Hypothesis Tests for One-Sample Means

October 2, 2019

Decision Errors

- It is entirely possible that we make the right conclusion based on our data... but the wrong conclusion based on the true (unknown) parameter!
- In a criminal court, sometimes people are wrongly convicted. Other times, guilty people are not convicted at all.
- Unlike in the courts, statistics gives us the tools to quantify how often we make these sorts of errors.

Decision Errors

- There are two competing hypotheses: null and alternative.
- In a hypothesis test, we make some statement about which might be true.
- There are four possible scenarios. We can
    1. Reject H_0 when H_0 is false.
    2. Fail to reject H_0 when H_0 is true.
    3. Reject H_0 when H_0 is true (error).
    4. Fail to reject H_0 when H_0 is false (error).

Decision Errors

		Test Conclusion	
		Do not reject H_0	Reject H_0
Truth	H_0 true	Correct Decision	**Type I Error**
	H_0 false	**Type II Error**	Correct Decision

- A **Type 1 Error** is rejecting H_0 when it is actually true.
- A **Type 2 Error** is failing to reject H_0 when the H_A is actually true.

Example

Let's think about criminal courts. The null hypothesis is innocence.

- A Type I error is when we decide that a person is guilty, even though they are innocent.
- A Type II error is when we decide that we do not have enough evidence to say that someone is guilty, but they are in fact guilty.

Significance Levels

- The significance level, α, indicates how often the data will lead us to incorrectly reject H_0
- This is how often we commit a Type I error!
- In fact, α is the probability of committing such an error

$$\alpha = P(\text{Type I error})$$

Significance Levels

If we use a 95% confidence interval for hypothesis testing and the null is true,

- The significance level is $\alpha = 0.05$.
- We make an error whenever the point estimate is at least 1.96 standard errors away from the population parameter.
- This happens about 5% of the time

Hypothesis Testing For One-Sample Means

We will start with the situation wherein we know that $X \sim N(\mu, \sigma)$ and the value of σ is known.

Confidence Interval for μ

This $(1-\alpha)100\%$ confidence interval for μ is

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

where $\sigma/\sqrt{n}$ is the SE and $z_{\alpha/2}$ is again the critical value.

Example

The following $n = 5$ observations are from a $N(\mu, 2)$ distribution. Find a 90% confidence interval for μ.

$$1.1, \quad 0.5, \quad 2, \quad 1.9, \quad 2.7$$

Example

Recall that when we say "90% confident", we mean:

- If we draw repeated samples of size 5 from this distribution, then 90% of the time the corresponding intervals will contain the true value of μ.

Confidence Interval for μ

- In practice, we typically do not know the population standard deviation σ.
- Instead, we have to estimate this quantity.
- We will use the sample statistic s to estimate σ.
- This strategy works quite well when $n \geq 30$

This works quite well because we expect large samples to give us precise estimates such that

$$SE = \frac{\sigma}{\sqrt{n}} \approx \frac{s}{\sqrt{n}}.$$

Confidence Interval for μ

When $n \geq 30$ and σ is unknown, a $(1-\alpha)100\%$ confidence interval for μ is

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

where we've plugged in s for σ.

Example

The average heart rate of a random sample of 60 students is found to be 74 with a standard deviation of 11. Find a 95% confidence interval for the true mean heart rate of the students.

Hypothesis Testing for a Population Mean

We begin with the setting where $n \geq 30$.

- It is certainly possible to use the confidence interval to complete a hypothesis test.
- However, we also want to be able to use the test statistic and p-value approaches.

Hypothesis Testing for a Population Mean

For $n \geq 30$, the test statistic is

$$ts = z = \frac{\bar{x} - \mu_0}{s/\sqrt{n}}$$

where again $s/\sqrt{n} \approx \sigma/\sqrt{n}$ because we are using a large sample.

Hypothesis Testing for a Population Mean

There are five steps to carrying out these hypothesis tests:

1. Write out the null and alternative hypotheses.
2. Calculate the test statistic.
3. Use the significance level to find the critical value

   OR

   use the test statistic to find the p-value.
4. Compare the critical value to the test statistic

   OR

   compare the p-value to α.
5. Conclusion.

Example

In its native habitat, the average density of giant hogweed is 5 plants per m^2. In an invaded area, a sample of 50 plants produced an average of 11.17 plants per m^2 with a standard deviation of 8.9. Does the invaded area have a different average density than the native area? Test at the 5% level of significance.

Hypothesis Testing for a Population Mean

We now move to the situation where $n < 30$.

If $n < 30$ but we are dealing with a normal distribution and σ is known,

$$ts = z = \frac{\bar{x} - \mu_0}{\sigma/\sqrt{n}}$$

but we know that this will rarely (if ever) occur in practice!

Introducing the t-Distribution

- With a small sample size, plugging in s for σ can result in some problems.
- Therefore less precise samples will require us to make some changes.
- This brings us to the t-distribution.

Introducing the t-Distribution

The ***t*-distribution** is a symmetric, bell-shaped curve like the normal distribution.

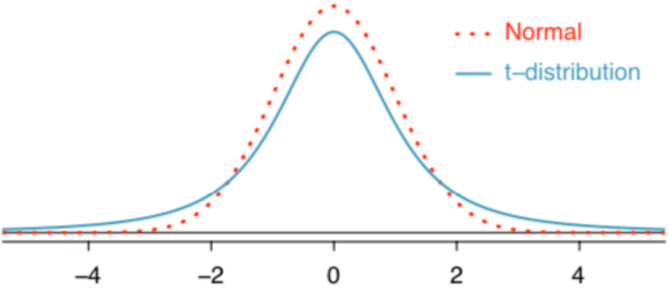

However, the t-distribution has more area in the tails.

The t-Distribution

The t-distribution:

- Is always centered at zero.
- Has one parameter: degrees of freedom (df).
- For our purposes,

$$df = n - 1$$

  where n is our sample size.

The t-Distribution

- The parameter df controls how fat the tails are.
- Higher values of df result in thinner tails.
  - I.e., larger sample sizes make the t-distribution look more normal.
- When $n \geq 30$, the t-distribution will be essentially equivalent to the normal distribution.
  - In practice, we often use t-tests even when $n \geq 30$.

Confidence Intervals for A Single Population Mean

When $n < 30$ and σ is unknown, we use the t-distribution for our confidence intervals. A $(1-\alpha)100\%$ confidence interval for μ is

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

Critical Values for the t-Distribution

Let's take a minute to look at the table of t-distribution critical values that we will use.

Test Statistics

The test statistic for the setting where $n < 30$ and σ is unknown is

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

(two-sided hypotheses)

P-Values

The p-value for two-sided hypotheses is then

$$2 \times P(t_{df} < -|ts|)$$

Example

The following data is on red blood cell counts (in 10^6 cells per microliter) for 9 people:

$$5.4, \quad 5.3, \quad 5.3, \quad 5.2, \quad 5.4, \quad 4.9, \quad 5.0, \quad 5.2, \quad 5.4$$

Test at the 5% level of significance if the average cell count is 5.