

Chapter 14 Introduction to Inference

Statistical Inference

Situation: We are interested in estimating some parameter (population mean, μ) that is unknown. We take a random sample from this population.

Goal: Draw **statistical inference** about the population parameter from the sample data.

Simple conditions for inference about a mean:

1. We have an SRS from the population of interest.
2. The variable we measure has an exactly normal distribution $\text{Normal}(\mu, \sigma)$ in the population.
3. The population mean μ is unknown, but the population standard deviation σ is known.

Confidence Intervals

The reasoning of Statistical Estimation

Example: Beetle cars

Suppose EPA wants to estimate μ , the average CO₂ emitted by all Beetle cars. 49 Beetle cars are randomly selected at the VW plant in Detroit and tested for CO₂ emissions. The test results show these 49 cars have an average of 1.5 grams CO₂ emission per mile.

What is our best guess about the unknown parameter μ ?

But we need to be careful in making a conclusion about μ because

- We know any other random sample of 49 cars would give a different value of

We do not expect \bar{x} to be exactly equal to μ , so we want to say how accurate this estimate is.

So instead of just giving one number (the value of \bar{x}) as our estimate of μ , it seems more desirable to give an _____ of values that may contain μ with some degree of _____.

Let's try to find an interval of values within which we can say that average CO₂ emission (μ) of Beetle cars lies with some high degree of confidence, say 95%.

For this, let's recall from Chapter 11 how \bar{x} behaves in repeated sampling.

In the "Beetle cars" example, suppose we repeatedly sample 49 cars and for each sample note their mean CO₂ emission. We know that the distribution of sample mean, \bar{x} is then

Assume we know that the standard deviation (σ) of CO₂ emissions of Beetle cars is 0.84 (it is unrealistic to assume σ to be known, we would get rid of this assumption in Chapter 17!).

Then,

Using the 68-95-99.7 rule, we know that \bar{x} will fall within 2 standard deviations $2\left(\frac{\sigma}{\sqrt{n}}\right) =$ _____ of the population mean μ with probability

This means that in 95% of all samples, the observed mean score \bar{x} will be within _____ points of the population mean μ . (Note: μ is unknown and fixed - it doesn't vary from sample to sample).

To say that \bar{x} is within 0.24 points of μ is equivalent to saying that μ is within 0.24 points of observed \bar{x} . This happens in 95% of all samples.

- ❖ Combining these facts we can say, "In 95% of all samples (of size $n = 49$) the true but unknown mean μ lies in the interval ". We can rewrite this interval as

Recall that mean of our sample was 1.5 grams. So we say that "*We are 95% confident that*

This interval we just calculated is a *95% confidence interval* for the unknown average CO₂ emission (μ) of all Beetle cars.

In general, **confidence intervals** for any parameter consists of two parts:

- 1) An interval calculated from the data, usually of the form

The **margin of error** conveys how accurate we believe our guess of the true parameter value is, based on the variability of the estimate.

- 2) A **confidence level**, C , gives the probability that the random interval captures the true parameter value in repeated samples. (Note that it is NOT the probability that any one specific interval calculated from a random sample captures the true parameter.)

Users can choose the confidence level, usually 90% or higher because we want to be quite sure of our conclusions. The most common confidence level is 95%.

Many types of confidence intervals exist for various kinds of parameters.

Ch. 14: Confidence Intervals for population mean μ (when population standard deviation σ is known).

Ch. 17: Confidence Intervals for population mean μ (when population standard deviation σ is unknown).

Ch. 18: Confidence Intervals for difference between two population means.

Ch. 19: Confidence Intervals for population proportion.

Ch. 20: Confidence Intervals for comparing two proportions.

We will concentrate on confidence intervals for the mean μ of a population in this chapter.

Interpreting a confidence interval

Note: We don't know if any of the confidence interval in the Beetle cars example contain μ or not! Then what do we mean by confidence?

The meaning of "Confidence": When we say "95% confident", we mean that

The confidence level is

IMPORTANT: We do not know whether the 95% confidence interval from a particular sample is one of the 95% that contain μ or one of the unlucky 5% that miss.

Our confidence is in the procedure, not in any one specific interval.

So it is **completely wrong to say:**

(because any one interval either contains the parameter or not, there is no randomness in it!)

Remember that probability (chance) is associated only with a random phenomenon. After you have constructed a CI from a random sample, there is no randomness left in it. Hence it doesn't make sense to attach any probability statement to a specific (numerical) CI.

Confidence Intervals for a Population Mean μ (standard deviation is known)

Confidence Intervals for a Population Mean

Choose an SRS of size n from a population having unknown mean μ and known standard deviation σ .

A level C confidence interval for μ is

Here, z^* is the value on the *standard normal curve* with area C between $-z^*$ and z^* . The interval is exact when the original population distribution is normal and is approximately correct for large n otherwise.

The most commonly used confidence levels are

C	90%	95%	99%
z^*			

z^* for other confidence levels can be found similarly from Table A or more conveniently from Table C.

Numbers like z^* that mark off specified areas are called **critical values** of the standard normal distribution.

Confidence intervals: The four-step process

- (1) STATE: What is the practical question that requires estimating a parameter?
- (2) PLAN: Identify the parameter, choose a level of confidence, and select the type of confidence interval that fits your situation.
- (3) SOLVE: Carry out the work in two phases:
 1. Check the conditions for the interval you plan to use.
 2. Calculate the confidence interval.
- (4) CONCLUDE: Return to the practical question to describe your results in this setting.

Ex: Beetle Cars

Recall that the CO₂ emissions of Beetle cars have a mean μ and standard deviation $\sigma = 0.84$. Our sample mean was $\bar{x} = 1.5$ and $n = 49$. We had calculated the 95% confidence interval for the mean CO₂ emission as

- a. Find a 90% confidence interval for the mean CO₂ emission. Follow the four-step process.

STATE:

PLAN:

SOLVE:

CONCLUDE:

- b. Find a 99% confidence interval for the mean CO₂ emission. Follow the four-step process.

STATE:

PLAN:

SOLVE:

CONCLUDE:

Tests of Significance

Statistical inference provides methods for drawing conclusions about a population from sample data.

Two of the most common types of statistical inference:

1) **Confidence intervals**

Goal is to estimate a population parameter.

2) **Tests of Significance**

Goal is to assess the evidence provided by the data about some claim concerning the population.

Basic Idea of Tests of Significance

The reasoning of statistical tests, like that of confidence intervals, is based on asking what would happen if we repeated the sample or experiment many times.

Example

Each day Tom and Mary decide who pays for lunch based on a toss of Tom's favorite quarter.

Heads - Tom pays

Tails - Mary pays

Tom says that tossing quarter has an even chance of landing heads/tails.

Mary thinks she pays more often.

Mary steals the quarter, tosses it 10 times, gets 7 tails (70% tails).

She is furious and claims that tossing this coin gives unbalanced results.

There are two possibilities:

1. Tom is telling truth – the chance of tails is 50% and the observed 7 tails out of 10 tosses was only due to sampling variability.
2. Tom is lying – the chance of tails is greater than 50%.

Suppose they call you to decide between 1 and 2 (maybe because they realized that they need a statistician to solve this problem!!).

To be fair to both of them, you toss the quarter 25 times. Suppose you get 21 tails.

If the coin is a fair coin, the actual probability of getting greater than or equal to 21 tails in 25 tosses is _____.

What would you conclude? Why?

Moral of the story: an outcome that would rarely happen if a claim were true is good evidence that the claim is not true.

Tests of Significance

Stating hypotheses

A **hypothesis** is a statement about the parameters in a population, e.g.

State your research question as two hypotheses, the null, and the alternative hypotheses. Remember that these are written in terms of the population parameters!!

The **null hypothesis** (H_0) is the statement being tested. This is assumed “true” and compared to the data to see if there is evidence against it. Typically, H_0 is a statement of “no difference” or “no effect”.

Suppose we want to test the null hypothesis that μ is some specified value, say μ_0 . Then

H_0 :

(**Note:** We would always express H_0 using equality sign)

The **alternative hypothesis** (H_a) is the statement about the population parameter that we hope or suspect is true. We are interested to see if the data supports this hypothesis.

H_a can be **one-sided** (e.g.) or **two-sided** (e.g.).

Statistical Significance

A **significance test** is a formal procedure for comparing observed data with a hypothesis whose truth we want to assess.

The results of a test are expressed in terms of a probability that measures how well the data and hypothesis agree.

Significance level (usually represented as α) is the value of probability below which we start consider significant differences. Typical α levels used are 0.1, 0.05 and 0.01.

Tests from Confidence Intervals

Issue: Is there any correspondence between test and CI?

A confidence interval for the population mean μ tells us which values of μ are plausible (those inside the interval) and which values are not plausible (those outside the interval) at the chosen level of confidence.

We can use this idea to carry out a test of any null hypothesis $H_0: \mu = \mu_0$ against a two-sided alternative hypothesis $H_a: \mu \neq \mu_0$.

Correspondence between CI & two-sided test

A level α two-sided test rejects a hypothesis $H_0: \mu = \mu_0$ exactly when the value μ_0 falls _____ a level $(1-\alpha)$ CI for μ .

Tests of significance (two-sided) from confidence intervals: The four-step process

- (1) STATE: What is the practical question that requires estimating a parameter?
- (2) PLAN: Identify the parameter of interest, state null and alternative hypotheses, fix the significance level α and choose the type of test (confidence interval) that fits your situation.
- (3) SOLVE: Carry out the work in three phases:
 1. Check the conditions for the test you plan to use.
 2. Calculate the corresponding level $100(1-\alpha)\%$ confidence interval.
 3. Reject H_0 if the hypothesized value under H_0 is outside the confidence interval in phase 2, otherwise, do not reject H_0 .
- (4) CONCLUDE: Return to the practical question to describe your results in this setting.

Inference for μ when σ is known

Ex: Home Depot sells concrete blocks. The store manager wants to estimate the average weight of all blocks in stock. A simple random sample of 64 blocks has a mean weight of 65.5 lbs. Assume that the weights of blocks are normally distributed with standard deviation 4.6 lbs.

- (a) Make a 95% CI for the mean weight of all blocks.
- (b) The store manager is interested in knowing if the mean weight of all blocks is 68 lbs or not (at 5% level). State the appropriate hypotheses. Using the above CI what do you conclude? Follow the four-step process.

STATE:

PLAN:

SOLVE:

CONCLUDE: