Chapter 16

Confidence Intervals

The Basics

Intended Learning Outcomes

- Explain why a point estimate alone is not enough
- Derive the logic of a confidence interval from the sampling distribution
- Interpret the confidence level correctly
- Find critical values z^* for common confidence levels
- Compute a confidence interval for μ when σ is known
- Describe how z^*, n, and σ affect the margin of error
- Determine the sample size needed for a target margin of error
- Use a confidence interval to evaluate a claim about μ
- Transform a CI using a linear or monotonic function

The Running Question

Context: CMHC (Canada Mortgage and Housing Corporation) reports the average monthly rent for a 1-bedroom apartment in London, ON is $\$1{,}200$ per month.

The Running Question

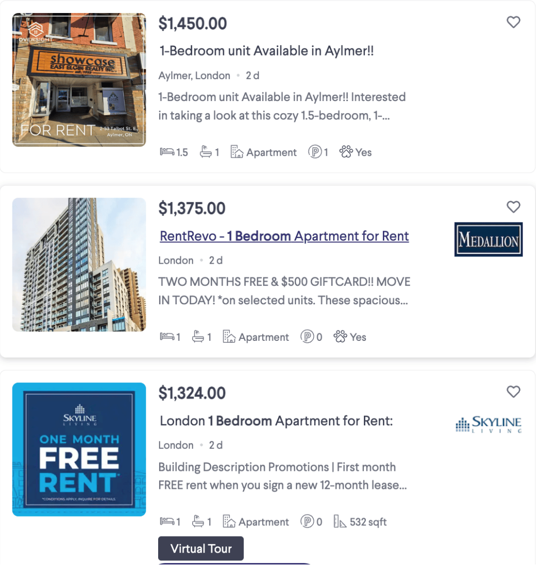

A sample of $n = 47$ current Kijiji listings finds:

$$\bar{x} = \$1{,}467 \qquad \sigma = \$250$$

Is the CMHC figure still accurate? By the end of this chapter, we will have a formal method to answer this question.

PART 1

Why a Point Estimate Is Not Enough

What can a single number really tell us?

The Problem with a Single Number

A university health clinic surveys 100 students about nightly sleep and reports: $\bar{x} = 6.8$ hours.

Sleeping an average of 6 or fewer hours per night is associated with adverse health outcomes. Is this a problem? We can't tell from the point estimate alone.

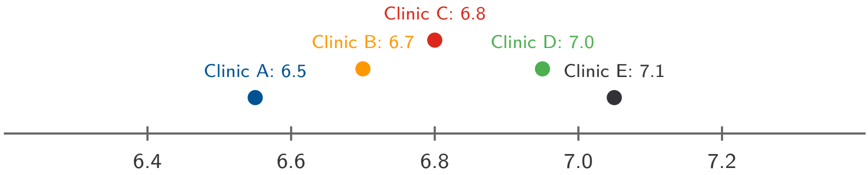

Every sample gives a different $\bar{x}$. A single number tells us nothing about *how close* it is to μ.

What We Really Want

Point estimate alone

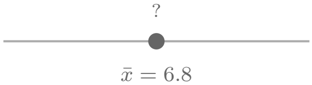

Best single guess for μ,
but no measure of precision.

$\Rightarrow$

Confidence interval

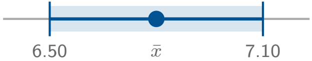

A **range of plausible values**
with a stated confidence level.

> **Principle:** A confidence interval adds precision to the point estimate: a range of plausible values for μ together with a stated confidence level.

PART 2

From Sampling Distribution to Confidence Interval

How does the sampling distribution help us build a range of plausible values?

Simple Conditions for Inference About a Mean

1. **Simple random sample (SRS)** from the population. The population is large relative to the sample.

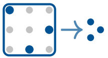

2. The variable has an exactly **Normal distribution** with mean μ and standard deviation σ.

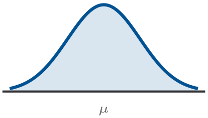

3. We don't know the population mean μ, but we **do know** the population standard deviation σ.

σ ✓

μ ?

Student Sleep Study

Example 16.1

A university health clinic wants to estimate mean nightly sleep among students.

Prior large-scale health surveys give a population standard deviation of $\sigma = 1.5$ hours.

The clinic surveys an SRS of $n = 100$ students and records each participant's nightly sleep duration.

Sample size	$n = 100$
Sample mean	$\bar{x} = 6.8$ hours
Population SD	$\sigma = 1.5$ hours
Goal	Interval estimate of μ

Recall: The Sampling Distribution

From our work on sampling distributions, we know:

If the population is Normal with mean μ and standard deviation σ, then the sample mean $\bar{X}$ from an SRS of size n follows

$$\bar{X} \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$$

Calculate: For the sleep study, the standard error of $\bar{X}$ is:

By the CLT, approximately 95% of sample means will fall within 0.30 units of the true mean μ.

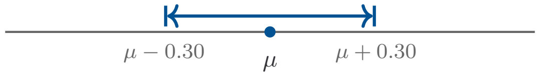

But distance is symmetric, so

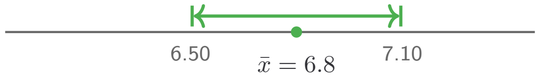

Why Does This Work?

The logic rests on one simple observation:

$$|\bar{x} - \mu| < 0.30 \quad \text{is the same statement as} \quad |\mu - \bar{x}| < 0.30$$

"$\bar{x}$ is close to μ" and "μ is close to $\bar{x}$" are the *same* thing.

Since the sampling distribution guarantees the first statement holds in 95% of samples, the second statement also holds in 95% of samples.

So the interval $\bar{x} \pm 2(\sigma/\sqrt{n})$ captures μ in about 95% of all possible samples.

Confidence Interval

Confidence Interval

A **level C confidence interval** for a parameter has three components:

1. A **point estimate**: a single number calculated from the data that serves as the best guess for the unknown parameter.
2. A **margin of error**: the amount added to and subtracted from the point estimate to form the interval:

$$\text{point estimate} \pm \text{margin of error}$$

3. A **confidence level** C: the probability that the method produces an interval capturing the true parameter in repeated samples.

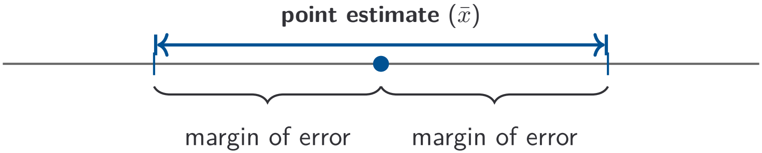

Probability vs. Confidence Language

$\bar{X}$ is a random variable

before data is collected

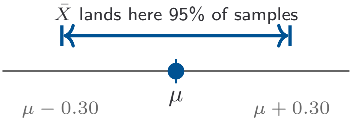

Fixed: μ

Varies: $\bar{X}$

Terminology:

$\Rightarrow$

collect

sample

$\bar{x}$ is a fixed number

after data is collected

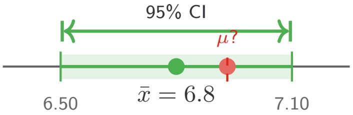

Fixed: $\bar{x} = 6.8$

Unknown: μ

Terminology:

What Does "95% Confident" Mean?

Confidence Level

Repeat many times: $C\%$ of all resulting intervals will contain μ.

❗ Common misconception:

Not: "μ has a 95% chance of being in *this* interval."

Since μ is fixed, it's either captured or it isn't.

The 95% describes the *method*, not this specific interval.

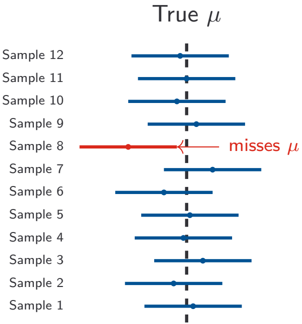

Each row = one sample's 95% CI

What the Interval Tells You

- The true average nightly sleep is **plausibly near the centre**: not 3 hours, not 9 hours, but somewhere around 6.8 hours.
- Values **outside** are inconsistent with your data
- **Width** signals precision:

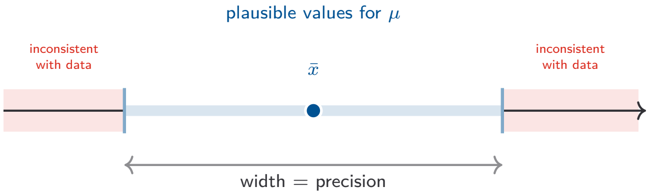

Acting on the Interval

Principle: Because the method is reliable, you are justified in **acting as if** the interval contains μ, not because you have proved it, but because you are using a trustworthy procedure.

Common misconception: Hedging does not mean paralysis. It stops you from **overclaiming**.

Act on the interval; do not treat it as certainty.

Paralysis

Cannot be certain, so refuses to decide

✗

Act on it

Plans and acts based on the interval

✓

Overclaiming

Claims the mean is definitely 6.8 hours

✗

CI in the Wild

Public Health

CI for mean daily sodium intake: $[2{,}800,\ 3{,}400]$ mg

Recommended limit: $2{,}300$ mg. Entire interval exceeds the limit.

Decision: Launch a sodium reduction campaign.

Justified regardless of where exactly in $[2{,}800,\ 3{,}400]$ the truth sits.

Engineering

CI for mean bolt diameter: $[9.97,\ 10.03]$ mm

Tolerance spec: 10.00 ± 0.05 mm. Entire interval within spec.

Decision: Sign off on the batch.

No need to prove the mean is exactly 10.00. The plausible range is acceptable.

CI in the Wild (cont.)

Urban Planning

CI for mean commute time: $[18,\ 26]$ minutes

Schedules bus frequency.
The whole interval shows the system is nowhere near the 45-minute threshold that would trigger a major infrastructure review.

Decision: No need to know whether the true mean is 21 or 23.

Clinical Medicine

CI for mean blood pressure reduction from a new drug: $[4,\ 12]$ mmHg

Even at the low end, 4 mmHg is clinically meaningful.

Decision: Prescribe the drug: acting on the interval, without needing to know whether the true effect is 4 or 12.

The Key Question

None of these researchers are claiming: “The true mean is X.”

They are asking:

Does the entire interval, or enough of it, support a decision?

Sometimes clear

Entire interval lies above or below a threshold

$\rightarrow$ act on it

Sometimes not

Interval straddles the threshold

$\rightarrow$ collect more data

Margin of Error

In the sleep study, find the precise values of the point estimate and margin of error, and write a sentence interpreting the interval in context.

Estimate

Margin of error

Student Sleep Study (cont.)

Example 16.3 — continued

Sleeping 6 or fewer hours per night is associated with increased health risks.

The 95% confidence interval for mean nightly sleep is **approximately** $[6.50,\ 7.10]$ hours.

Use the interval to assess whether the mean nightly sleep duration for the student population falls at or below the at-risk threshold. What can we conclude?

Check Your Understanding

A researcher surveys 40 students and reports a 95% CI for mean weekly study time: $(12.4,\ 15.8)$ hours. The sample mean was $\bar{x} = 14.1$ hours. What quantity are they hoping falls inside the interval?

a) The sample mean $\bar{x} = 14.1$ hours

b) The population mean μ (unknown)

c) The confidence level (95%)

d) The study times of all individual students

Check Your Understanding

A campus transit study reports:

> *We are 95% confident that the mean one-way commute time for Western students is between 18 and 26 minutes.*

Which interpretation is correct?

a) There is a 95% probability that μ lies between 18 and 26 minutes.

b) If we repeated sampling many times, about 95% of the resulting intervals would contain μ.

c) 95% of individual students commute between 18 and 26 minutes.

d) The sample mean falls between 18 and 26 minutes in 95% of samples.

PART 3

Making It Precise

How do we move from “roughly 2 standard deviations” to an exact formula?

The Critical Value z^*

In the sleep example we used “2 standard deviations” for roughly 95% confidence. The exact multiplier depends on the confidence level we want.

Critical Value z^*

The **critical value** z^* is the number such that the standard Normal curve has area C between $-z^*$ and z^*:

$$P(-z^* < Z < z^*) = C$$

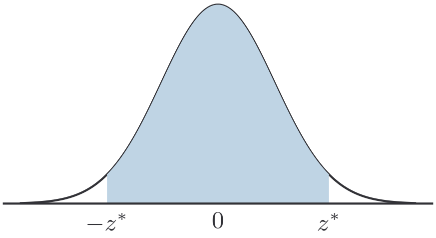

Common Critical Values

Example 16.4

Confidence level C	90%	95%	99%
Critical value z^*			

Principle: Higher confidence $\Rightarrow$ larger z^* $\Rightarrow$ wider interval.

More confidence costs precision.

The Confidence Interval Formula

Context: SRS of size n from a Normal population with unknown mean μ and known standard deviation σ.

Confidence Interval for μ (σ Known)

1. Choose a confidence level C and find the corresponding critical value z^*.
2. Compute the margin of error: $\text{MOE} = z^* \cdot \dfrac{\sigma}{\sqrt{n}}$
3. The level C confidence interval for μ is:

$$\bar{x} \pm \text{MOE} \;=\; \bar{x} \pm z^* \cdot \frac{\sigma}{\sqrt{n}}$$

Applying the Formula: London Rent

Context: $n = 47$, $\bar{x} = \$1{,}467$, $\sigma = \$250$.

Calculate the 95% confidence interval for the true mean rent μ:

Campus Café Wait Times

Example 16.5 — Setup

Context: A campus café uses its ordering system to track how long customers wait during the lunch rush. Historical data show that wait times have a standard deviation of $\sigma = 2.0$ minutes.

On a randomly selected day, an SRS of $n = 64$ customers gives $\bar{x} = 4.2$ minutes.

Compute a 90% confidence interval for the true mean wait time μ.

Practice: Starbucks Cup Volumes

Example 16.6

Context: A Starbucks Grande is supposed to contain 473 mL (16 oz).
The population standard deviation is known to be $\sigma = 6$ mL.
Quality control takes an SRS of 50 cups and finds $\bar{x} = 471.2$ mL.

Task: Construct a 99% CI for the true mean fill volume μ.

Check Your Understanding

Context: An instructor administers a midterm to $n = 401$ students and finds $\bar{x} = 79.5$ with $\sigma = 15$.

Which expression gives the correct 95% confidence interval for the true mean exam score μ?

a) $79.5 \pm 1.96 \times 15$

b) $79.5 \pm (15/20)$

c) $79.5 \pm 1.96 \times 15 \times 20$

d) $79.5 \pm 1.96 \times (15/20)$

PART 4

How Confidence Intervals Behave

What makes a confidence interval wider or narrower?

What Controls the Margin of Error?

We want high confidence and a small margin of error. The margin of error is:

$$\text{MOE} = z^* \cdot \frac{\sigma}{\sqrt{n}}$$

When does the interval get narrower?

> **Applies when:** We can control C and n, but we usually cannot control σ; it is a property of the population.

Effect of Changing the Confidence Level

Example 16.7

Context: Café example with $\bar{x} = 4.2$, $\sigma = 2.0$, $n = 64$.

Task: Compute the margin of error at three confidence levels.

Confidence level	z^*	Margin of error	Interval
90%	1.64		
95%	1.96		
99%	2.58		

Principle: Higher confidence $\Rightarrow$ wider interval.

Check Your Understanding

Context: In computing a 95% confidence interval, you found that the MOE is 1.5. Suppose you instead wanted a 90% confidence interval with the same data. What would the new MOE be?

a) 1.96

b) 1.50

c) 1.64

d) 1.26

e) 0.75

f) Impossible to tell.

Quick Check: Effects on Margin of Error

For each change, does the margin of error **increase** or **decrease**?

1. Increasing the confidence level from 90% to 99%:

2. Increasing the sample size from 50 to 200:

3. Increasing the population standard deviation:

PART 5

Choosing the Sample Size

How many observations do we need?

Solving for Sample Size

If you want a specific margin of error MOE, you can solve for the required sample size n.

Sample Size for a Desired Margin of Error

1. Start with the margin of error formula: $\text{MOE} = z^* \cdot \frac{\sigma}{\sqrt{n}}$
2. Solve for n and round up:
$$n = \left\lceil \left(\frac{z^* \cdot \sigma}{\text{MOE}} \right)^2 \right\rceil$$
3. The ceiling $\lceil\ \rceil$ means **always round up**, because rounding down gives a margin of error slightly above the target.

How Large a Sample?

Example 16.8

Context: The health clinic from the sleep study ($\sigma = 1.5$ hours) wants to estimate mean nightly sleep with 95% confidence.

Problem: What sample size is needed for a margin of error of at most MOE $= 0.2$ hours (about 12 minutes)?

Calculate the required sample size.

Practice: Sample Size

Example 16.9

Context: A campus transportation office wants to estimate the average time students spend looking for parking. Prior studies give $\sigma = 3$ minutes. They want 95% confidence.

Task:

a) What sample size is needed for a margin of error of at most 1 minute?

b) How would the required sample size change if they wanted 99% confidence instead?

Back to the Running Question: Sample Size

Context: In our London rent study, we had $n = 47$ listings and found MOE $\approx \$72$.

How many listings would reduce the margin of error to $\$50$?

PART 6

Using CIs to Evaluate Claims

Can a confidence interval tell us whether a claim about μ is plausible?

Evaluating a Claim with a Confidence Interval

Suppose someone claims the population mean is μ_0. We can use a confidence interval to assess that claim.

If μ_0 is inside the CI:

The claimed value is among the plausible values for μ. The data are **consistent** with the claim.

If μ_0 is outside the CI:

The claimed value is **not** among the plausible values. The data provide evidence **against** the claim.

> **Principle:** If the CI excludes μ_0, the data provide evidence against the claim $\mu = \mu_0$. The CI also shows the full range of plausible values, not just a yes/no verdict.

Juice Bottle Fill Volumes

Example 16.10 — Setup

Context: A juice company labels each bottle as containing 250 mL. A consumer protection agency suspects the company is under-filling bottles. They test an SRS of $n = 36$ bottles and find $\bar{x} = 246$ mL. The population standard deviation is known to be $\sigma = 12$ mL.
Goal: Assess the claim $\mu = 250$ using a confidence interval.
Calculate: the 95% CI for μ.

Determine: Is $\mu = 250$ plausible?

Exercise: What If $\bar{x} = 249$?

Context: Same juice setting ($\mu_0 = 250$, $\sigma = 12$, $n = 36$), but now suppose $\bar{x} = 249$.

a) Build a 95% confidence interval for μ.

b) Does the interval contain $\mu_0 = 250$?

c) Do the data provide evidence against the company's claim?

London Rent and the CMHC Benchmark

Example 16.11 — Setup

In Part 3, we built a 95% CI for average asking rents in London, ON using a Kijiji sample ($n = 47$, $\bar{x} = \$1{,}467$, $\sigma = \$250$).

Our 95% confidence interval was:

$$\text{CI} = \left[\$1{,}395.50,\ \$1{,}538.50\right]$$

CMHC reports the average 1-bedroom rent in London is $\$1{,}200$.

Synthesize: What can we conclude about the CMHC claim based on our confidence interval?

Coin Tosses: A Confidence Interval

Example 16.12 — Callback from Chapter 15

Context: In Chapter 15 (Ex 15.10), we tossed a coin $n = 100$ times and observed $\hat{p} = 0.60$ heads. The population standard deviation is $\sigma = 0.5$.

Calculate: Build a 95% confidence interval for the true proportion of heads p.

Conclude: Is the claim that the coin is fair $(p_0 = 0.50)$ plausible?

PART 7

Transforming Confidence Intervals

What if the quantity of interest is a function of μ?

Transforming a Confidence Interval

Suppose you have a CI for μ_X, but the quantity of interest is $Y = g(X)$, where g is a **monotonic** (always increasing or always decreasing) function.

CI Under a Monotonic Transformation

1. Construct the level C CI for μ_X as usual. Call the endpoints L and U.
2. Apply g to **both endpoints**:

$$g \text{ increasing: } \big(g(L),\ g(U)\big) \qquad\qquad g \text{ decreasing: } \big(g(U),\ g(L)\big)$$

3. The confidence level is **unchanged**.

Principle: When g is decreasing, the endpoints **flip**: the smaller input maps to the larger output.

Linear Transformations of a CI

The most common case: $Y = aX + b$ (scaling and/or shifting).

CI Under a Linear Transformation

1. Build a level C CI for μ_X: $[L,\ U]$.
2. Apply the transformation to **both endpoints**:

$$\text{CI for } \mu_Y \ = \ [aL + b,\ \ aU + b]$$

3. If $a < 0$, the endpoints **swap order**.

Intuition: A linear function stretches and shifts the number line uniformly, so the confidence level is preserved exactly. The new standard deviation is $\sigma_Y = |a|\,\sigma_X$, which the transformation handles automatically.

Common examples: unit conversions (Celsius to Fahrenheit), rate calculations (steps to calories), currency exchange.

Daily Steps and Caloric Expenditure

Example 16.14 — Building the CI

Context: A campus health researcher uses fitness trackers to study physical activity.

An SRS of $n = 64$ students records daily step counts. Large-scale health surveys give $\sigma = 2{,}400$ steps/day. The sample mean is $\bar{x} = 7{,}500$ steps/day.

Construct a 95% CI for the mean daily step count μ_S.

Daily Steps and Caloric Expenditure

Example 16.14 — Transforming the CI

Metabolic studies show students burn about **0.04 calories per step**, so $C = 0.04 \times S$.

From Part (a): 95% CI for mean steps is $[6{,}912,\ 8{,}088]$.

b) Use $C = 0.04 \times S$ to construct a 95% CI for the mean daily caloric expenditure μ_C.

c) A campus wellness program claims students burn on average 300 calories per day from walking. Do the data support this?

Fleet Fuel Efficiency

Example 16.15 — Setup

Context: A municipal fleet manager measures fuel consumption X (L/100 km) for city buses. To report to council, efficiency must be expressed in km/L:

$$Y = \frac{100}{X}$$

An SRS of $n = 49$ buses gives $\bar{x} = 8$ L/100 km. Industry fleet data give $\sigma = 1.4$ L/100 km.

a) Construct a 95% CI for the mean fuel consumption μ_X.

b) Use $Y = 100/X$ to construct a 95% CI for the fleet's mean fuel efficiency in km/L.

c) The manufacturer advertises an efficiency of 12.5 km/L. Do the data support this claim?

Fleet Fuel Efficiency

Example 16.15 — Solution

CHAPTER 16

Summary

Key formulas and concepts from this chapter

Chapter 16 Summary — Formulas & Reference

Key formulas

- CI for μ: $\bar{x} \pm z^* \frac{\sigma}{\sqrt{n}}$
- Margin of error: $\text{MOE} = z^* \frac{\sigma}{\sqrt{n}}$
- Sample size: $n = \left\lceil \left(\frac{z^* \sigma}{\text{MOE}} \right)^2 \right\rceil$

Transformations

- If CI for μ is (L, U), then CI for $g(\mu)$ is

$$\big(g(L),\ g(U)\big)$$

when g is increasing, and

$$\big(g(U),\ g(L)\big)$$

when g is decreasing.

Common z^* values

- 90%: $z^* = 1.64$
- 95%: $z^* = 1.96$
- 99%: $z^* = 2.58$

Chapter 16 Summary — Key Ideas

Core concepts

- A point estimate alone gives no sense of precision; a CI adds that information
- The confidence *level* is the method's long-run success rate, not the probability for one interval
- Higher confidence $\Rightarrow$ wider interval
- Larger n $\Rightarrow$ narrower interval

Evaluating claims

- If $\mu_0 \notin$ CI, the data provide evidence against the claim $\mu = \mu_0$
- If $\mu_0 \in$ CI, we do not have evidence against $\mu = \mu_0$ (but we have not proven it)

Comprehensive Practice

Example 16.13

Context: A university claims the average student commute is 25 minutes. The population standard deviation is $\sigma = 10$ min. An SRS of $n = 100$ students gives $\bar{x} = 27$.

1. What is the 95% confidence interval?
2. Do the data provide evidence against the university's claim?
3. Would a 99% confidence interval change your conclusion? Compute it.
4. What sample size is needed for a margin of error of at most 2 minutes (at 95% confidence)?

Practice: Water Quality Monitoring

Example 16.16

Context: Ontario's Ministry of the Environment monitors dissolved oxygen (DO) in a lake used for recreation. Long-term sensor data give $\sigma = 1.2$ mg/L.

An SRS of $n = 36$ water samples yields $\bar{x} = 7.8$ mg/L.

1. Construct a 95% confidence interval for the mean dissolved oxygen level.
2. The provincial standard for supporting aquatic life is 8.0 mg/L. Do the data provide evidence that the lake's mean DO falls below this standard?

Practice: Varsity Vertical Jump

Example 16.17

Context: A sports science lab measures vertical jump heights for varsity basketball players. A national database gives $\sigma = 8$ cm.

An SRS of $n = 25$ players on a university team gives $\bar{x} = 62$ cm.

1. Construct a 95% confidence interval for the team's mean vertical jump.
2. The national average for varsity players is 58 cm. Do the data provide evidence that this team jumps higher than the national average?
3. How many players would need to be tested to achieve a margin of error of at most 2 cm?

Practice: Greenhouse Temperature

Example 16.18

Context: A greenhouse monitors hourly air temperature (°C) for tomato seedlings. Long-term sensor data give $\sigma = 2.5°$C.

An SRS of $n = 25$ hourly readings gives $\bar{x} = 23°$C.

1. Construct a 95% CI for the mean temperature in °C.
2. The seed packet lists the ideal temperature in Fahrenheit: $F = 1.8\,C + 32$. Transform the CI to °F.
3. The recommended mean temperature is 72°F. Do the data support the greenhouse meeting this target?