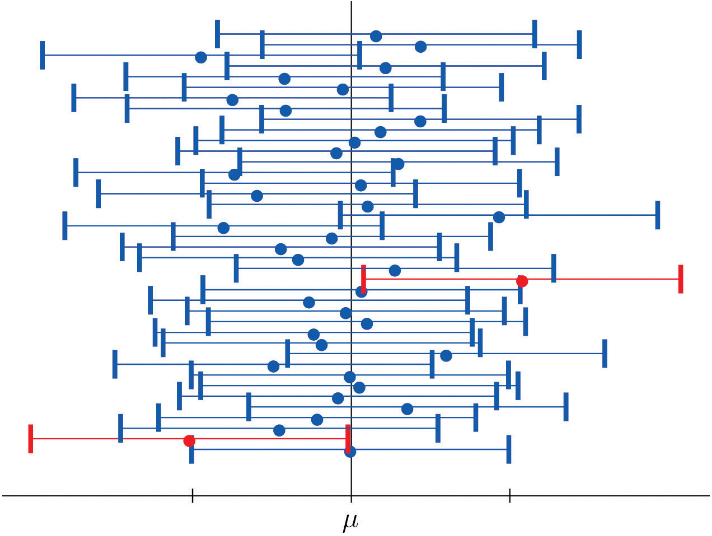

Interpretations of confidence intervals

We can say that

- $(64.4, 66.6)$ is a 95% confidence interval for μ, or
- We are 95% confident that the true μ is contained by this interval (i.e., between 64.4 and 66.6 grams).

We cannot say that

- The probability that μ is contained by this interval is 0.95,

as both μ and this interval are fixed and there is only one truth: "contain" or "not contain". We just do not know which one is true (when μ is unknown).

Confidence is not probability!

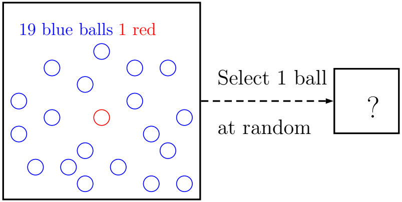

- Probability describes the chance of selecting a blue ball before you actually do it (or if you do it many times)
- Confidence is, after you selected one ball, how certain you believe the ball you got is blue (without looking at it).

Relationship between m and n, α

(m: margin of error, n: sample size, $1-\alpha$: confidence level)

$$m = z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

- The larger the sample size n, the smaller the margin of error m (the shorter the confidence interval);
- The larger the confidence level $1-\alpha$, the bigger the margin of error m (the wider the confidence interval).

Example 0.2 (Continuation of the brown egg example)**.** For another sample from the same population with the same mean $\bar{x} = 65.5$ but a larger size $n = 48$, a 95% confidence interval is

$$\bar{x} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}} = 65.5 \pm 1.96 \cdot \frac{2}{\sqrt{48}} = 65.5 \pm 0.55.$$

How large should the sample size be in order for the margin of error to be 0.2 (at level 95%)?

$$n = \left(z_{\alpha/2}\frac{\sigma}{m}\right)^2 = \left(1.96 \cdot \frac{2}{0.2}\right)^2 = 384.2.$$

The smallest sample size thus is 385.

Example 0.3 (Continuation of the brown egg example)**.** Using the same sample, a 99% confidence interval is

$$\bar{x} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}} = 65.5 \pm 2.576 \cdot \frac{2}{\sqrt{12}} = 65.5 \pm 1.5 = (64.0, 67.0),$$

and a 90% confidence interval is

$$\bar{x} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}} = 65.5 \pm 1.645 \cdot \frac{2}{\sqrt{12}} = 65.5 \pm 0.95$$

Remark. 99% CI $>$ (longer than) 95% CI $>$ 90% CI

What if we do not know σ?

Assuming a normal population $N(\mu, \sigma^2)$, with both μ, σ^2 unknown, we can still construct a $1-\alpha$ confidence intervals for

(1) μ

(2) σ^2

We present the details next.

Confidence interval for μ (when σ is unknown)

Recall when σ was assumed to be known, to derive a $1-\alpha$ confidence interval for μ, we started with

$$P(\bar{X} - m < \mu < \bar{X} + m) = 1 - \alpha$$

and got (after rearranging terms)

$$P(-m < \bar{X} - \mu < m) = 1 - \alpha.$$

In order to solve for m, we then standardized $\bar{X} \sim N(\mu, \sigma^2/n)$:

$$P\left(-\frac{m}{\sigma/\sqrt{n}} < \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} < \frac{m}{\sigma/\sqrt{n}}\right) = 1 - \alpha.$$

When σ is unknown, we can use its estimator S in place of σ: Dividing all sides of the inequalities in the equation

$$P(-m < \bar{X} - \mu < m) = 1 - \alpha.$$

by $S/\sqrt{n}$ gives that

$$P\left(-\frac{m}{S/\sqrt{n}} < \frac{\bar{X} - \mu}{S/\sqrt{n}} < \frac{m}{S/\sqrt{n}}\right) = 1 - \alpha$$

To determine m, we need to know the distribution of the middle quantity. It turns out that it follows a t **distribution with** $n-1$ **degrees of freedom**:

$$\frac{\bar{X} - \mu}{S/\sqrt{n}} \sim t(n-1) = t_{n-1}.$$

Student's t distributions

Def 0.3. The t distribution with ν degrees of freedom is a continuous distribution whose pdf has the following form

$$f(x) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\,\Gamma(\frac{\nu}{2})}\left(1+\frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}, \quad -\infty < x < \infty.$$

Properties:

(1) The graphs are also symmetric, unimodal and bell-shaped.

(2) $\mathrm{E}(X) = 0$.

(3) $\mathrm{Var}(X) = \frac{\nu}{\nu-2}$ (when $\nu > 2$).

(4) $t(\nu) \to N(0,1)$ as $\nu \to +\infty$.

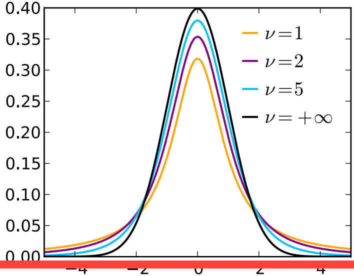

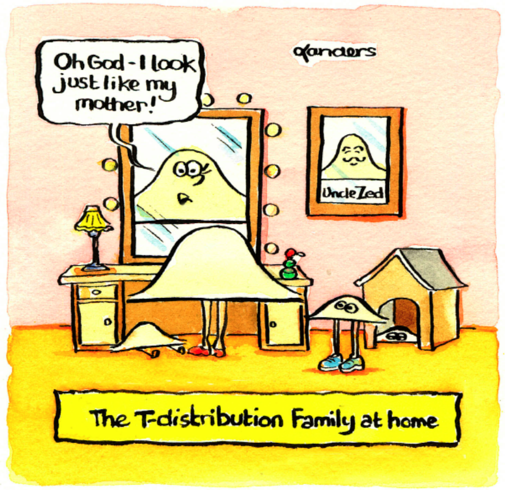
Oh God - I look just like my mother!
©anders
Uncle Zed
The T-distribution Family at home

Confidence interval for μ (when σ unknown)

Theorem 0.2. A $1-\alpha$ confidence interval for μ in the case of a normal population

$$X_1, \ldots, X_n \overset{iid}{\sim} N(\mu, \sigma^2),$$

where σ is unknown, is

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

Remark. Compare with:

$$\bar{x} \pm z_{\alpha/2} \frac{\sigma}{\sqrt{n}} \text{ (when } \sigma \text{ known).}$$

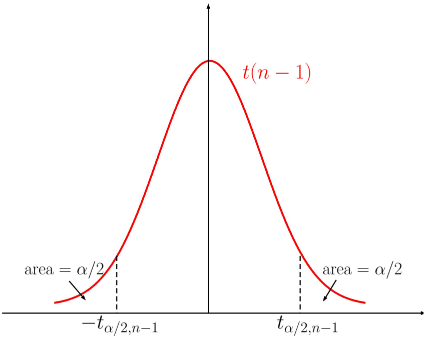

(Use the t table to find the t critical value $t_{\alpha/2,n-1}$)

Example 0.4. In the brown egg example, we selected a sample of 12 eggs (in a carton) and obtained that $\bar{x} = 65.5$ and $s^2 = 4.69$. Assuming normal population (with unknown variance), we obtain a 95% confidence interval

$$\bar{x} \pm t_{\alpha/2,n-1}\frac{s}{\sqrt{n}} = 65.5 \pm t_{0.025,11}\frac{\sqrt{4.69}}{\sqrt{12}} = 65.5 \pm 2.201\sqrt{\frac{4.69}{12}} = 65.5 \pm 1.4.$$

Remark. Previously, when $\sigma = 2$ was used, we obtained the following 95% confidence interval

$$\bar{x} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}} = 65.5 \pm 1.96 \cdot \frac{2}{\sqrt{12}} = 65.5 \pm 1.1,$$

which is shorter. Why?

Confidence interval for σ^2

Assume the same setting of a random sample from a normal population:

$$X_1, \ldots, X_n \overset{\text{iid}}{\sim} N(\mu, \sigma^2),$$

where neither μ nor σ^2 is known.

We already know that

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

is an (unbiased) estimator for σ^2.

We can further use S^2 to construct a $1 - \alpha$ confidence interval for σ^2.

Theorem 0.3. A $1-\alpha$ confidence interval for σ^2 in the case of a normal population $X_1, \ldots, X_n \overset{iid}{\sim} N(\mu, \sigma^2)$ is

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

where $\chi^2_{\alpha/2,\, n-1},\ \chi^2_{1-\alpha/2,\, n-1}$ denote the critical values associated to the chi-square distribution with $n-1$ degrees of freedom:

For any $X \sim \chi^2(n-1)$,

$$P(X > \chi^2_{\alpha/2,\, n-1}) = \alpha/2$$
$$P(X > \chi^2_{1-\alpha/2,\, n-1}) = 1-\alpha/2.$$

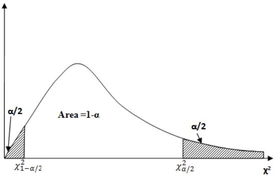

In the brown egg example, suppose we did not know the true value of σ^2. Let us find a 95% confidence interval for σ^2 based on the specific sample we have been using: $n = 12, s^2 = 4.69$.

We need to find the two χ^2 critical values (by using table):

- $\chi^2_{\alpha/2,\, n-1} = \chi^2_{.025,\, 11} = 21.92$;
- $\chi^2_{1-\alpha/2,\, n-1} = \chi^2_{.975,\, 11} = 3.82$.

Therefore, a 95% confidence interval for σ^2 is

$$\left(\frac{(n-1)s^2}{\chi^2_{\alpha/2,\, n-1}},\ \frac{(n-1)s^2}{\chi^2_{1-\alpha/2,\, n-1}} \right) = \left(\frac{11 \cdot 4.69}{21.92},\ \frac{11 \cdot 4.69}{3.82} \right) = (2.35, 13.51).$$

One-sided confidence intervals

Sometimes there is a need for only one-sided confidence intervals:

- Lower confidence bound

$$1 - \alpha = P(\mu > \bar{X} - m)$$

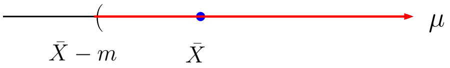

- Upper confidence bound

$$1 - \alpha = P(\mu < \bar{X} + m)$$

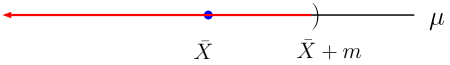

Theorem 0.4. Assuming a random sample $X_1, \ldots, X_n \overset{\text{iid}}{\sim} N(\mu, \sigma^2)$ with unknown μ but known σ^2. Then

- A $1 - \alpha$ lower confidence bound for μ is

$$\mu > \bar{x} - z_\alpha \frac{\sigma}{\sqrt{n}}$$

- A $1 - \alpha$ upper confidence bound for μ is

$$\mu < \bar{x} + z_\alpha \frac{\sigma}{\sqrt{n}}$$

Remark. For each confidence bound, $m = z_\alpha \frac{\sigma}{\sqrt{n}}$.

Proof. We derive only the lower confidence bound (the other part is similar). Rewrite $1-\alpha = P(\mu > \bar{X} - m)$ as

$$1-\alpha = P(\bar{X} - \mu < m)$$

Standardize $\bar{X}$ to have

$$1-\alpha = P\left(\frac{\bar{X}-\mu}{\sigma/\sqrt{n}} < \frac{m}{\sigma/\sqrt{n}}\right)$$

Since $\frac{\bar{X}-\mu}{\sigma/\sqrt{n}} \sim N(0,1)$, we obtain that

$$\frac{m}{\sigma/\sqrt{n}} = z_\alpha \quad \longrightarrow m = z_\alpha \frac{\sigma}{\sqrt{n}}.$$

Consequently, a $1-\alpha$ lower confidence bound for μ is $\bar{X} - z_\alpha \frac{\sigma}{\sqrt{n}}$. □

Example 0.5. In the brown egg example (where $\bar{x} = 65.5, \sigma = 2$), a 95% upper confidence bound for μ is

$$\mu < \bar{x} + z_\alpha \frac{\sigma}{\sqrt{n}} = 65.5 + z_{.05} \frac{2}{\sqrt{12}} = 65.5 + 1.645 \frac{2}{\sqrt{12}} = 66.45.$$

Similarly, a 95% lower confidence bound for μ is

$$\mu > \bar{x} - z_\alpha \frac{\sigma}{\sqrt{n}} = 65.5 - 1.645 \frac{2}{\sqrt{12}} = 64.55.$$

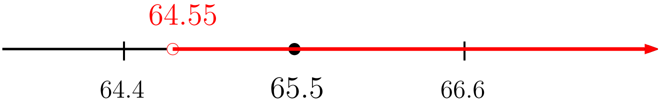

Remark. When σ is unknown, the one-sided confidence intervals for μ can be obtained by using the t distribution instead:

- A $1-\alpha$ lower confidence bound for μ is

$$\mu > \bar{x} - t_{\alpha,n-1}\frac{s}{\sqrt{n}}$$

- A $1-\alpha$ upper confidence bound for μ is

$$\mu < \bar{x} + t_{\alpha,n-1}\frac{s}{\sqrt{n}}$$

Similarly, the one-sided confidence intervals for σ^2 are

- A $1-\alpha$ lower confidence bound for σ^2 is

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

- A $1-\alpha$ upper confidence bound for σ^2 is

$$0 < \sigma^2 < \frac{(n-1)s^2}{\chi^2_{1-\alpha,\,n-1}}$$

Summary

Assume a random sample from a distribution, $X_1, \ldots, X_n \overset{\text{iid}}{\sim} f(x)$, with an unknown parameter θ.

- **Basic concepts**
  - **Interval estimator**: a random interval of the form $\hat{\theta} \pm m = (\hat{\theta} - m, \hat{\theta} + m)$, where m is called the margin of error.
  - A desired property of an interval estimator is the high *coverage probability*:

$$P(\hat{\theta} - m < \theta < \hat{\theta} + m) = 1 - \alpha$$

- For any specific sample, the corresponding specific interval is called a **confidence interval** for θ (at level $1-\alpha$).

- **Important results**
  - For a normal population $N(\mu, \sigma^2)$ with known μ but *known* σ^2, a $1-\alpha$ confidence interval for μ is

$$\bar{x} \pm z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$$

Pay attention to how the margin of error $m = z_{\alpha/2}\frac{\sigma}{\sqrt{n}}$ depends on the sample size n and confidence level $1-\alpha$.

– For a normal population $N(\mu, \sigma^2)$ with both μ, σ^2 *unknown*, a $1-\alpha$ confidence interval for μ is

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

In this case, a $1-\alpha$ confidence interval for σ^2 is

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