Math 3081 (Probability and Statistics)

Lecture #22 of 27 $\sim$ August 11th, 2021

---

t Distributions and Confidence Intervals

- t Distributions
- Confidence Intervals with t Statistics

This material represents §4.2.1-4.2.2 from the course notes, and problems 16-20 from WeBWorK 6.

This also represents the last of the material for Midterm 3.

Overview of §5

We now move into the fifth and final chapter of the course, which deals with some additional types of hypothesis tests based on the normal distribution.

- In contrast with the z-tests, which require knowing the population standard deviation, our goal in this chapter is to construct hypothesis tests for normally-distributed quantities whose parameters are unknown.
- First we will discuss t tests and confidence intervals, for the mean of approximately normally distributed variables with an unknown standard deviation.
- Then we will discuss χ^2 tests and confidence intervals, for the variance of a sum of approximately normally distributed variables, which can also be used to to assess independence and goodness of fit.

t Distributions, I

In our discussion of hypothesis testing so far, we have relied on z tests, which require an approximately normally distributed test statistic whose standard deviation is known.

- However, in most situations, it is unlikely that we would actually know the population standard deviation.
- [Insert your own example of some quantity you'd want to estimate here, and then explain why you probably don't know the population standard deviation.]
- Instead, in such cases, we must estimate the population standard deviation from the sample standard deviation.

We already discussed the problem of estimating the population standard deviation from a sample back in our discussion of estimators.

t Distributions, II

Suppose values $x_1, \ldots, x_n$ are drawn from a normal distribution with unknown mean μ and unknown standard deviation σ.

- Let $\overline{x} = \frac{1}{n}(x_1 + \cdots + x_n)$ be the sample mean.
- We showed that the maximum likelihood estimate $\hat{\sigma} = \sqrt{\frac{1}{n}\left[(x_1 - \overline{x})^2 + (x_2 - \overline{x})^2 + \cdots + (x_n - \overline{x})^2\right]}$ for the standard deviation is biased.
- Instead of $\hat{\sigma}$, we use the sample standard deviation $S = \sqrt{\frac{1}{n-1}\left[(x_1 - \overline{x})^2 + (x_2 - \overline{x})^2 + \cdots + (x_n - \overline{x})^2\right]}$, whose square S^2 is an unbiased estimator of the population variance σ^2.

t Distributions, III

It may seem reasonable to say that if we use the estimated standard deviation S in place of the unknown population σ, then we should be able to use a z test with the resulting approximation.

- However, this turns out not to be the case!
- This might be surprising, because in other situations, such as the normal approximation to the binomial distribution, we have been able to adapt z tests with estimated standard deviations.
- However (if you recall) when I gave those explanations, I included a careful analysis of how far off the standard deviation estimate was, and showed it introduced a very small error.
- As we will see, that is not what happens here!

t Distributions, IV

To make things more explicit, we convert the discussion to a distribution with a single unknown parameter.

- We do this by looking at the normalized ratio $\dfrac{\overline{x} - \mu}{S/\sqrt{n}}$, which has mean 0 and standard deviation 1.
- This ratio is analogous to the z-score $\dfrac{\overline{x} - \mu}{\sigma/\sqrt{n}}$, whose distribution (under the assumptions of the null hypothesis that the true mean is μ) is the standard normal distribution of mean 0 and standard deviation 1.
- If we take $\dfrac{\overline{x} - \mu}{S/\sqrt{n}}$ as our test statistic, then (as we will show) this test statistic is not normally distributed!
- The distribution is similar in shape to the normal distribution, but it is in fact different, and is called the t distribution.

t Distributions, V

We can illustrate visually the lack of normality of the normalized test statistic $\dfrac{\overline{x}-\mu}{S/\sqrt{n}}$ by simulating a sampling procedure.

- Explicitly, suppose that X is normally distributed with mean $\mu = 0$ and standard deviation $\sigma = 1$, and we want to test the hypothesis that the mean actually is equal to 0 using the normalized test statistic $\dfrac{\overline{x}-\mu}{S/\sqrt{n}}$ with $n = 3$.
- To understand the behavior of $\dfrac{\overline{x}-\mu}{S/\sqrt{n}}$, we sample the standard normal distribution to obtain 3 data points x_1, x_2, x_3 and then compute $\dfrac{\overline{x}-\mu}{S/\sqrt{n}}$ using the sample mean $\overline{x}$ and estimated standard deviation S.

t Distributions, VI

Here is a histogram showing a total of 10000 samples, along with the pdf of the standard normal distribution:

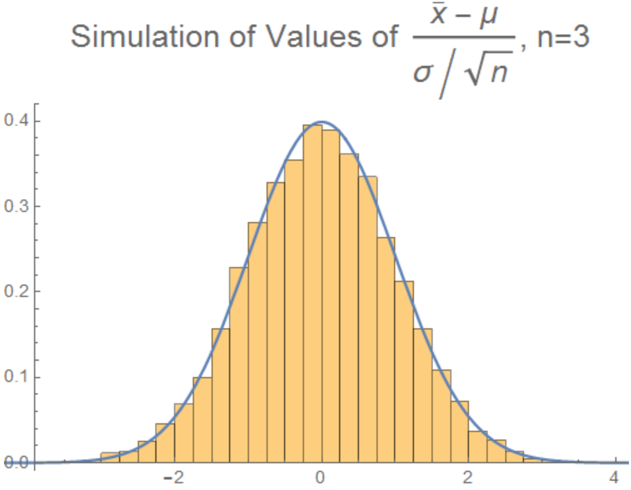

t Distributions, VII

Here is a histogram showing a total of 10000 samples, along with the pdf of the standard normal distribution:

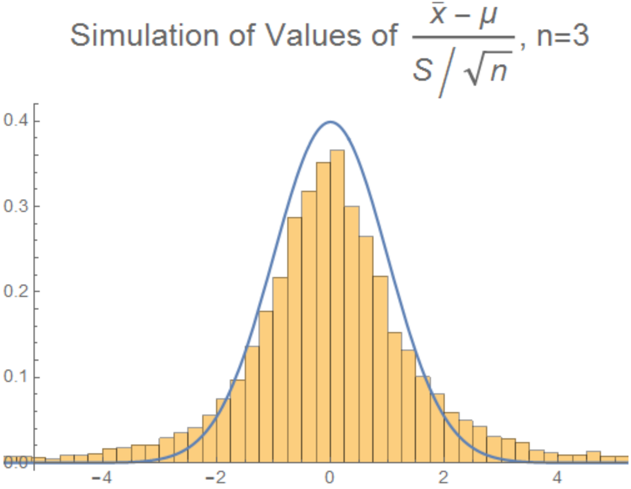

t Distributions, VIII

Compare the results of those two histograms:

- The first one simulates the test statistic $\dfrac{\overline{x} - \mu}{\sigma/\sqrt{n}}$.
- It matches the normal distribution very closely, which it should, since $\overline{x}$ actually is normally distributed with mean μ and standard deviation $\sigma/\sqrt{n}$!
- The second one simulates the test statistic $\dfrac{\overline{x} - \mu}{S/\sqrt{n}}$.
- It differs quite a bit from the normal distribution: there are values occurring in the tails of the distribution far more often than they do for the normal distribution, while the values near the center occur slightly less often than predicted.

t Distributions, X

Here are some more plots of simulations ($n = 5$):

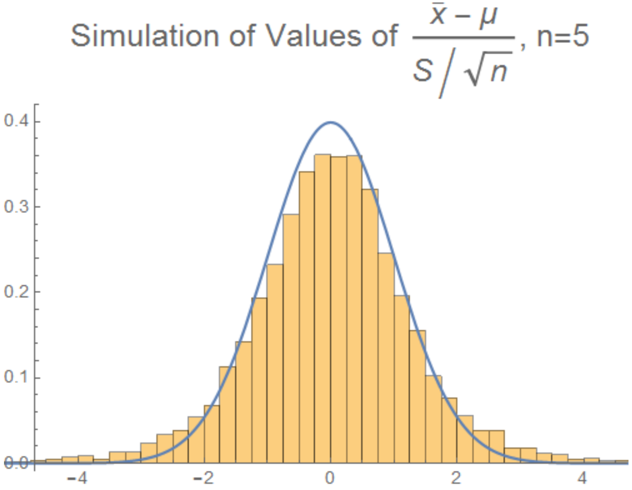

t Distributions, XI

Here are some more plots of simulations ($n = 10$):

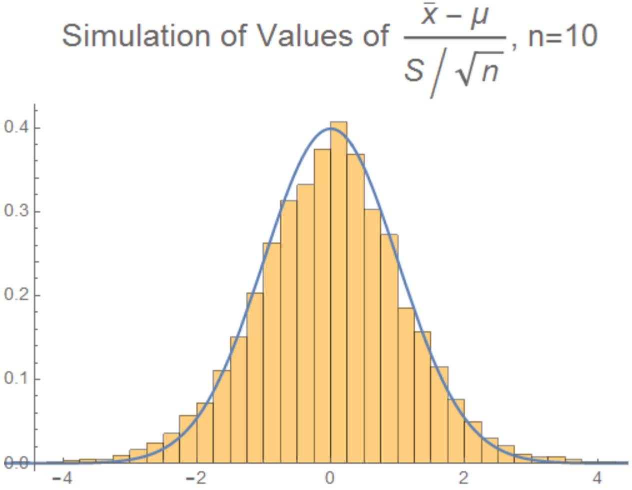

t Distributions, XII

Here are some more plots of simulations ($n = 20$):

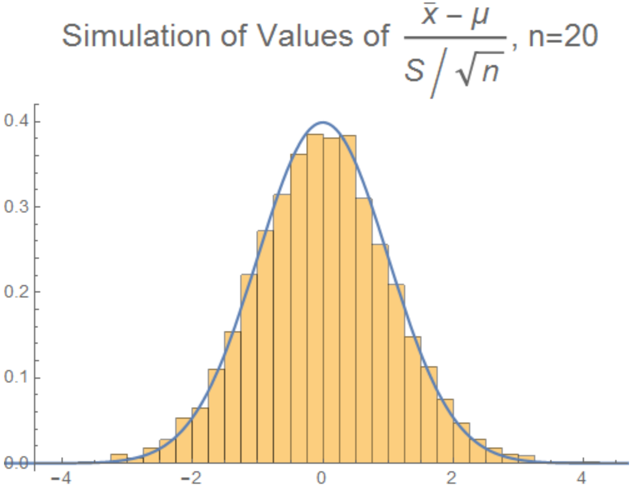

t Distributions, XIII

Here are some more plots of simulations ($n = 100$):

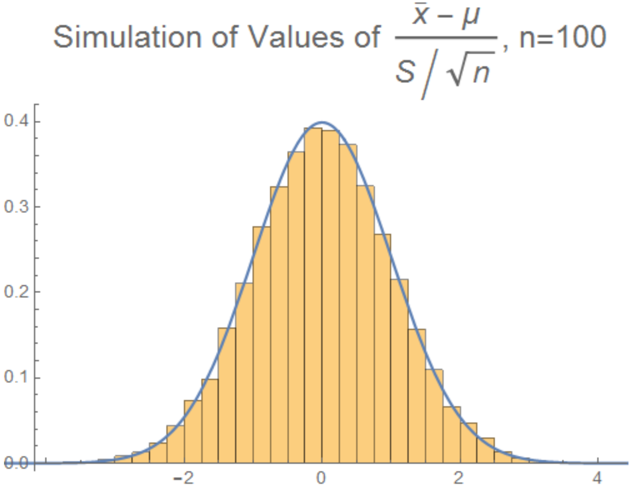

The Gamma Function, I

I will give the definition of the correct model, called the t-distribution, in a moment, but to do so we first need a few facts about the gamma function:

Definition

If z is a positive real number[a], the <u>gamma function</u> $\Gamma(z)$ is defined to be the value of the improper integral $\Gamma(z) = \int_0^\infty x^{z-1} e^{-x}\,dx$.

[a] In fact, this definition also makes perfectly good sense if z is a complex number whose real part is positive (which is why I used the letter z here).

- The gamma function arises naturally in complex analysis, number theory, and combinatorics, in addition to our use here in statistics.
- By integrating by parts, one may see that $\Gamma(z+1) = z\Gamma(z)$ for all z. Combined with the easy observation that $\Gamma(1) = 1$, we can see that $\Gamma(n) = (n-1)!$ for all positive integers n.

The Gamma Function, II

In addition to the values $\Gamma(n) = (n-1)!$, the values of the gamma function at half-integers can also be computed explicitly.

- To compute $\Gamma(1/2)$, we may substitute $u = \sqrt{x}$ to see $\Gamma(1/2) = 2\int_0^\infty e^{-u^2}\,du = \sqrt{\pi}$, as we calculated before when analyzing the normal distribution.
- Then, by using the identity $\Gamma(z+1) = z\Gamma(z)$, we can calculate $\Gamma(n+\frac{1}{2}) = (n-\frac{1}{2})(n-\frac{3}{2})\cdot\cdots\cdot\frac{1}{2}\sqrt{\pi} = \dfrac{(2n)!}{2^{2n}n!}\sqrt{\pi}$.
- If you like, you can confuse your friends by telling them that $\frac{1}{2}! = \sqrt{\pi}/2$, which also follows from this calculation.

t Distributions, I

With Γ taken care of, here is the official definition of the t distribution:

Definition

The t distribution with k degrees of freedom is the continuous random variable T_k whose probability density function

$$p_{T_k}(x) = \frac{\Gamma(\frac{k+1}{2})}{\sqrt{k\pi}\Gamma(\frac{k}{2})} \cdot (1 + x^2/k)^{-(k+1)/2}$$ *for all real numbers x.*

We will show in a moment that the t distribution with $n-1$ degrees of freedom is the proper model for the test statistic $\dfrac{\overline{x}-\mu}{S/\sqrt{n}}$.

t Distributions, II

Some history:

- The t distribution was first derived in 1876 by Helmert and Lüroth, and then appeared in several other papers.
- It is often referred to as Student's t distribution, because an analysis was published under the pseudonym "Student" by William Gosset in 1908, who because of his work at Guinness did not publish the results under his own name.
- The standard version of the story holds that Guinness wanted all its staff to publish using pseudonyms to protect its brewing methods and related data, since a paper had been previously published by one of its statisticians that inadvertently revealed some of its trade secrets.

t Distributions, III

<u>Examples</u>:

- The t distribution with 1 degree of freedom has probability density function $p_{T_1}(x) = \dfrac{1}{\pi(1+x^2)}$, which is the Cauchy distribution.
- The t distribution with 2 degrees of freedom has probability density function $p_{T_2}(x) = \dfrac{1}{(2+x^2)^{3/2}}$.
- The t distribution with 3 degrees of freedom has probability density function $p_{T_3}(x) = \dfrac{6\sqrt{3}}{\pi(3+x^2)^2}$.

t Distributions, IV

Here are a few basic properties of the t distributions (remember that the pdf is $p_{T_k}(x) = \dfrac{\Gamma(\frac{k+1}{2})}{\sqrt{k\pi}\Gamma(\frac{k}{2})} \cdot (1 + x^2/k)^{-(k+1)/2}$):

- Since $p_{T_k}(-x) = p_{T_k}(x)$, the pdf is symmetric about 0 (just like the normal distribution).
- Per the symmetry about 0, we would typically expect that the expected value of the distribution would be 0. This is true when $k \geq 2$, but in fact the expected value is undefined when $k = 1$ (the integral is a non-convergent improper integral).
- It is more difficult to compute the variance, but by manipulating the integrals appropriately, one can eventually show that the variance is undefined for $k = 1$ (expected value is undefined), ∞ for $k = 2$, and $\dfrac{k}{k-2}$ for $k > 2$.

t Distributions, V

Another very important property:

Proposition (t Distributions and Normal Distributions)

As $k \to \infty$, the probability density function
$p_{T_k}(x) = \frac{\Gamma(\frac{k+1}{2})}{\sqrt{k\pi}\Gamma(\frac{k}{2})} \cdot (1 + x^2/k)^{-(k+1)/2}$ *for the t distribution*
approaches the standard normal distribution.

t Distributions, V

Another very important property:

Proposition (t Distributions and Normal Distributions)

As $k \to \infty$, the probability density function
$$p_{T_k}(x) = \frac{\Gamma(\frac{k+1}{2})}{\sqrt{k\pi}\,\Gamma(\frac{k}{2})} \cdot (1 + x^2/k)^{-(k+1)/2} \text{ for the } t \text{ distribution}$$
approaches the standard normal distribution.

Proof:

- Using the fact that $\lim_{k\to\infty}(1 + y/k)^k = e^y$, we can see that $\lim_{k\to\infty}(1 + x^2/k)^{-(k+1)/2} = e^{-x^2/2}$.
- Thus, $\displaystyle\lim_{k\to\infty} \frac{\Gamma(\frac{k+1}{2})}{\sqrt{k\pi}\,\Gamma(\frac{k}{2})} \cdot (1 + x^2/k)^{-(k+1)/2} = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$.
- The limit of the constant follows from Stirling's formula, our calculation of $\Gamma(k/2)$ five slides ago, or by observing that the limit function must be a pdf.

t Distributions, VI

Our main result is the following:

Theorem (Modeling Property of the t Distribution)

Suppose $n \geq 2$ and that $X_1, X_2, \ldots, X_n$ are independent, identically normally distributed random variables with mean μ and standard deviation σ. If $\overline{X} = \frac{1}{n}(X_1 + \cdots + X_n)$ denotes the sample mean and $S = \sqrt{\frac{1}{n-1}\left[(X_1 - \overline{X})^2 + (X_2 - \overline{X})^2 + \cdots + (X_n - \overline{X})^2\right]}$ denotes the sample standard deviation, then the distribution of the normalized test statistic $\frac{\overline{X} - \mu}{S/\sqrt{n}}$ is the t distribution T_{n-1} with $n - 1$ degrees of freedom.

We will only outline the proof, since most of the actual calculations are rather technical and unenlightening.

t Distributions, VIII

Proof (outline):

- First, we show that the sample mean $\overline{X}$ and the sample standard deviation S are independent. This is relatively intuitive, but the proof requires the observation that orthogonal changes of variable preserve independence.
- Next, we compute the probability density functions for the numerator $\overline{X} - \mu$ (which is normal with mean 0 and standard deviation $\sigma/\sqrt{n}$) and the denominator.

t Distributions, IX

Proof (outline) (continued):

- We are finding the PDF of the denominator term $S/\sqrt{n}$.
- First, we compute the probability density of $\dfrac{n-1}{\sigma^2}S^2 = \dfrac{1}{\sigma^2}\left[(X_1 - \overline{X})^2 + (X_2 - \overline{X})^2 + \cdots + (X_n - \overline{X})^2\right] = (\frac{X_1-\mu}{\sigma})^2 + (\frac{X_2-\mu}{\sigma})^2 + \cdots + (\frac{X_{n-1}-\mu}{\sigma})^2$.
- This last expression is the sum of the squares of $n-1$ independent standard normal distributions, which is known as a χ^2 distribution (which we discuss in more detail next week).
- The pdf of the denominator $\dfrac{1}{S/\sqrt{n}}$ can then be computed using the pdf above, using standard techniques for computing the pdf of a function of a random variable.

t Distributions, X

Proof (outline) (continued) (continued):

- Now we have the PDFs of both $\overline{X} - \mu$ and $\dfrac{1}{S/\sqrt{n}}$.
- Because since $\overline{X} - \mu$ and $S/\sqrt{n}$ were shown to be independent, the joint pdf for $\overline{X} - \mu$ and $S/\sqrt{n}$ is simply the product of their individual pdfs.
- Then, at last, we can the probability density function for $\dfrac{\overline{X} - \mu}{S/\sqrt{n}} = (\overline{X} - \mu) \cdot \dfrac{1}{S/\sqrt{n}}$ can be calculated by evaluating an appropriate integral of the joint pdf of $\overline{X} - \mu$ and $S/\sqrt{n}$.

All that's left is to actually perform all the calculations!

t Distributions, XVIII

To illustrate, here's the normal distribution:

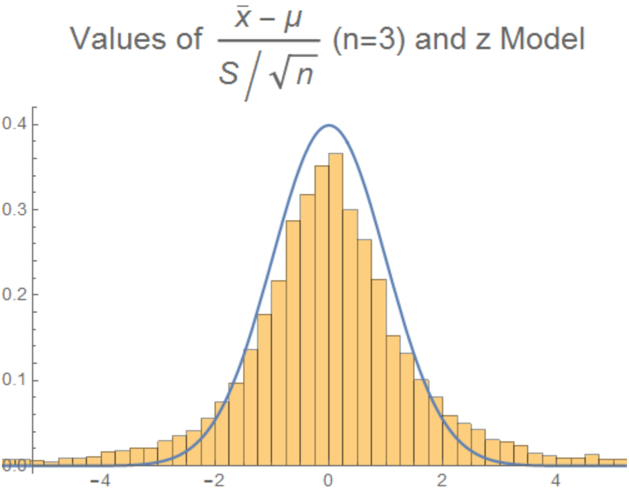

t Distributions, XIX

... and here's the t distribution (much better!):

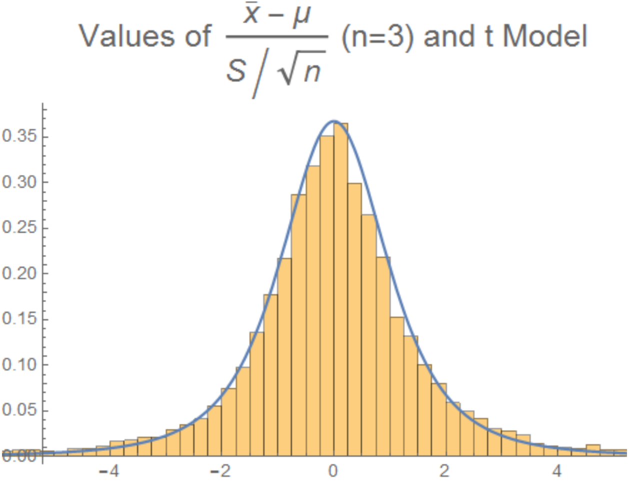

t Distributions, XVIII

Here's the sample for $n = 5$:

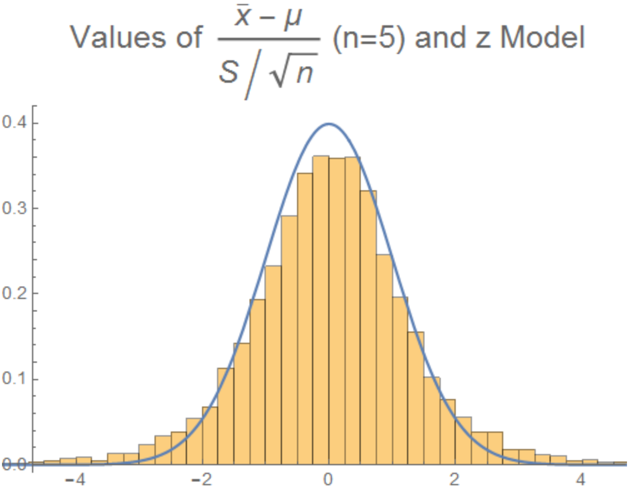

t Distributions, XIX

And the t distribution with the same dataset:

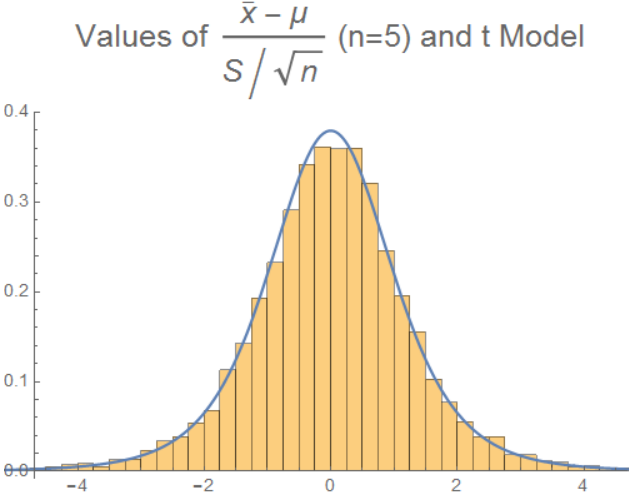

t Confidence Intervals, I

Before we discuss how to use the t distribution for hypothesis testing (next lecture), we will explain how to use t statistics to find confidence intervals (the rest of this lecture).

t Confidence Intervals, II

The idea is quite simple: if we want to find a confidence interval for the unknown mean of a normal distribution whose standard deviation is also unknown, we can estimate the mean using the t distribution.

- Specifically, since the normalized statistic $\dfrac{\overline{x} - \mu}{S/\sqrt{n}}$ is modeled by the t-distribution T_{n-1} with $n - 1$ degrees of freedom, we can compute a $100(1 - \alpha)\%$ confidence interval using a t-statistic in place of the z-statistic that we used for normally distributed random variables whose standard deviation was known.

t Confidence Intervals, III

Here is the picture for the normal distribution:

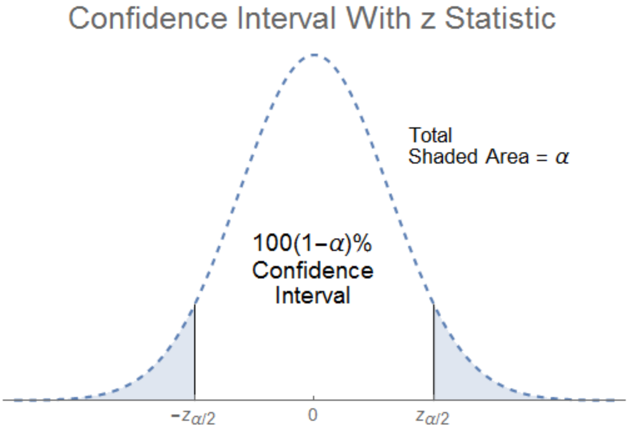

Note $z_{\alpha/2}$ (called c last week) has $P(N_{0,1} \geq z_{\alpha/2}) = \alpha/2$.

t Confidence Intervals, IV

Here is the picture for the t distribution:

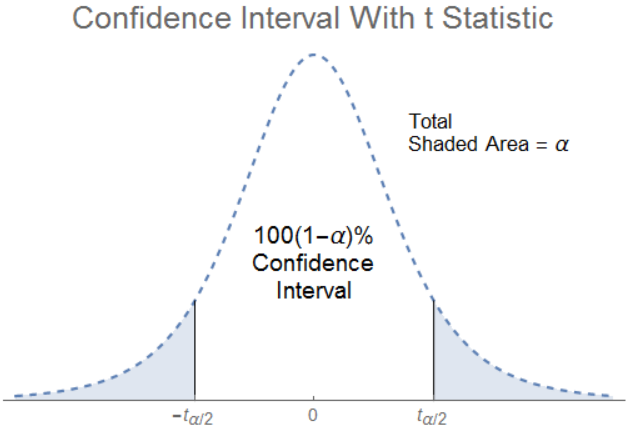

Here $t_{\alpha/2,df}$ has $P(T_{n-1} \geq t_{\alpha/2,df}) = \alpha/2$.

t Confidence Intervals, V

So what we want is to find the value $t_{\alpha/2,df}$ that plays the role of $z_{\alpha/2}$ for the t distribution.

- You can think of this value as the number of standard deviations in the margin of error for the confidence interval.
- If we find $t_{\alpha/2,df}$ such that $P(-t_{\alpha/2,df} < T_{n-1} < t_{\alpha/2,df}) = 1 - \alpha$, then this yields the $100(1-\alpha)\%$ confidence interval $\hat{\mu} \pm t_{\alpha/2,df}\frac{S}{\sqrt{n}} = (\hat{\mu} - t_{\alpha/2,df}\frac{S}{\sqrt{n}}, \hat{\mu} + t_{\alpha/2,df}\frac{S}{\sqrt{n}})$.
- Using the symmetry of the t distribution, $P(-t_{\alpha/2,df} < T_{n-1} < t_{\alpha/2,df}) = 1 - \alpha$ is equivalent to $P(T_{n-1} < -t_{\alpha/2,df}) = \alpha/2$, or also to $P(t_{\alpha/2,df} < T_{n-1}) = 1 - (\alpha/2)$, which allows us to compute the value of $t_{\alpha/2,df}$ by evaluating the inverse cumulative distribution function for T_{n-1}.

t Confidence Intervals, VI

We can summarize this discussion as follows:

Proposition

A $100(1-\alpha)\%$ *confidence interval for the unknown mean* μ *of a normal distribution with unknown standard deviation is given by* $\hat{\mu} \pm t_{\alpha/2,df} \frac{S}{\sqrt{n}} = (\hat{\mu} - t_{\alpha/2,df} \frac{S}{\sqrt{n}}, \hat{\mu} + t_{\alpha/2,df} \frac{S}{\sqrt{n}})$ *where* n *sample points* $x_1, \ldots, x_n$ *are taken from the distribution,* $\hat{\mu} = \frac{1}{n}(x_1 + \cdots + x_n)$ *is the sample mean,* $S = \sqrt{\frac{1}{n-1}[(x_1-\hat{\mu})^2 + \cdots + (x_n-\hat{\mu})^2]}$ *is the sample standard deviation, and* $t_{\alpha/2,df}$ *is the constant satisfying* $P(-t_{\alpha/2,df} < T_{n-1} < t_{\alpha/2,df}) = 1-\alpha$*.*

Note that the number of degrees of freedom is $df = n-1$, not n.

t Confidence Intervals, VII

Some specific values of $t_{\alpha/2,df}$ for various common values of n and α are given in the table below (note that the last row for $n = \infty$ represents the entry for the normal distribution):

Entries give $t_{\alpha/2,df}$ such that $P(-t_{\alpha/2,df} < T_{df} < t_{\alpha/2,df}) = 1 - \alpha$								
df	50%	80%	90%	95%	98%	99%	99.5%	99.9%
1	1	3.0777	6.3138	12.706	31.820	63.657	127.32	636.62
2	0.8165	1.8856	2.9200	4.3027	6.9646	9.9248	14.089	31.599
3	0.7649	1.6477	2.3534	3.1824	4.5407	5.8409	7.4533	12.924
4	0.7407	1.5332	2.1318	2.7764	3.7469	4.6041	5.5976	8.6103
5	0.7267	1.4759	2.0150	2.5706	3.3649	4.0321	4.7733	6.8688
10	0.6998	1.3722	1.8125	2.2281	2.7638	3.1693	3.5814	4.5869
20	0.6870	1.3253	1.7247	2.0860	2.5280	2.8453	3.1534	3.8495
50	0.6794	1.2987	1.6759	2.0086	2.4033	2.6778	2.9370	3.4960
100	0.6770	1.2901	1.6602	1.9840	2.3642	2.6259	2.8707	3.3905
∞	0.6745	1.2816	1.6449	1.9600	2.3263	2.5758	2.8070	3.2905

t Confidence Intervals, VIII

<u>Example</u>: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

1. A 90% confidence interval for the average score on the exam.
2. A 95% confidence interval for the average score on the exam.
3. A 99.5% confidence interval for the average score on the exam.
4. The 90%, 95%, and 99.5% confidence intervals for the average score if the population standard deviation were known to be 9.1 points.
5. Compare the z and t confidence intervals calculated above.

t Confidence Intervals, VIII

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

1. A 90% confidence interval for the average score on the exam.
2. A 95% confidence interval for the average score on the exam.
3. A 99.5% confidence interval for the average score on the exam.
4. The 90%, 95%, and 99.5% confidence intervals for the average score if the population standard deviation were known to be 9.1 points.
5. Compare the z and t confidence intervals calculated above.

- We need to look up the appropriate entries from the t table, with $df = 21 - 1 = 20$ (20 degrees of freedom).

t Confidence Intervals, IX

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

1. A 90% confidence interval for the average score on the exam.

t Confidence Intervals, IX

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

1. A 90% confidence interval for the average score on the exam.

- We have $\hat{\mu} = 78.2$, $S = 9.1$, and $n = 15$.
- The confidence interval is given by $\hat{\mu} \pm t_{\alpha/2,df}(S/\sqrt{n})$.
- The entry in the table for $t_{\alpha/2,df}$ is 1.7247.
- Thus, the 90% confidence interval is $\hat{\mu} \pm 1.7247 \cdot S/\sqrt{n} = (74.77, 81.62)$.

t Confidence Intervals, X

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

2. A 95% confidence interval for the average score on the exam.

t Confidence Intervals, X

<u>Example</u>: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

2. A 95% confidence interval for the average score on the exam.

- We have $\hat{\mu} = 78.2$, $S = 9.1$, and $n = 15$.
- The confidence interval is given by $\hat{\mu} \pm t_{\alpha/2,df}(S/\sqrt{n})$.
- The entry in the table for $t_{\alpha/2,df}$ is 2.0860.
- Thus, the 95% confidence interval is $\hat{\mu} \pm 2.0860 \cdot S/\sqrt{n} = (74.06, 82.34)$.

t Confidence Intervals, XI

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

4. A 99.5% confidence interval for the average score on the exam.

t Confidence Intervals, XI

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

4. A 99.5% confidence interval for the average score on the exam.

- We have $\hat{\mu} = 78.2$, $S = 9.1$, and $n = 15$.
- The confidence interval is given by $\hat{\mu} \pm t_{\alpha/2,df}(S/\sqrt{n})$.
- The entry in the table for $t_{\alpha/2,df}$ is 3.1534.
- Thus, the 99.5% confidence interval is $\hat{\mu} \pm 3.1534 \cdot S/\sqrt{n} = (71.94, 84.46)$.

t Confidence Intervals, XII

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

5. The 90%, 95%, and 99.5% confidence intervals for the average score if the population standard deviation were known to be 9.1 points.

t Confidence Intervals, XII

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

5. The 90%, 95%, and 99.5% confidence intervals for the average score if the population standard deviation were known to be 9.1 points.

- We now have $\hat{\mu} = 78.2$, $\sigma = 9.1$, and $n = 15$.
- The confidence interval is given by $\hat{\mu} \pm z_{\alpha/2}(\sigma/\sqrt{n})$.
- We can find the entries for the z scores in the bottom row of the table, with $n = \infty$.
- The 90% confidence interval is $\hat{\mu} \pm 1.6449 \cdot \sigma/\sqrt{n} = (74.93, 81.46)$, the 95% confidence interval is $\hat{\mu} \pm 1.9600 \cdot \sigma/\sqrt{n} = (74.31, 82.09)$, and the 99.5% confidence interval is $\hat{\mu} \pm 2.8070 \cdot \sigma/\sqrt{n} = (72.63, 83.77)$.

t Confidence Intervals, XII

Example: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

6. Compare the z and t confidence intervals calculated above.

t Confidence Intervals, XII

<u>Example</u>: The exam scores in a statistics class are expected to be normally distributed. 21 students' scores are sampled, and the average score is 78.2 points with a sample standard deviation of 9.1 points. Find the following:

6. Compare the z and t confidence intervals calculated above.

Here's a table:

Confidence	t	z
90%	(74.77,81.62)	(74.93, 81.46)
95%	(74.06,82.34)	(74.31, 82.09)
99.5%	(71.94,84.46)	(72.63, 83.77)

Unsurprisingly, knowing the population standard deviation gives us narrower confidence intervals, but they're pretty close to the ones using the sample standard deviation: that's because 21 is a reasonably big sample.

t Confidence Intervals, XIII

Example: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

1. Find the sample mean and sample standard deviation.
2. Find 80%, 90%, 95%, and 99.9% confidence intervals for the true mean of the distribution.
3. Find confidence intervals for a normal distribution whose standard deviation is the same as this sample estimate.
4. Compare the two sets of confidence intervals.

t Confidence Intervals, XIV

<u>Example</u>: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

1. Find the sample mean and sample standard deviation.

t Confidence Intervals, XIV

<u>Example</u>: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

1. Find the sample mean and sample standard deviation.

- The sample mean is $\hat{\mu} = \frac{1}{5}(1.21 + 4.60 + 4.99 - 2.21 + 3.21) = 2.36$.
- The sample variance is $S^2 = \dfrac{1}{4}\Big[(1.21-2.36)^2+(4.60-2.36)^2 +(4.99-2.36)^2+(-2.21-2.36)^2+(3.21-2.36)^2\Big] = 8.7161$.
- The sample standard deviation is then $S = \sqrt{8.7161} = 2.9523$.

Most calculators (and basically all software) has functions that will compute these values for you. The mean is not so bad, but the standard deviation is rather annoying to evaluate by hand.

t Confidence Intervals, XV

Example: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

2. Find 80%, 90%, 95%, and 99.9% confidence intervals for the true mean of the distribution.

t Confidence Intervals, XV

Example: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

2. Find 80%, 90%, 95%, and 99.9% confidence intervals for the true mean of the distribution.

- We have $\hat{\mu} = 2.36$, $S = 2.9523$, $df = n - 1 = 4$.
- The confidence interval is $\hat{\mu} \pm t_{\alpha/2,df}(S/\sqrt{n})$.
- We just need to use the table / a calculator to get $t_{\alpha/2,df}$.
- The 80% CI is $\hat{\mu} \pm 1.5332 \cdot S/\sqrt{n} = (0.3357, 4.3843)$.
- The 90% CI is $\hat{\mu} \pm 2.1318 \cdot S/\sqrt{n} = (-0.4546, 5.1746)$.
- The 95% CI is $\hat{\mu} \pm 2.7764 \cdot S/\sqrt{n} = (-1.3057, 6.0257)$.
- The 99.9% CI is $\hat{\mu} \pm 8.6103 \cdot S/\sqrt{n} = (-9.0083, 13.7283)$.

t Confidence Intervals, XVI

Example: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

3. Find confidence intervals for a normal distribution whose standard deviation is the same as this sample estimate.

t Confidence Intervals, XVI

Example: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, −2.21, and 3.21.

3. Find confidence intervals for a normal distribution whose standard deviation is the same as this sample estimate.

- Now $\hat{\mu} = 2.36$ and $\sigma = 2.9523$, with $n = \infty$.
- The confidence interval is $\hat{\mu} \pm z_{\alpha/2}(\sigma/\sqrt{n})$.
- The 80% CI is $\hat{\mu} \pm 1.2816 \cdot \sigma/\sqrt{n} = (0.6679, 4.0521)$.
- The 90% CI is $\hat{\mu} \pm 1.6449 \cdot \sigma/\sqrt{n} = (0.1882, 4.5118)$.
- The 95% CI is $\hat{\mu} \pm 1.9600 \cdot \sigma/\sqrt{n} = (-0.2278, 4.9478)$.
- The 99.9% CI is $\hat{\mu} \pm 3.2905 \cdot \sigma/\sqrt{n} = (-1.9845, 6.7045)$.

t Confidence Intervals, XVII

Example: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

4. Compare the two sets of confidence intervals.

t Confidence Intervals, XVII

Example: A normal distribution with unknown mean and standard deviation is sampled five times, yielding the values 1.21, 4.60, 4.99, -2.21, and 3.21.

4. Compare the two sets of confidence intervals.

Confidence	t	z
80%	(0.3357,4.3843)	(0.6679,4.0521)
90%	(-0.4546,5.1746)	(0.1882,4.5118)
95%	(-1.3057,6.0257)	(-0.2278,4.9478)
99.9%	(-9.0083,13.7283)	(-1.9845,6.7045)

- Note how much narrower the z CIs are!
- For example, if we erroneously quoted the 80% normal confidence interval, by using the cdf for the t distribution we can see that it is actually only a 64% confidence interval for the t statistic: quite a bit lower!

t Confidence Intervals, XVIII

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%.

1. Find the sample mean and sample standard deviation.
2. Find 50%, 80%, 90%, and 95% confidence intervals for the true reaction yield, under the assumption that the reaction yield is approximately normally distributed.

t Confidence Intervals, XVIII

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%.

1. Find the sample mean and sample standard deviation.
2. Find 50%, 80%, 90%, and 95% confidence intervals for the true reaction yield, under the assumption that the reaction yield is approximately normally distributed.

- Since the reaction yield is approximately normally distributed, but we do not know the standard deviation, it is appropriate to use the t distribution here.

t Confidence Intervals, XIX

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%.

1. Find the sample mean and sample standard deviation.

t Confidence Intervals, XIX

<u>Example</u>: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.3%, 52.6%, and 56.1%.

1. Find the sample mean and sample standard deviation.

- The sample average is $\hat{\mu} = \frac{1}{3}(41.3\% + 52.6\% + 56.1\%) = 50\%$.
- The sample standard deviation is $S = \sqrt{\frac{1}{2}\left[(41.3\% - 50\%)^2 + (52.6\% - 50\%)^2 + (56.1\% - 50\%)^2\right]} = 7.7350\%$.

t Confidence Intervals, XX

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.0%, 52.6%, and 56.1%.

2. Find 50%, 80%, 90%, and 95% confidence intervals for the true reaction yield, under the assumption that the reaction yield is approximately normally distributed.

t Confidence Intervals, XX

Example: To estimate the reaction yield, a new chemical synthesis is run three times, giving yields of 41.0%, 52.6%, and 56.1%.

2. Find 50%, 80%, 90%, and 95% confidence intervals for the true reaction yield, under the assumption that the reaction yield is approximately normally distributed.

- We have $\hat{\mu} = 50\%$, $S = 7.7350\%$, $n = 3$.
- The confidence interval is $\hat{\mu} \pm t_{\alpha/2,df}(S/\sqrt{n})$.
- We just need to use the table / a calculator to get $t_{\alpha/2,df}$.
- The 50% CI is $\hat{\mu} \pm 0.8165 \cdot S/\sqrt{n} = (46.35\%, 53.65\%)$.
- The 80% CI is $\hat{\mu} \pm 1.8856 \cdot S/\sqrt{n} = (41.58\%, 58.42\%)$.
- The 90% CI is $\hat{\mu} \pm 2.9200 \cdot S/\sqrt{n} = (36.96\%, 63.04\%)$.
- The 95% CI is $\hat{\mu} \pm 4.3027 \cdot S/\sqrt{n} = (30.79\%, 69.21\%)$.

t Confidence Intervals, XXI

Example: Your intrepid Math 3081 instructor has written five chapters of notes, totaling 25pp, 36pp, 18pp, 21pp, and 31pp. Assume that the number of pages per chapter of notes is approximately normally distributed.

1. Find the sample mean and sample standard deviation.
2. Find 50%, 80%, 90%, and 99% confidence intervals for the average number of pages per chapter.
3. Given that I have written 56 chapters' worth of notes for my courses over the last eight years, find a point estimate for the total number of pages in these chapters, as well as 50%, 80%, 90%, and 99% confidence intervals.

t Confidence Intervals, XXII

Example: Your intrepid Math 3081 instructor has written five chapters of notes, totaling 25pp, 36pp, 18pp, 21pp, and 31pp. Assume that the number of pages per chapter of notes is approximately normally distributed.

1. Find the sample mean and sample standard deviation.

t Confidence Intervals, XXII

Example: Your intrepid Math 3081 instructor has written five chapters of notes, totaling 25pp, 36pp, 18pp, 21pp, and 31pp. Assume that the number of pages per chapter of notes is approximately normally distributed.

1. Find the sample mean and sample standard deviation.

- The sample mean is $\hat{\mu} = \frac{1}{5}(25 + 36 + 18 + 21 + 31) = 26.2$.
- The sample variance is $S^2 = \dfrac{1}{4}\Big[(25 - 26.2)^2 + (36 - 26.2)^2 + (18 - 26.2)^2 + (21 - 26.2)^2 + (31 - 26.2)^2\Big] = 53.7$, so the sample standard deviation is $S = \sqrt{53.7} = 7.3280$.

t Confidence Intervals, XXIII

Example: Your intrepid Math 3081 instructor has written five chapters of notes, totaling 25pp, 36pp, 18pp, 21pp, and 31pp. Assume that the number of pages per chapter of notes is approximately normally distributed.

2. Find 50%, 80%, 90%, and 99% confidence intervals for the average number of pages per chapter.

t Confidence Intervals, XXIII

Example: Your intrepid Math 3081 instructor has written five chapters of notes, totaling 25pp, 36pp, 18pp, 21pp, and 31pp. Assume that the number of pages per chapter of notes is approximately normally distributed.

2. Find 50%, 80%, 90%, and 99% confidence intervals for the average number of pages per chapter.

- We have $\hat{\mu} = 26.2$, $S = 7.3280$, and $df = n - 1 = 4$.
- We just need to use the table / a calculator to get $t_{\alpha/2,df}$.
- The 50% CI is $\hat{\mu} \pm 0.7407 \cdot S/\sqrt{n} = (23.8, 28.6)$.
- The 80% CI is $\hat{\mu} \pm 1.5332 \cdot S/\sqrt{n} = (21.2, 31.2)$.
- The 90% CI is $\hat{\mu} \pm 2.1318 \cdot S/\sqrt{n} = (19.2, 33.2)$.
- The 99% CI is $\hat{\mu} \pm 4.6041 \cdot S/\sqrt{n} = (11.1, 41.3)$.

t Confidence Intervals, XXIV

Example: Your intrepid Math 3081 instructor has written five chapters of notes, totaling 25pp, 36pp, 18pp, 21pp, and 31pp. Assume that the number of pages per chapter of notes is approximately normally distributed.

3. Given that I have written 56 chapters' worth of notes for my courses over the last eight years, find a point estimate for the total number of pages in these chapters, as well as 50%, 80%, 90%, and 99% confidence intervals.

t Confidence Intervals, XXIV

Example: Your intrepid Math 3081 instructor has written five chapters of notes, totaling 25pp, 36pp, 18pp, 21pp, and 31pp. Assume that the number of pages per chapter of notes is approximately normally distributed.

3. Given that I have written 56 chapters' worth of notes for my courses over the last eight years, find a point estimate for the total number of pages in these chapters, as well as 50%, 80%, 90%, and 99% confidence intervals.

- We just scale the confidence intervals we just calculated by 56.
- The 50% CI is $56(\hat{\mu} \pm 0.7407 \cdot S/\sqrt{n}) = (1331, 1603)$.
- The 80% CI is $56(\hat{\mu} \pm 1.5332 \cdot S/\sqrt{n}) = (1186, 1749)$.
- The 90% CI is $56(\hat{\mu} \pm 2.1318 \cdot S/\sqrt{n}) = (1076, 1858)$.
- The 95% CI is $56(\hat{\mu} \pm 4.6041 \cdot S/\sqrt{n}) = (622, 2312)$.

In case you were wondering, the actual number of pages is 1,239. (Do you believe this number? How would you test it?)

Summary

We introduced the t distributions and some of their properties.

We discussed how to construct confidence intervals using t statistics.

Next lecture: One-sample t tests, two-sample t tests.