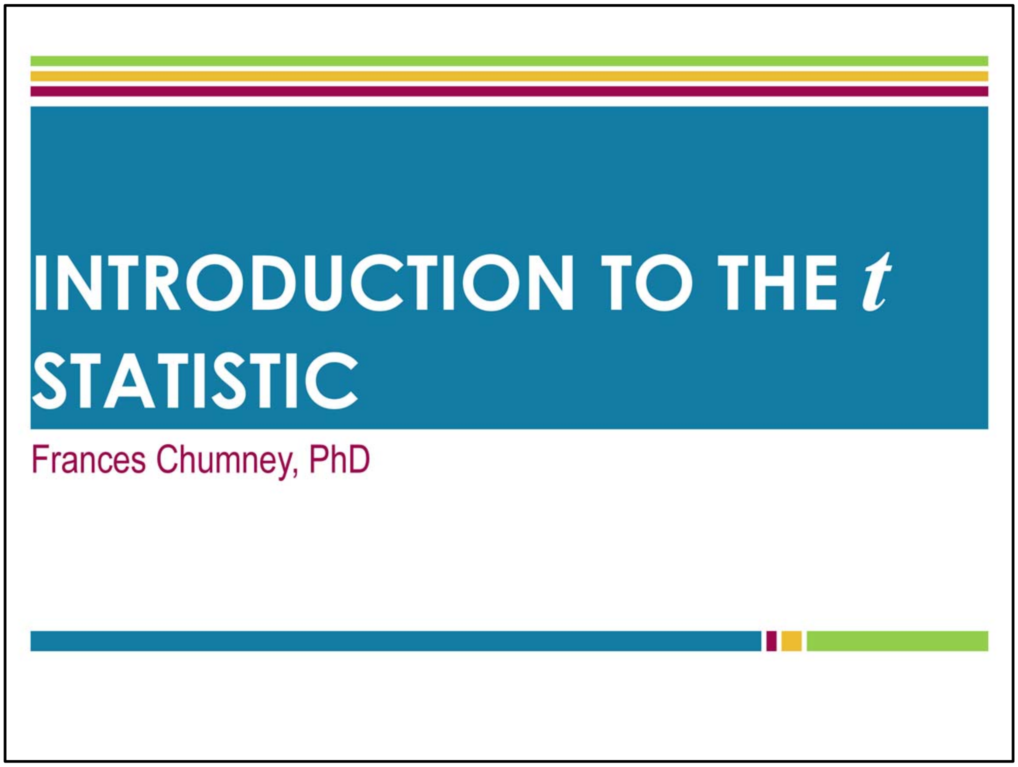

TRANSCRIPT:

This is Dr. Chumney with an introduction to the *t* statistic.

z-SCORES REVISITED

- Basic Assumptions
  - M approximates μ
  - σ_M estimates how well M approximates μ
  - Quantifies inferences made about the population
- The Problem with z-scores
  - Requires knowledge of σ

TRANSCRIPT:

Before we talk about the t statistic, it is helpful to have a quick reminder about z scores, and their assumptions in particular. First of all, the sample mean is assumed to approximate the population mean.

Second, the standard error of the mean estimates how well the sample mean approximates the population mean, and tells us how much difference is reasonable to expect between the sample mean and the population mean.

The third assumption is that the z-statistic quantifies inferences about the population that we make from the sample.

The primary problem with *z*-scores is that the computation of z-scores requires knowledge of the population standard deviation. This is a problem because most of the time we do not know enough about the population we are sampling from to know what its standard deviation is. So, this is a circular problem: we need the population values to compute z-statistics, but the purpose of a z-statistic is to make inferences about a population from a sample of it.

ESTIMATE OF STANDARD ERROR

estimate of the standard error when the population standard deviation is unknown; estimate of standard difference between M and μ

$$s_M = \sqrt{\frac{s^2}{n}}$$

TRANSCRIPT:

Because we do not know the population standard deviation, we cannot calculate the standard error of the mean. Instead, we estimate the standard error of the mean. The estimate of the standard error is an estimate of what the error is when the population standard deviation is not known. It is an **estimate** of the standard difference between the sample mean and the population mean.

The formula, shown here, is the square root of the sample variance divided by the sample size.

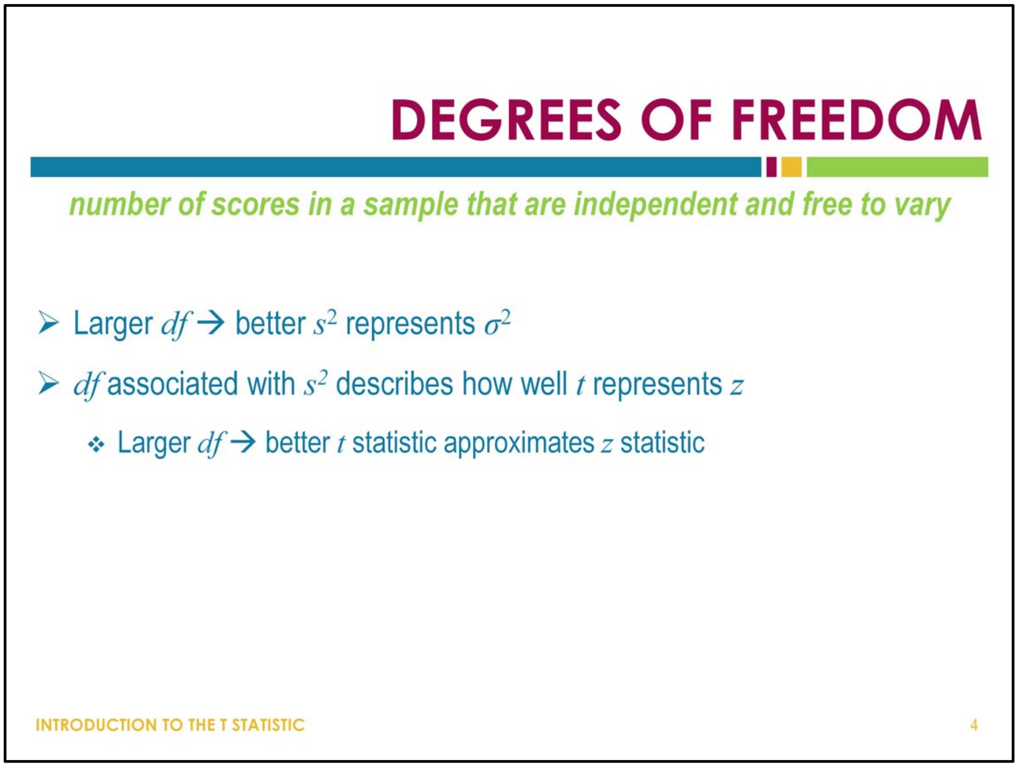

TRANSCRIPT:

The more degrees of freedom we have to work with, the better job the sample variance does of representing the population variance because bigger samples are better representatives of the population in general.

t STATISTIC

statistic used to test hypotheses about unknown population μ when the value of σ is unknown

- May be positive or negative (absolute value = magnitude)
- No upper- or lower- bound values

$$z = \frac{M-\mu}{\sigma_M} = \frac{M-\mu}{\sqrt{\frac{\sigma^2}{n}}} \qquad t = \frac{M-\mu}{s_M} = \frac{M-\mu}{\sqrt{\frac{s^2}{n}}}$$

TRANSCRIPT:

The t statistic is a statistic we use to test hypotheses about unknown population means when the value of the population standard deviation is also unknown.

The formula for the t statistic has the same structure as the formula for the z-statistic, except that it uses the estimated standard error of the mean in the denominator.

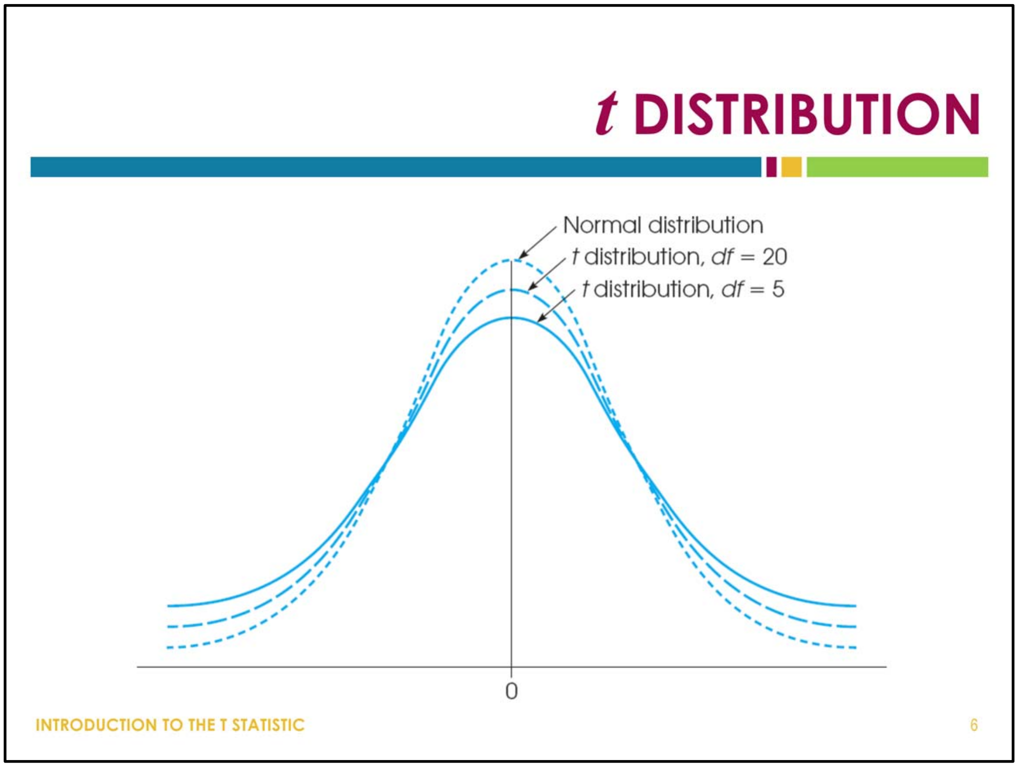

TRANSCRIPT:

The t-distribution approximates a normal curve in the same way that the t-statistic approximates a z-statistic.

How well the t-distribution approximates a normal distribution is determined by *df*. The larger the sample size, the larger the *df,* and the larger the *df*, the better job the t-distribution does of approximating a normal distribution.

So, as *df* increases, the t distribution looks more and more “normal” in shape.

t DISTRIBUTION: PROPORTIONS

	Proportion in One Tail					
	0.25	0.10	0.05	0.025	0.01	0.005
	Proportion in Two Tails Combined					
df	0.50	0.20	0.10	0.05	0.02	0.01
1	1.000	3.078	6.314	12.706	31.821	63.657
2	0.816	1.886	2.920	4.303	6.965	9.925
3	0.765	1.638	2.353	3.182	4.541	5.841
4	0.741	1.533	2.132	2.776	3.747	4.604
5	0.727	1.476	2.015	2.571	3.365	4.032
6	0.718	1.440	1.943	2.447	3.143	3.707

TRANSCRIPT:

With the z-statistic, we used the Unit Normal Table. With the t-statistic, we will use the t-distribution table.

In the t-distribution table, we identify the p value in the top 2 rows, look for the correct number of *df* in the first column, and then look for the critical value of t – the value of t that serves as the boundary for the critical region.

For example, if *df* = 3, 5% of the t-distribution is located in the tail beyond a value of *t* = 2.353.

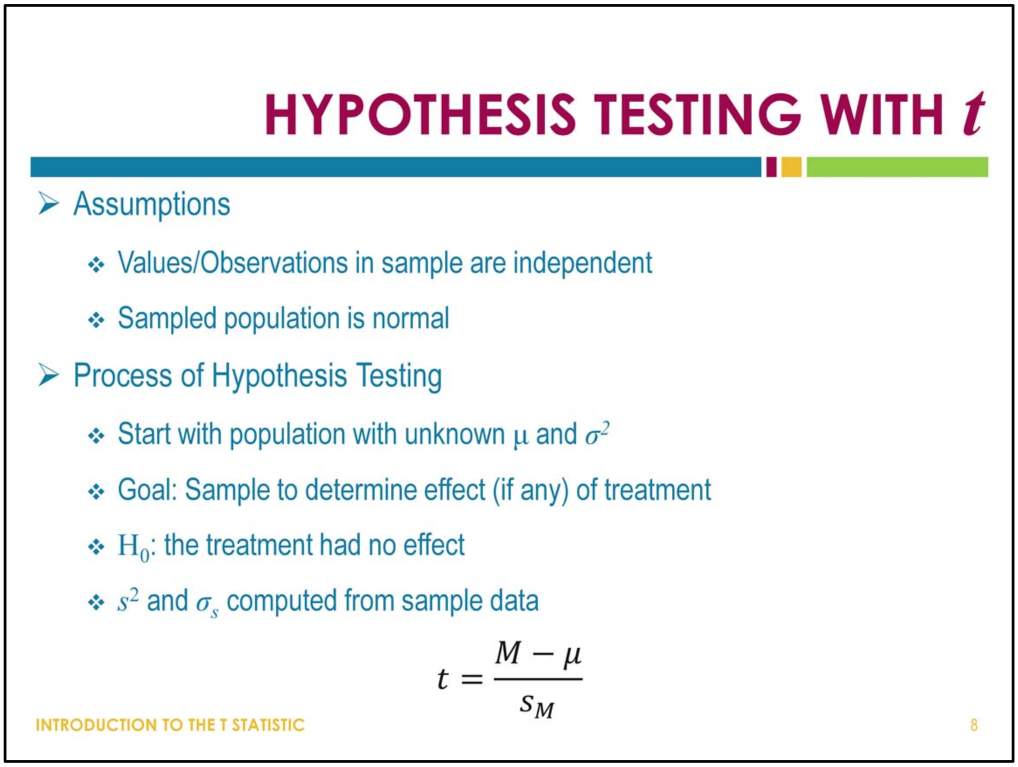

TRANSCRIPT:

There are a few assumptions associated with the t-statistic. First, it is assumed that all observations are independent. Second, the t-statistic assumes that the data come from a population for which the data would be normally distributed if data were obtained from all members of the population.

We do hypothesis tests with the t-statistic the same as with the z-statistic. To test hypotheses with the t-statistic, we start with the population for which we do not know the mean and variance. Usually, this population is a treatment group of some sort, from which we have data on a sample.

The goal is to use a sample of the treated population to determine if the treatment had any effect on the population.

The null hypothesis for testing with the t-statistic is the same as with the z-statistic. The null hypothesis is always that the treatment had no effect; or, that the population mean is unchanged. Just like in unit 2, the null hypothesis provides specific values for the population mean.

Different from z-statistic hypothesis tests, is that t-statistic tests use the sample data to provide a value for the sample mean, and the variance and estimated standard error are computed from the sample data instead of from the population parameters.

So, from just looking at the formula for the t-statistic, we can make a few inferences about the relationship between the difference in the means (the numerator) and the estimated standard error of the mean. For instance, when the difference in means is much larger than the estimated standard error, we get a larger t value (positive or negative).

If the difference between the sample and population means is large relative to the estimated standard error, we are more likely to reject the null hypothesis because the differences indicates the sample mean is very different from the population mean.

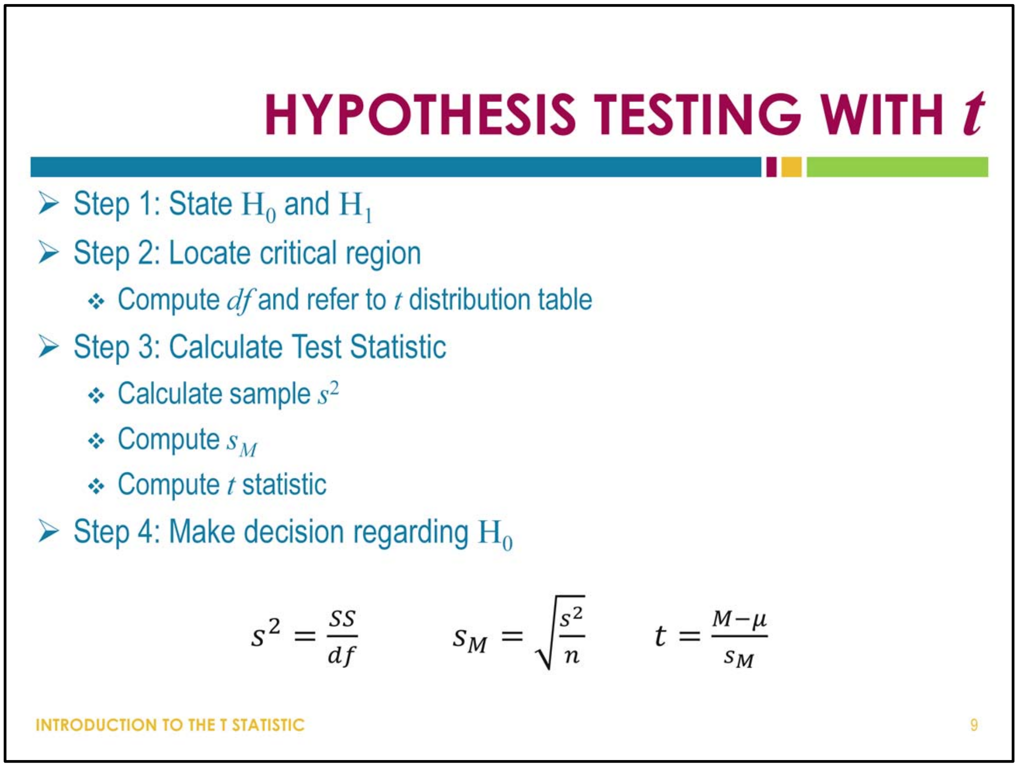

TRANSCRIPT:

We do hypothesis tests with the t-statistic the same as we did in chapter 8 with the z-statistic. To test hypotheses with the t-statistic, we start with the population for which we do not know the mean and variance. Usually, this population is a treatment group of some sort, from which we have data on a sample.

The goal is to use a sample of the treated population to determine if the treatment had any effect on the population.

The null hypothesis for testing with the t-statistic is the same as with the z-statistic. The null hypothesis is always that the treatment had no effect; or, that the population mean is unchanged. Just like in unit 2, the null hypothesis provides specific values for the population mean.

Different from z-statistic hypothesis tests, is that t-statistic tests use the sample data to provide a value for the sample mean, and the variance and estimated standard error are computed from the sample data instead of from the population parameters.

So, from just looking at the formula for the t-statistic, we can make a few inferences about the relationship between the difference in the means (the numerator) and the estimated standard error of the mean. For instance, when the difference in means is much larger than the estimated standard error, we get a larger t value (positive or negative).

If the difference between the sample and population means is large relative to the estimated standard error, we are more likely to reject the null hypothesis because the differences indicates the sample mean is very different from the population mean.

EFFECT SIZES

- Cohen's d

$$d = \frac{M - \mu}{s}$$

  - d = 0.2 – Small Effect (mean difference ≈ 0.2 standard deviation)
  - d = 0.5 – Medium Effect (mean difference ≈ 0.5 standard deviation)
  - d = 0.8 – Large Effect (mean difference ≈ 0.8 standard deviation)

- r^2

$$r^2 = \frac{t^2}{t^2 + df}$$

  - r^2 = 0.01 – Small Effect
  - r^2 = 0.09 – Medium Effect
  - r^2 = 0.25 – Large Effect

INTRODUCTION TO THE T STATISTIC 10

TRANSCRIPT:

In addition to an estimation of Cohen's d, r-squared can also be used as an effect size for t-tests. R2 determines how much of the variability in scores is explained by the treatment effect. Because the treatment causes scores to increase/decrease, the treatment is causing the varying scores.

If we identify how much of that variation is explained by the treatment, we then have a measure of the size of the treatment effect.

The amount of variability accounted for is usually reported as a proportion or percentage of the total variability.

INFLUENCE OF n AND s ON s_M

- Larger s_M → smaller values of t (closer to zero)
- Any factor that increases s_M will reduce likelihood of rejecting H_0
- Large Variance
  - Difference less likely to be significant
  - Scores widely scattered, so harder to see consistent patterns in data
  - Reduced effect size
- Large Sample
  - Difference more likely to be significant
  - smaller s_M

INTRODUCTION TO THE T STATISTIC 11

TRANSCRIPT:

The estimate of the standard error has a lot to do with the magnitude of a *t* statistic, and whether or not we decide to reject a null hypothesis. Understanding how the sample size and standard deviation can impact the *t* statistic is important because it helps to identify potential explanations for surprising results.

Larger values of the standard error typically result in *t* statistics closer to zero. The larger the variance of a sample, the less likely the t statistic will be significant, and the smaller the effect size will be.

Finally, with a large sample, the standard error is typically going to be smaller, which means the statistic is more likely to be significant.