Lesson 11:
Chapter 7 Section 1

Caleb Moxley

BSC Mathematics

12 October 15

§7.1 Inference for the Mean of a Population

When we first started looking at graphical representations of data, we first described a single data set and then we compared the graphs of two data sets. We did similar things with probability distributions — we looked at one distribution, and then we compared distributions.

§7.1 Inference for the Mean of a Population

When we first started looking at graphical representations of data, we first described a single data set and then we compared the graphs of two data sets. We did similar things with probability distributions — we looked at one distribution, and then we compared distributions. It only makes sense that our investigation of statistical inference will involve this same process: First we'll look at claims about a single population, and then we'll move on to claims about two populations.

§7.1 Inference for the Mean of a Population

When we first started looking at graphical representations of data, we first described a single data set and then we compared the graphs of two data sets. We did similar things with probability distributions — we looked at one distribution, and then we compared distributions. It only makes sense that our investigation of statistical inference will involve this same process: First we'll look at claims about a single population, and then we'll move on to claims about two populations. Investigating claims about more than two populations involves sophisticated statistical inference tests which we'll look at in Chapters 12 & 13.

§7.1 Inference for the Mean of a Population

The claims we're going to investigate in this chapter involve population means. We'll be able to answer questions like these below.

§7.1 Inference for the Mean of a Population

The claims we're going to investigate in this chapter involve population means. We'll be able to answer questions like these below.

- Is a population mean greater than, less than, or equal to a particular value?

§7.1 Inference for the Mean of a Population

The claims we're going to investigate in this chapter involve population means. We'll be able to answer questions like these below.

- Is a population mean greater than, less than, or equal to a particular value?
- Is one population's mean larger than, less than, or equal to another population's mean?

§7.1 Inference for the Mean of a Population

The claims we're going to investigate in this chapter involve population means. We'll be able to answer questions like these below.

- Is a population mean greater than, less than, or equal to a particular value?
- Is one population's mean larger than, less than, or equal to another population's mean?
- Is this population's standard deviation the same as another population's standard deviation?

§7.1 Inference for the Mean of a Population

The claims we're going to investigate in this chapter involve population means. We'll be able to answer questions like these below.

- Is a population mean greater than, less than, or equal to a particular value?
- Is one population's mean larger than, less than, or equal to another population's mean?
- Is this population's standard deviation the same as another population's standard deviation?

Because we'll be working with populations we ostensibly know little about, we won't be able to assume that we know the standard deviation of the population σ

§7.1 Inference for the Mean of a Population

The claims we're going to investigate in this chapter involve population means. We'll be able to answer questions like these below.

- Is a population mean greater than, less than, or equal to a particular value?
- Is one population's mean larger than, less than, or equal to another population's mean?
- Is this population's standard deviation the same as another population's standard deviation?

Because we'll be working with populations we ostensibly know little about, we won't be able to assume that we know the standard deviation of the population σ — this means our sampling distributions cannot be assumed to be $N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$ via the Central Limit Theorem!

§7.1 Inference for the Mean of a Population

When we don't know σ, we estimate it by s — the standard deviation of the sample. We might naïvely try to use $N\left(\mu, \frac{s}{\sqrt{n}}\right)$ as the sampling distribution of the mean, but this is not the sampling distribution when we use the **standard error** $\frac{s}{\sqrt{n}}$ instead of $\frac{\sigma}{\sqrt{n}}$. The sampling distribution is a t-distribution with $n-1$ degrees of freedom!

§7.1 Inference for the Mean of a Population

When we don't know σ, we estimate it by s — the standard deviation of the sample. We might naïvely try to use $N\left(\mu, \frac{s}{\sqrt{n}}\right)$ as the sampling distribution of the mean, but this is not the sampling distribution when we use the **standard error** $\frac{s}{\sqrt{n}}$ instead of $\frac{\sigma}{\sqrt{n}}$. The sampling distribution is a t-distribution with $n-1$ degrees of freedom! Thus we have the following.

§7.1 Inference for the Mean of a Population

When we don't know σ, we estimate it by s — the standard deviation of the sample. We might naïvely try to use $N\left(\mu, \frac{s}{\sqrt{n}}\right)$ as the sampling distribution of the mean, but this is not the sampling distribution when we use the **standard error** $\frac{s}{\sqrt{n}}$ instead of $\frac{\sigma}{\sqrt{n}}$. The sampling distribution is a t-distribution with $n-1$ degrees of freedom! Thus we have the following.

Definition (one-sample t test statistic)

Suppose that a SRS of size n is drawn from a normally distributed population with mean μ. Then the **one-sample t statistic**

$$t = \frac{\bar{x} - \mu}{s/\sqrt{n}}$$

is $t(k)$ distributed, where $k = n - 1$ is the degrees of freedom.

§7.1 Inference for the Mean of a Population

A couple of notes on the t distribution with k degrees of freedom:

§7.1 Inference for the Mean of a Population

A couple of notes on the t distribution with k degrees of freedom:

- The $t(k)$ distribution is very nearly normal.

§7.1 Inference for the Mean of a Population

A couple of notes on the t distribution with k degrees of freedom:

- The $t(k)$ distribution is very nearly normal.
- It is a little flatter than the standard normal distribution.

§7.1 Inference for the Mean of a Population

A couple of notes on the t distribution with k degrees of freedom:

- The $t(k)$ distribution is very nearly normal.
- It is a little flatter than the standard normal distribution.
- As $k \to \infty$, $t(k)$ approximates a standard normal distribution.

§7.1 Inference for the Mean of a Population

A couple of notes on the t distribution with k degrees of freedom:

- The $t(k)$ distribution is very nearly normal.
- It is a little flatter than the standard normal distribution.
- As $k \to \infty$, $t(k)$ approximates a standard normal distribution.

The t distribution's extra variation is due to the increased variability caused by substituting a random s for a fixed σ.

§7.1 Inference for the Mean of a Population

A couple of notes on the t distribution with k degrees of freedom:

- The $t(k)$ distribution is very nearly normal.
- It is a little flatter than the standard normal distribution.
- As $k \to \infty$, $t(k)$ approximates a standard normal distribution.

The t distribution's extra variation is due to the increased variability caused by substituting a random s for a fixed σ.

By the way, you can thank beer (specifically Guinness) for the t distribution.

§7.1 Inference for the Mean of a Population

Once you have the t test statistic, you can find P-values for significance tests.

§7.1 Inference for the Mean of a Population

Once you have the t test statistic, you can find P-values for significance tests.

Definition (one-sample P-value tests for mean with unknown σ)

To test the hypothesis $H_0 : \mu = \mu_0$ based on a SRS of size n from a population with unknown μ and unknown σ, compute the **test statistic**

$$t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}}.$$

For a $t(n-1)$ RV, the P-value for a test of H_0 against

$$\begin{array}{lcr} H_a : \mu > \mu_0 & \text{is} & P(T > t) \\ H_a : \mu, < \mu_0 & \text{is} & P(T < t) \\ H_a : \mu \neq \mu_0 & \text{is} & 2P(T > |t|) \end{array}$$

§7.1 Inference for the Mean of a Population

In the previous test, these P-values are exact when the underlying population is normally distributed and are approximate when the underlying population is not normally distributed but the sample size n is large.

§7.1 Inference for the Mean of a Population

In the previous test, these P-values are exact when the underlying population is normally distributed and are approximate when the underlying population is not normally distributed but the sample size n is large.

Basically, this process is exactly the same as it was for populations with known σ, but we simply use a different test statistic.

In the previous test, these P-values are exact when the underlying population is normally distributed and are approximate when the underlying population is not normally distributed but the sample size n is large.

Basically, this process is exactly the same as it was for populations with known σ, but we simply use a different test statistic. This will be our theme throughout most of the course: we use the same basic hypothesis tests (P-value, critical value, and/or confidence interval), but we change the way we compute the test statistic and the distributions we use.

§7.1 Inference for the Mean of a Population

What's changed in our computation of the confidence interval?

§7.1 Inference for the Mean of a Population

What's changed in our computation of the confidence interval?

Definition (one-sample t confidence interval)

Suppose a SRS is drawn from a population having unknown μ and unknown σ. A level C confidence interval for μ is

$$\bar{x} \pm t^* \frac{s}{\sqrt{n}},$$

where t^* is the value for the $t(n-1)$ density curve having area C between $-t^*$ and t^*. We call $t^* \frac{s}{\sqrt{n}}$ the margin of error.

Again, these intervals are exact when the population is normal and are approximate when the population is not normal but the sample size is large.

§7.1 Inference for the Mean of a Population

And finally, we can also conduct a critical value test.

§7.1 Inference for the Mean of a Population

And finally, we can also conduct a critical value test.

Definition (one-sample t critical value test)

Suppose a SRS is drawn from a population having unknown μ and unknown σ. A level α critical value (denoted t^*) for μ is the value satisfying

- $P(T < t^*) = \alpha$ if $H_a : \mu < \mu_0$.
- $P(T > t^*) = \alpha$ if $H_a : \mu > \mu_0$.
- $P(T > t^*) = \alpha/2$ if $H_a : \mu \neq \mu_0$.

And finally, we can also conduct a critical value test.

Definition (one-sample t critical value test)

Suppose a SRS is drawn from a population having unknown μ and unknown σ. A level α critical value (denoted t^*) for μ is the value satisfying

- $P(T < t^*) = \alpha$ if $H_a : \mu < \mu_0$.
- $P(T > t^*) = \alpha$ if $H_a : \mu > \mu_0$.
- $P(T > t^*) = \alpha/2$ if $H_a : \mu \neq \mu_0$.

We can use these critical values to conduct a critical value test.

Example (monthly rates of return on a portfolio)

An investment owner is unhappy with the performance of his investment and believes that it's being managed poorly. He would like to compare the average rate of return month-over-month with the rate of return of the S&P 500, which was 4.5% annually. With 5% significance, he tests the claim that his investment performed worse than the S&P. A sample of 36 month-over-month returns had an average of 4.35% annually with standard deviation 1.25% annually. Does this sample support the claim?

Example (monthly rates of return on a portfolio)

An investment owner is unhappy with the performance of his investment and believes that it's being managed poorly. He would like to compare the average rate of return month-over-month with the rate of return of the S&P 500, which was 4.5% annually. With 5% significance, he tests the claim that his investment performed worse than the S&P. A sample of 36 month-over-month returns had an average of 4.35% annually with standard deviation 1.25% annually. Does this sample support the claim?

The P-value for this test is 0.2382, which is greater than α. Thus, we don't reject H_0, and the data doesn't support his claim. There's not enough evidence to say that the portfolio is doing worse than the S&P.

Example (changing means)

A candy company started using a new type of sweetener, and they were concerned that they might have to change the nutritional label to reflect this. They would like to know if the average amount of calories has changed from the previous average of 250 per package. To check this claim, they take a SRS of 49 candy packs. They get the data linked below. Construct a 95% CI for the mean μ of the number of calories in each package.

Data Set 1

Example (changing means)

A candy company started using a new type of sweetener, and they were concerned that they might have to change the nutritional label to reflect this. They would like to know if the average amount of calories has changed from the previous average of 250 per package. To check this claim, they take a SRS of 49 candy packs. They get the data linked below. Construct a 95% CI for the mean μ of the number of calories in each package.

Data Set 1

The CI is $(242.6, 249.3)$. What does this suggest about the new mean as compared to the old mean?

Example (changing means)

The time to drive between two cities has a normal distribution. A trucking company claims that the average time to drive between these two cities is no more than 3 hours. They want to check that this claim is still true after a change to the traffic signals is implemented. They take a SRS of 41 drive times after the changes are made. This sample had a mean of 2.91 hours and a standard deviation of 0.25 hours. Does this sample support the claim with $\alpha = 0.02$?

§7.1 Inference for the Mean of a Population

Example (changing means)

The time to drive between two cities has a normal distribution. A trucking company claims that the average time to drive between these two cities is no more than 3 hours. They want to check that this claim is still true after a change to the traffic signals is implemented. They take a SRS of 41 drive times after the changes are made. This sample had a mean of 2.91 hours and a standard deviation of 0.25 hours. Does this sample support the claim with $\alpha = 0.02$?

Well, the critical value for this (right-tailed) test is $t^* = 2.123$. And the test statistic is $t = -2.305$, which is not more extreme than t^*, so we fail to reject H_0 and support our claim.

§7.1 Inference for the Mean of a Population

Some populations we investigate could be paired in some way! Think about this: When we do a study that measures body temperature after a short exercise, what might be the population of this study, i.e. what data are we really looking at?

§7.1 Inference for the Mean of a Population

Some populations we investigate could be paired in some way! Think about this: When we do a study that measures body temperature after a short exercise, what might be the population of this study, i.e. what data are we really looking at? We're actually looking at the **differences** in the body temperatures of subjects before an after the short exercise.

Some populations we investigate could be paired in some way! Think about this: When we do a study that measures body temperature after a short exercise, what might be the population of this study, i.e. what data are we really looking at? We're actually looking at the **differences** in the body temperatures of subjects before an after the short exercise. This is a common way of conducting a study. It's called a matched-pair procedure, and it could be used in a t or a z setting, but we'll discuss it in the t setting here.

Example

Data Set 2 below contains the heart rates of subjects before and after walking briskly for 10 minutes. Create a 90% CI for the mean of the differences, where the difference is the new heart rate minus the old. Assume these differences are normally distributed. Use this 90% CI to test the claim that the mean of the differences is more than 20 beats per minute.

Example

Data Set 2 below contains the heart rates of subjects before and after walking briskly for 10 minutes. Create a 90% CI for the mean of the differences, where the difference is the new heart rate minus the old. Assume these differences are normally distributed. Use this 90% CI to test the claim that the mean of the differences is more than 20 beats per minute.

Data Set 2

Example

Data Set 2 below contains the heart rates of subjects before and after walking briskly for 10 minutes. Create a 90% CI for the mean of the differences, where the difference is the new heart rate minus the old. Assume these differences are normally distributed. Use this 90% CI to test the claim that the mean of the differences is more than 20 beats per minute.

Data Set 2

The 90% CI is (23.4,24.8). We can, thus reject the null hypothesis and support our claim. What's the significance of this test?

§7.1 Inference for the Mean of a Population

Example

Data Set 2 below contains the heart rates of subjects before and after walking briskly for 10 minutes. Create a 90% CI for the mean of the differences, where the difference is the new heart rate minus the old. Assume these differences are normally distributed. Use this 90% CI to test the claim that the mean of the differences is more than 20 beats per minute.

Data Set 2

The 90% CI is (23.4,24.8). We can, thus reject the null hypothesis and support our claim. What's the significance of this test? 0.05! Because it's a one-tailed test, and confidence intervals are inherently two-tailed.

§7.1 Inference for the Mean of a Population

Here are a few notes about the t-test procedure.

§7.1 Inference for the Mean of a Population

Here are a few notes about the t-test procedure.

- The t procedure depends on how normally distributed the data is.

§7.1 Inference for the Mean of a Population

Here are a few notes about the t-test procedure.

- The t procedure depends on how normally distributed the data is.
- Large samples are always fine ($n \geq 40$), but when sample sizes are small, we need to be more careful. A sample size between 15 and 39 is safe as long as there isn't strong skewness/outliers. A sample less than 15 must have nearly-normal data to use the t procedure.

Here are a few notes about the *t*-test procedure.

- The *t* procedure depends on how normally distributed the data is.
- Large samples are always fine ($n \geq 40$), but when sample sizes are small, we need to be more careful. A sample size between 15 and 39 is safe as long as there isn't strong skewness/outliers. A sample less than 15 must have nearly-normal data to use the *t* procedure.
- Be wary of discarding outliers!

What might we do when we can't directly use a t procedure?

§7.1 Inference for the Mean of a Population

What might we do when we can't directly use a t procedure? You might try the **sign test**, which is the simplest **distribution-free** test, i.e. it does not rely on any assumptions about the underlying distribution of the data.

What might we do when we can't directly use a t procedure? You might try the **sign test**, which is the simplest **distribution-free** test, i.e. it does not rely on any assumptions about the underlying distribution of the data.

Definition (sign test for matched pairs)

Ignore pairs with difference 0; the number of trails n is the count of the remaining pairs. The test statistic is the count X of pairs with positive differences. P-values for X are based on the binomial $B(n, 0.5)$ distribution.

§7.1 Inference for the Mean of a Population

What might we do when we can't directly use a t procedure? You might try the **sign test**, which is the simplest **distribution-free** test, i.e. it does not rely on any assumptions about the underlying distribution of the data.

Definition (sign test for matched pairs)

Ignore pairs with difference 0; the number of trails n is the count of the remaining pairs. The test statistic is the count X of pairs with positive differences. P-values for X are based on the binomial $B(n, 0.5)$ distribution.

This test can be performed in Minitab.