Inference with One Sample Means and Small Samples

Dr. Ab Mosca (they/them)

Slides based off slides courtesy of OpenIntro and John McGreedy of Johns Hopkins University

Plan for Today

- Hypothesis testing with small samples

Friday the 13th

Between 1990 - 1992 researchers in the UK collected data on traffic flow, accidents, and hospital admissions on Friday 13th and the previous Friday, Friday 6th. Below is an excerpt from this data set on traffic flow. We can assume that traffic flow on given day at locations 1 and 2 are independent.

	type	date	6th	13th	diff	location
1	traffic	1990, July	139246	138548	698	loc 1
2	traffic	1990, July	134012	132908	1104	loc 2
3	traffic	1991, September	137055	136018	1037	loc 1
4	traffic	1991, September	133732	131843	1889	loc 2
5	traffic	1991, December	123552	121641	1911	loc 1
6	traffic	1991, December	121139	118723	2416	loc 2
7	traffic	1992, March	128293	125532	2761	loc 1
8	traffic	1992, March	124631	120249	4382	loc 2
9	traffic	1992, November	124609	122770	1839	loc 1
10	traffic	1992, November	117584	117263	321	loc 2

Friday the 13th

- We want to investigate if people's behavior is different on Friday the 13th compared to Friday 6th.

Friday the 13th

- We want to investigate if people's behavior is different on Friday the 13th compared to Friday 6th.
- One approach is to compare the traffic flow on these two days.

Friday the 13th

- We want to investigate if people's behavior is different on Friday the 13th compared to Friday 6th.
- One approach is to compare the traffic flow on these two days.

What is our research question?

What are H_o and H_a?

Friday the 13^{th}

- We want to investigate if people's behavior is different on Friday the 13^{th} compared to Friday 6^{th}.
- One approach is to compare the traffic flow on these two days.

What is our research question?

Does traffic flow change on Friday the 13^{th} compared to Friday the 6^{th}?

What are H_o and H_a?

Friday the 13^{th}

- We want to investigate if people's behavior is different on Friday the 13^{th} compared to Friday 6^{th}.
- One approach is to compare the traffic flow on these two days.
- H_0 : Average traffic flow on Friday 6^{th} and 13^{th} are equal.

  H_A : Average traffic flow on Friday 6^{th} and 13^{th} are different.

Friday the 13^{th}

- We want to investigate if people's behavior is different on Friday the 13^{th} compared to Friday 6^{th}.
- One approach is to compare the traffic flow on these two days.
- H_0 : Average traffic flow on Friday 6^{th} and 13^{th} are equal.

  H_A : Average traffic flow on Friday 6^{th} and 13^{th} are different.

Each row in the data set represents traffic flow recorded at one location in the same month of the same year: one count from Friday 6^{th} and the other Friday 13^{th}. Are these two counts independent?

	type	date	6^{th}	13^{th}	diff	location
1	traffic	1990, July	139246	138548	698	loc 1
2	traffic	1990, July	134012	132908	1104	loc 2
3	traffic	1991, September	137055	136018	1037	loc 1
4	traffic	1991, September	133732	131843	1889	loc 2
5	traffic	1991, December	123552	121641	1911	loc 1
6	traffic	1991, December	121139	118723	2416	loc 2

Friday the 13th

- We want to investigate if people's behavior is different on Friday the 13th compared to Friday 6th.
- One approach is to compare the traffic flow on these two days.
- H_0 : Average traffic flow on Friday 6th and 13th are equal.

  H_A : Average traffic flow on Friday 6th and 13th are different.

Each case in the data set represents traffic flow recorded at the same location in the same month of the same year: one count from Friday 6th and the other Friday 13th. Are these two counts independent?

No! Why not?

Friday the 13th

- We want to investigate if people's behavior is different on Friday the 13th compared to Friday 6th.
- One approach is to compare the traffic flow on these two days.
- H_0 : Average traffic flow on Friday 6th and 13th are equal.

  H_A : Average traffic flow on Friday 6th and 13th are different.

Each case in the data set represents traffic flow recorded at the same location in the same month of the same year: one count from Friday 6th and the other Friday 13th. Are these two counts independent?

No! If there's a lot of traffic on the 6th, I might take a different route on the 13th.

Hypotheses

What are the hypotheses (in symbols) for testing for a difference between the average traffic flow between Friday 6th and 13th?

A. $H_0 : \boldsymbol{\mu}_{6th} = \boldsymbol{\mu}_{13th}$

$H_A : \boldsymbol{\mu}_{6th} \neq \boldsymbol{\mu}_{13th}$

B. $H_0 : \boldsymbol{p}_{6th} = \boldsymbol{p}_{13th}$

$H_A : \boldsymbol{p}_{6th} \neq \boldsymbol{p}_{13th}$

C. $H_0 : \boldsymbol{\mu}_{diff} = 0$

$H_A : \boldsymbol{\mu}_{diff} \neq 0$

D. $H_0 : \bar{x}_{diff} = 0$

$H_A : \bar{x}_{diff} \neq 0$

H_0 : Average traffic flow on Friday 6th and 13th are equal.

H_A : Average traffic flow on Friday 6th and 13th are different.

Hypotheses

What are the hypotheses (in symbols) for testing for a difference between the average traffic flow between Friday 6th and 13th?

A. $H_0 : \mu_{6th} = \mu_{13th}$

$H_A : \mu_{6th} \neq \mu_{13th}$

B. $H_0 : p_{6th} = p_{13th}$

$H_A : p_{6th} \neq p_{13th}$

C. $H_0 : \mu_{diff} = 0$

$H_A : \mu_{diff} \neq 0$

D. $H_0 : \bar{x}_{diff} = 0$

$H_A : \bar{x}_{diff} \neq 0$

H_0 : Average traffic flow on Friday 6th and 13th are equal.

H_A : Average traffic flow on Friday 6th and 13th are different.

Conditions

- *Independence*
- Sample size / skew

Conditions

- *Independence*: We are told to assume that cases (rows) are independent
- Sample size / skew

Conditions

- *Independence*: We are told to assume that cases (rows) are independent
- Sample size / skew:

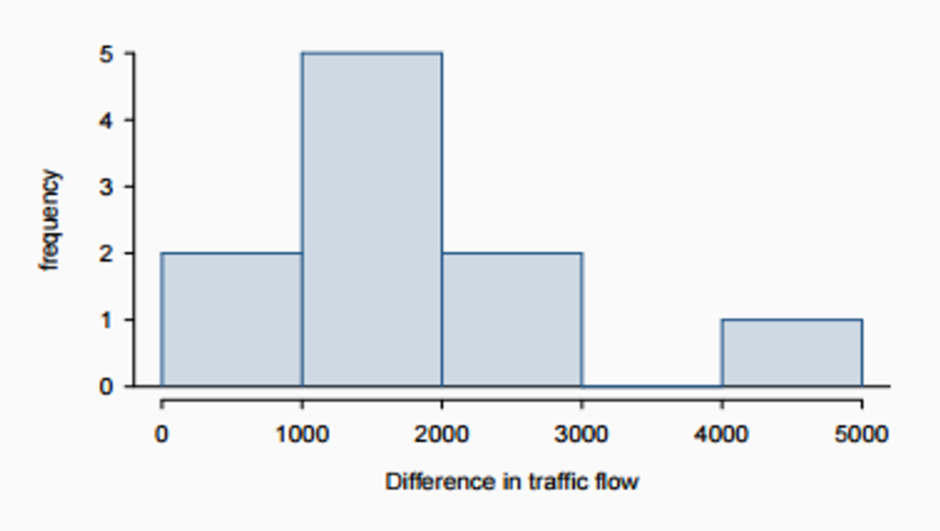

- The sample distribution does not appear to be extremely skewed, but it's very difficult to assess with such a small sample size. We might want to think about whether we would expect the population distribution to be skewed or not (probably not, it should be equally likely to have days with lower than average traffic and higher than average traffic)
- We do not know $\boldsymbol{\sigma}$ and *n* is too small to assume *s* is reliable estimate for $\boldsymbol{\sigma}$

Conditions

- *Independence*: We are told to assume that cases (rows) are independent
- Sample size / skew:

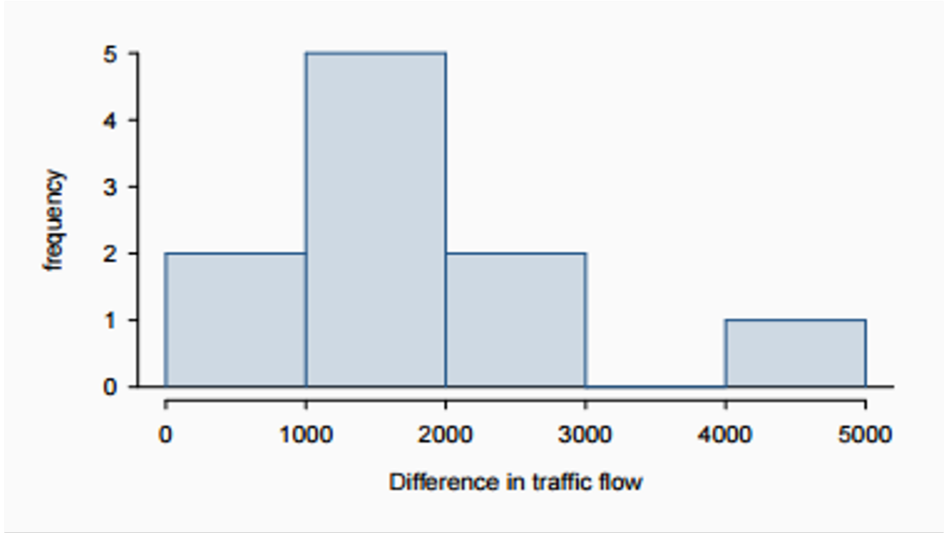

- The sample distribution does not appear to be extremely skewed, but it's very difficult to assess with such a small sample size. We might want to think about whether we would expect the population distribution to be skewed or not (probably not, it should be equally likely to have days with lower than average traffic and higher than average traffic)
- We do not know $\boldsymbol{\sigma}$ and n is too small to assume s is reliable estimate for $\boldsymbol{\sigma}$

What do we do when the sample size is small?

Review: what purpose does a large sample server?

As long as observations are independent, and the population distribution is not extremely skewed, a large sample would ensure that...

- the sampling distribution of the mean is nearly normal
- the estimate of the standard error, as $\frac{s}{\sqrt{n}}$, is reliable

The normality condition

- The CLT, which states that sampling distributions will be nearly normal, holds true for *any* sample size as long as the population distribution is nearly normal

The normality condition

- The CLT, which states that sampling distributions will be nearly normal, holds true for *any* sample size as long as the population distribution is nearly normal
- While this is a helpful special case, it's inherently difficult to verify normality in small data sets

The normality condition

- The CLT, which states that sampling distributions will be nearly normal, holds true for *any* sample size as long as the population distribution is nearly normal
- While this is a helpful special case, it’s inherently difficult to verify normality in small data sets
- We should exercise caution when verifying the normality condition for small samples. It is important to not only examine the data but also think about where the data come from
  - For example, ask: would I expect this distribution to be symmetric, and am I confident that outliers are rare?

The *t* distribution

- When the population standard deviation is unknown (almost always), the uncertainty of the standard error estimate is addressed by using a new distribution: the t distribution.

The *t* distribution

- When the population standard deviation is unknown (almost always), the uncertainty of the standard error estimate is addressed by using a new distribution: the *t* distribution.
- This distribution also has a bell shape, but its tails are thicker than the normal model's.

The *t* distribution

- When the population standard deviation is unknown (almost always), the uncertainty of the standard error estimate is addressed by using a new distribution: the *t* distribution.
- This distribution also has a bell shape, but its tails are thicker than the normal model's.
- Therefore observations are more likely to fall beyond two SDs from the mean than under the normal distribution

The *t* distribution

- When the population standard deviation is unknown (almost always), the uncertainty of the standard error estimate is addressed by using a new distribution: the ***t*** **distribution**.
- This distribution also has a bell shape, but its tails are thicker than the normal model's.
- Therefore observations are more likely to fall beyond two SDs from the mean than under the normal distribution
- These extra thick tails are helpful for resolving our problem with a less reliable estimate for the standard error (since *n* is small)

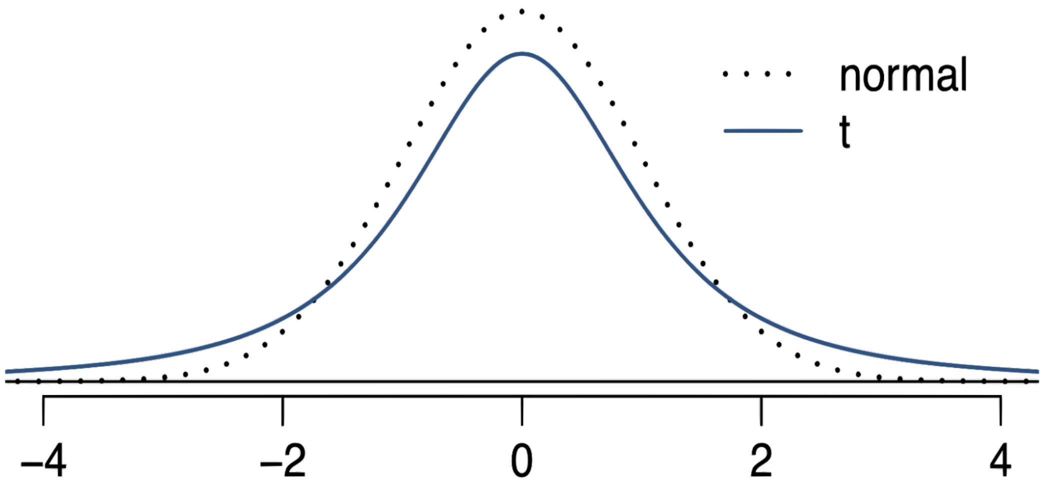

The *t* distribution (cont.)

- Always centered at zero, like the standard normal (z) distribution
- Has a single parameter: *degrees of freedom* (*df*).

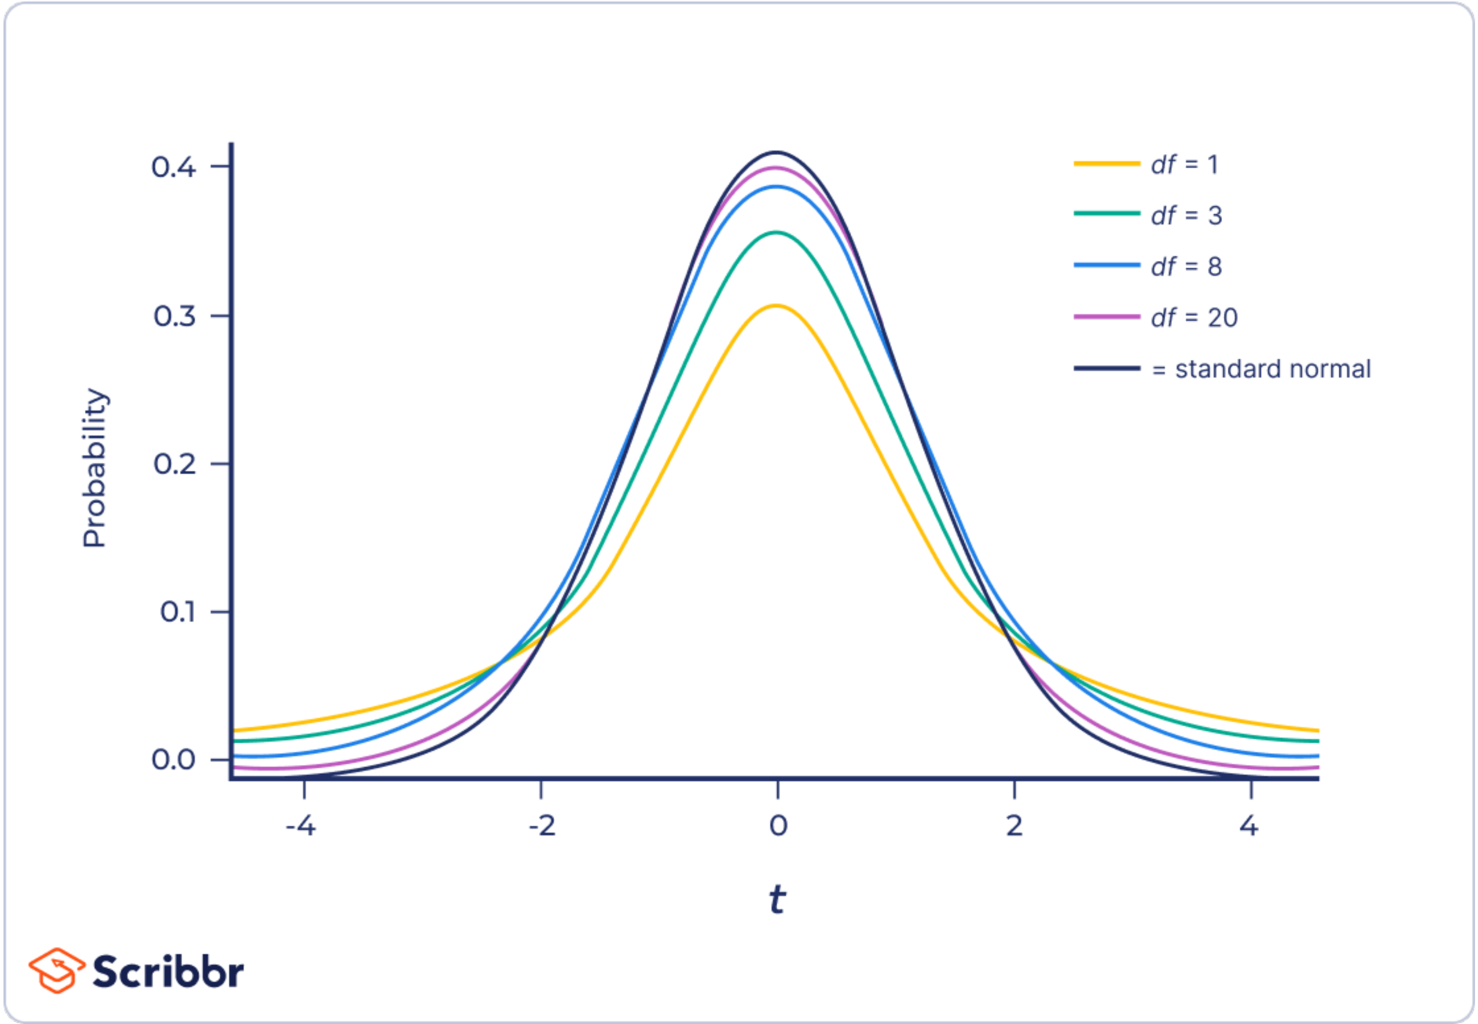

The *t* distribution (cont.)

- Always centered at zero, like the standard normal (z) distribution
- Has a single parameter: *degrees of freedom* (*df*).

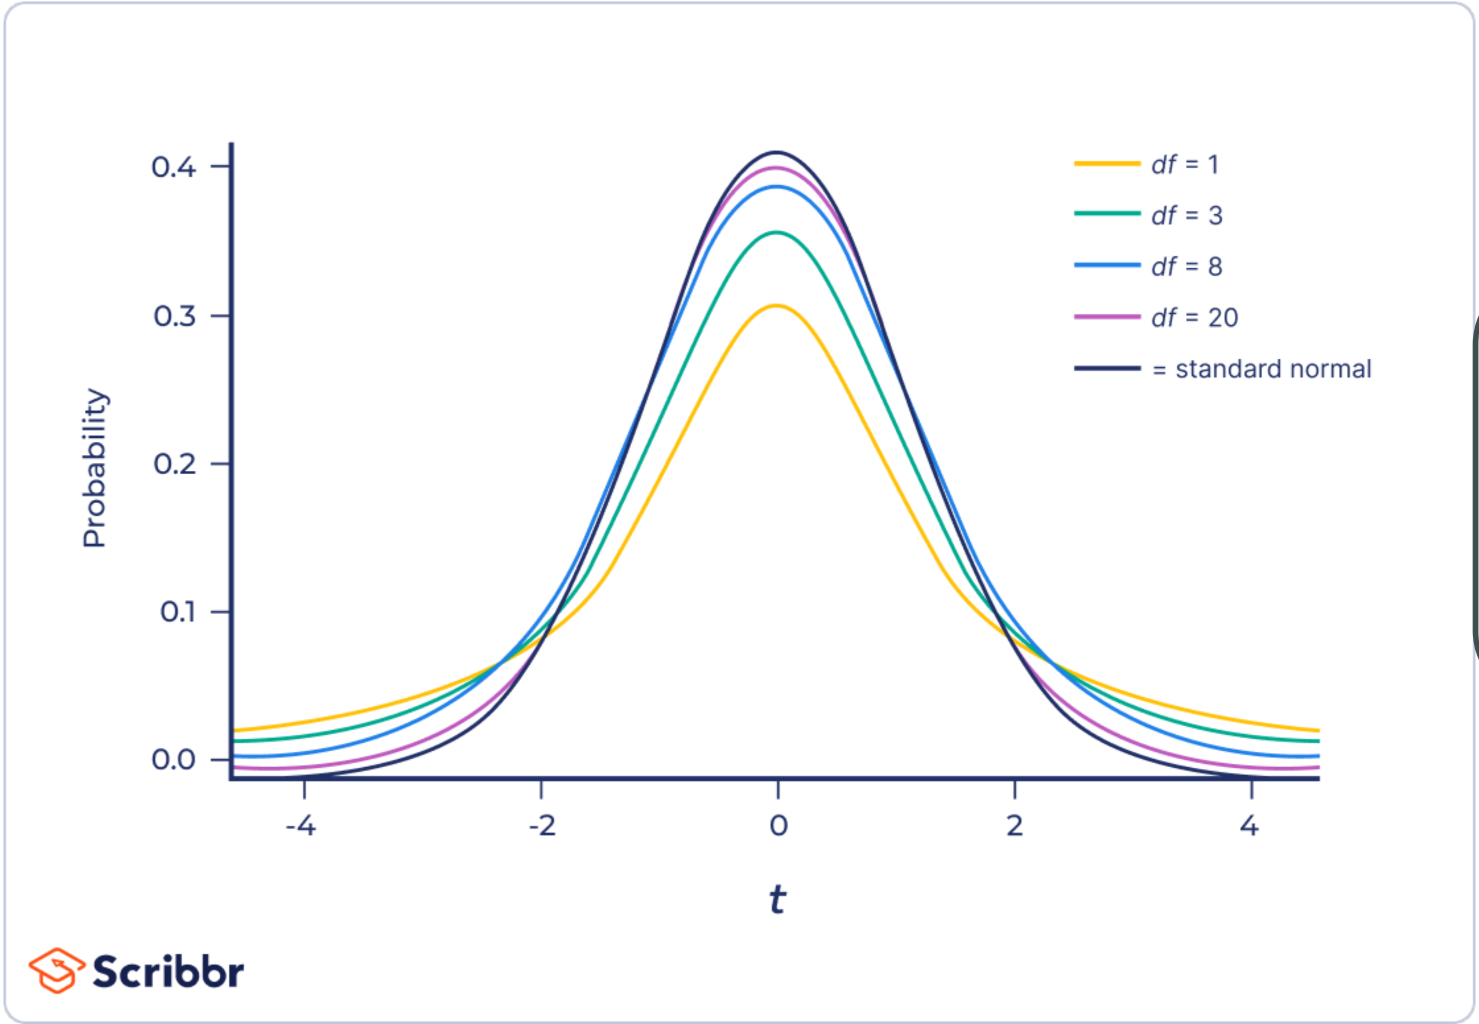

What happens to the shape of the *t* distribution as *df* increases?

The *t* distribution (cont.)

- Always centered at zero, like the standard normal (z) distribution
- Has a single parameter: *degrees of freedom* (*df*).

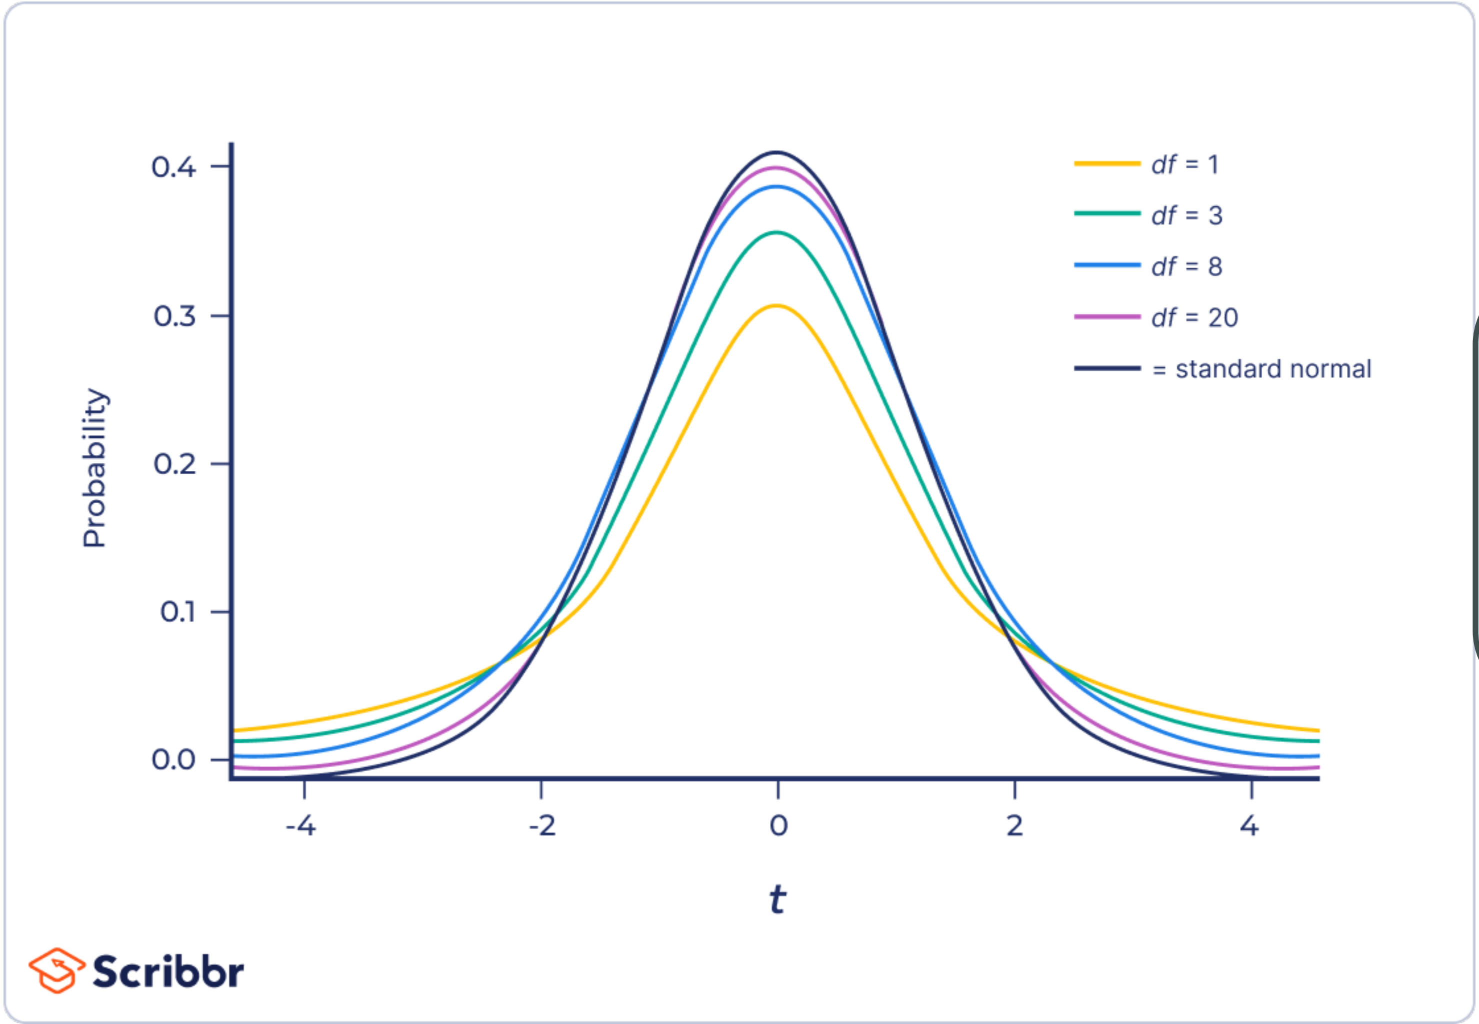

What happens to the shape of the ***t*** distribution as *df* increases?

Approaches normal

Back to Friday the 13th

	type	date	6th	13th	diff	location
1	traffic	1990, July	139246	138548	698	loc 1
2	traffic	1990, July	134012	132908	1104	loc 2
3	traffic	1991, September	137055	136018	1037	loc 1
4	traffic	1991, September	133732	131843	1889	loc 2
5	traffic	1991, December	123552	121641	1911	loc 1
6	traffic	1991, December	121139	118723	2416	loc 2
7	traffic	1992, March	128293	125532	2761	loc 1
8	traffic	1992, March	124631	120249	4382	loc 2
9	traffic	1992, November	124609	122770	1839	loc 1
10	traffic	1992, November	117584	117263	321	loc 2

$\bar{x}_{diff} = 1836$

$s_{diff} = 1176$

Find the test statistic

Test statistic for inference on a small sample mean

The test statistic for inference on a small sample ($n < 50$) mean is the T statistic with $df = n - 1$

$$T_{df} = \frac{\text{point estimate} - \text{null value}}{SE}$$

Note: Null value is 0 because in the null hypothesis we set $\mu_{diff} = 0$

Find the test statistic

Test statistic for inference on a small sample mean

The test statistic for inference on a small sample (n < 50) mean is the T statistic with $df = n - 1$

$$T_{df} = \frac{\text{point estimate} - \text{null value}}{SE}$$

in context...

$$point\ estimate = \bar{x}_{diff} = 1836$$

Note: Null value is 0 because in the null hypothesis we set $\mu_{diff} = 0$

Find the test statistic

Test statistic for inference on a small sample mean

The test statistic for inference on a small sample ($n < 50$) mean is the T statistic with $df = n - 1$

$$T_{df} = \frac{\text{point estimate} - \text{null value}}{SE}$$

in context...

$$point\ estimate = \bar{x}_{diff} = 1836$$

$$SE = \frac{s_{diff}}{\sqrt{n}} = \frac{1176}{\sqrt{10}} = 372$$

Note: Null value is 0 because in the null hypothesis we set $\mu_{diff} = 0$

Find the test statistic

Test statistic for inference on a small sample mean

The test statistic for inference on a small sample ($n < 50$) mean is the T statistic with $df = n - 1$

$$T_{df} = \frac{\text{point estimate} - \text{null value}}{SE}$$

in context...

$$point\ estimate = \bar{x}_{diff} = 1836$$

$$SE = \frac{s_{diff}}{\sqrt{n}} = \frac{1176}{\sqrt{10}} = 372$$

$$T = \frac{1836 - 0}{372} = 4.94$$

Note: Null value is 0 because in the null hypothesis we set $\mu_{diff} = 0$

Find the test statistic

Test statistic for inference on a small sample mean

The test statistic for inference on a small sample (n < 50) mean is the T statistic with $df = n - 1$

$$T_{df} = \frac{\text{point estimate} - \text{null value}}{SE}$$

in context...

$$\begin{aligned} \textit{point estimate} &= \bar{x}_{diff} = 1836 \\ SE &= \frac{s_{diff}}{\sqrt{n}} = \frac{1176}{\sqrt{10}} = 372 \\ T &= \frac{1836 - 0}{372} = 4.94 \\ df &= 10 - 1 = 9 \end{aligned}$$

Note: Null value is 0 because in the null hypothesis we set $\mu_{diff} = 0$

Finding the p-value

Finding the p-value

Locate the calculated T statistic on the appropriate *df* row, obtain the p-value from the corresponding column heading (one or two tail, depending on the alternative hypothesis).

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 1	3.08	6.31	12.71	31.82	63.66
2	1.89	2.92	4.30	6.96	9.92
3	1.64	2.35	3.18	4.54	5.84
.	.	.	.	.	
17	1.33	1.74	2.11	2.57	2.90
18	1.33	1.73	2.10	2.55	2.88
19	1.33	1.73	2.09	2.54	2.86
20	1.33	1.72	2.09	2.53	2.85
:	:	:	:	.	
400	1.28	1.65	1.97	2.34	2.59
500	1.28	1.65	1.96	2.33	2.59
∞	1.28	1.64	1.96	2.33	2.58

Finding the p-value (cont.)

one tail		0.100	0.050	0.025	0.010	0.005
two tails		0.200	0.100	0.050	0.020	0.010
df	6	1.44	1.94	2.45	3.14	3.71
	7	1.41	1.89	2.36	3.00	3.50
	8	1.40	1.86	2.31	2.90	3.36
	9	1.38	1.83	2.26	2.82	3.25
	10	1.37	1.81	2.23	2.76	3.17

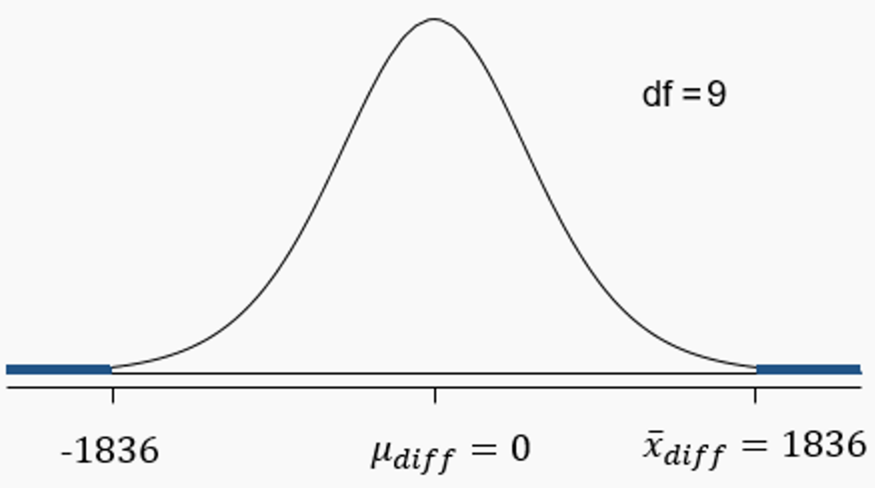

$T = 4.94$

Finding the p-value (cont.)

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.71
7	1.41	1.89	2.36	3.00	3.50
8	1.40	1.86	2.31	2.90	3.36
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

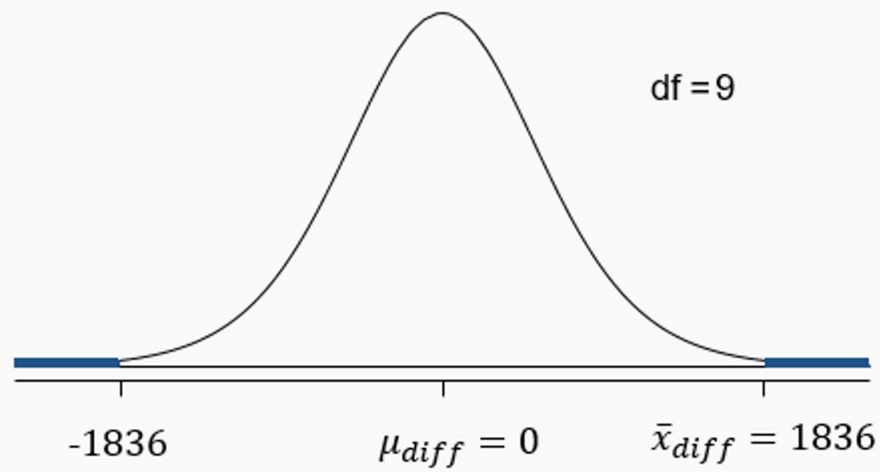

$T = 4.94$

Finding the p-value (cont.)

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.[illegible]1
7	1.41	1.89	2.36	3.00	3.50
8	1.40	1.86	2.31	2.90	3.36
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

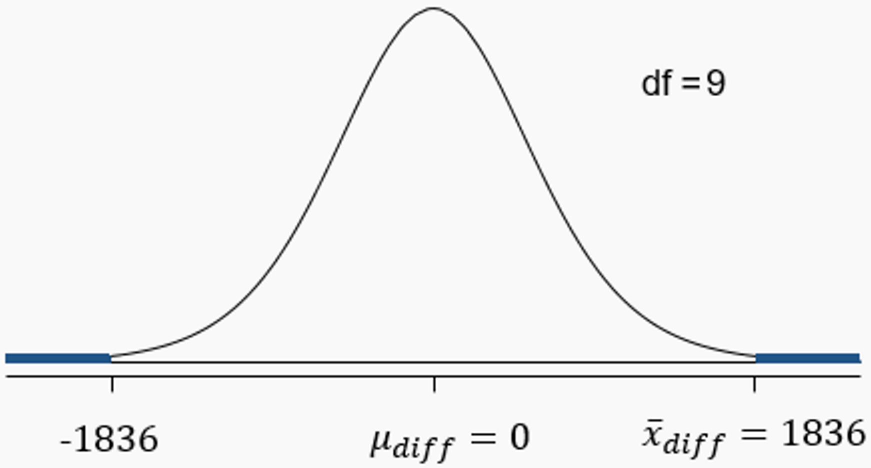

$T = 4.94$

Finding the p-value (cont.)

one tail		0.100	0.050	0.025	0.010	0.005
two tails		0.200	0.100	0.050	0.020	0.010
df	6	1.44	1.94	2.45	3.14	3.[illegible]1
	7	1.41	1.89	2.36	3.00	3.50
	8	1.40	1.86	2.31	2.90	3.36
	9	1.38	1.83	2.26	2.82	3.25
	10	1.37	1.81	2.23	2.76	3.17

$p < 0.010$

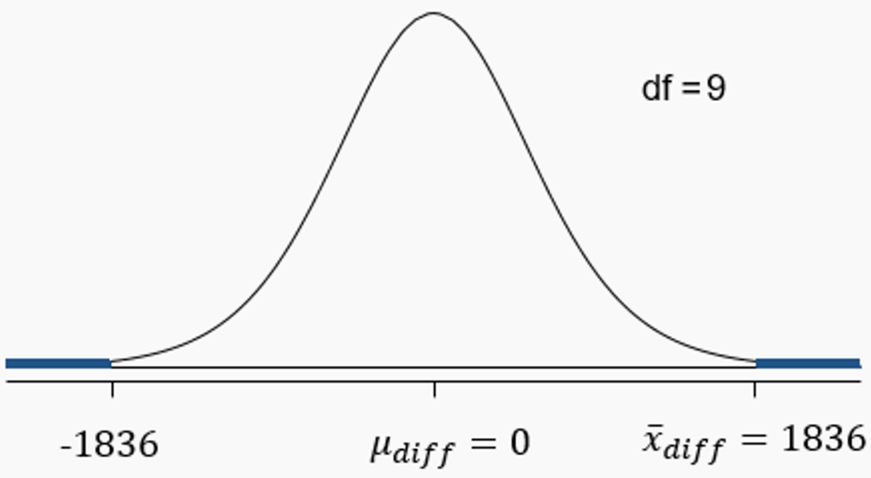

$T = 4.94$

Finding the p-value (cont.)

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.[illegible]1
7	1.41	1.89	2.36	3.00	3.[illegible]0
8	1.40	1.86	2.31	2.90	3.[illegible]6
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

$p < 0.010$

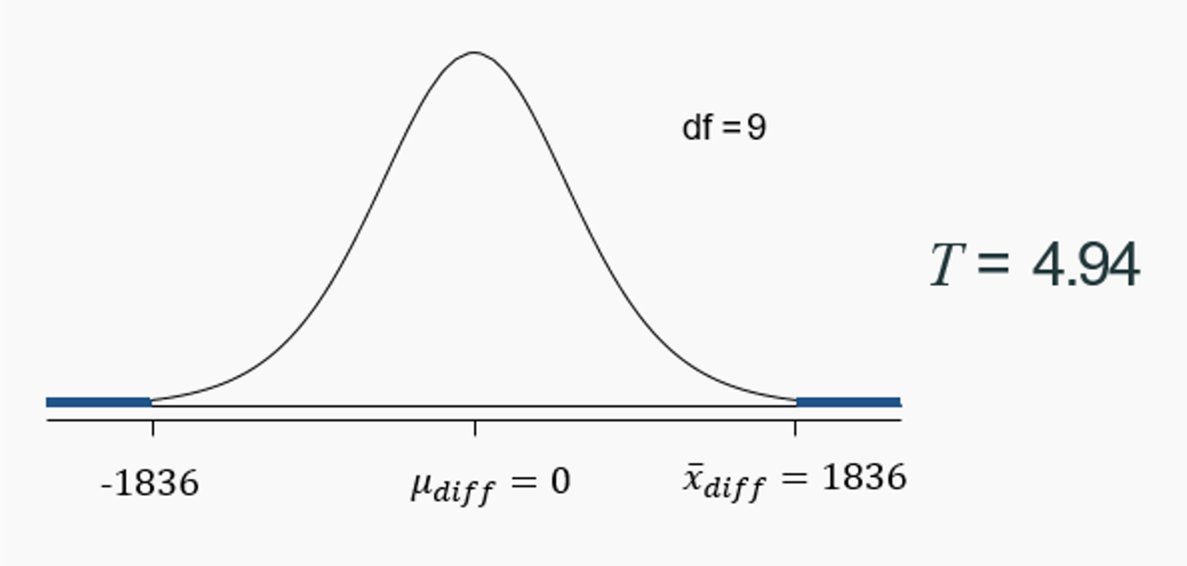

What is the conclusion of the hypothesis test?

Finding the p-value (cont.)

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.71
7	1.41	1.89	2.36	3.00	3.50
8	1.40	1.86	2.31	2.90	3.36
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

$p < 0.010$

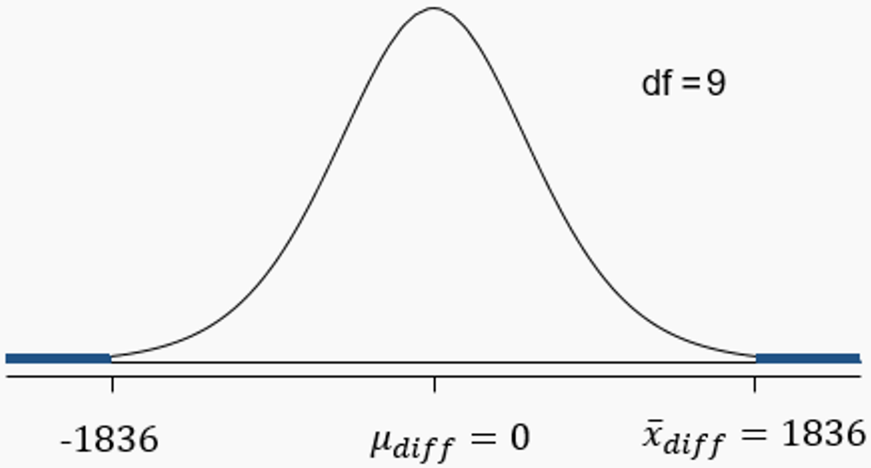

$T = 4.94$

What is the conclusion of the hypothesis test?

The data provides convincing evidence of a difference between traffic flow on Friday 6th and 13th

Conclusion of the test

What is the conclusion of this hypothesis test?

Since the p-value is quite low (< 0.01), we conclude that the data provide strong evidence of a difference between traffic flow on Friday 6th and 13th.

What is the difference?

- We concluded that there is a difference in the traffic flow between Friday 6th and 13th

What is the difference?

- We concluded that there is a difference in the traffic flow between Friday 6th and 13th
- But it would be more interesting to find out what exactly this difference is

What is the difference?

- We concluded that there is a difference in the traffic flow between Friday 6th and 13th
- But it would be more interesting to find out what exactly this difference is
- We can use a confidence interval to estimate this difference

Confidence interval for a small sample mean

- Confidence intervals are always of the form

$$point\ estimate \pm ME$$

Confidence interval for a small sample mean

- Confidence intervals are always of the form

$$point\ estimate \pm ME$$

- ME is always calculated as the product of a critical value and SE

Confidence interval for a small sample mean

- Confidence intervals are always of the form

$$point\ estimate \pm ME$$

- ME is always calculated as the product of a critical value and SE
- Since small sample means follow a ***t*** distribution (and not a z distribution), the critical value is a t^* (as opposed to a z^*?).

$$point\ estimate \pm t^*\ x\ SE$$

Finding the critical t (t^*)

- We can find the critical t (t*) that our calculated t would need to exceed to result in a p-value less than α

Finding the critical t (t^*)

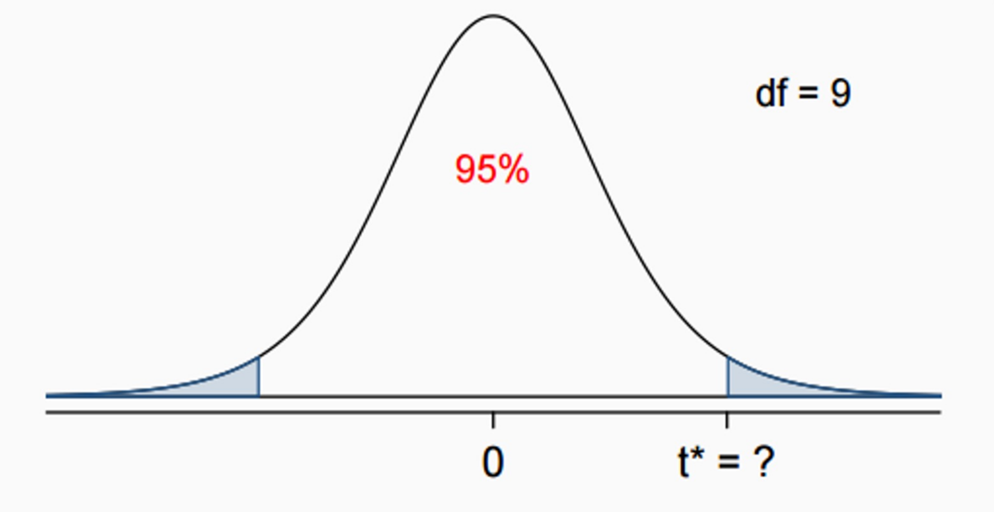

n = 10, df = 10 - 1 = 9, t^* is at the intersection of rows df = 9 and two tail probability 0.05

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.71
7	1.41	1.89	2.36	3.00	3.50
8	1.40	1.86	2.31	2.90	3.36
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

Finding the critical t (t^*)

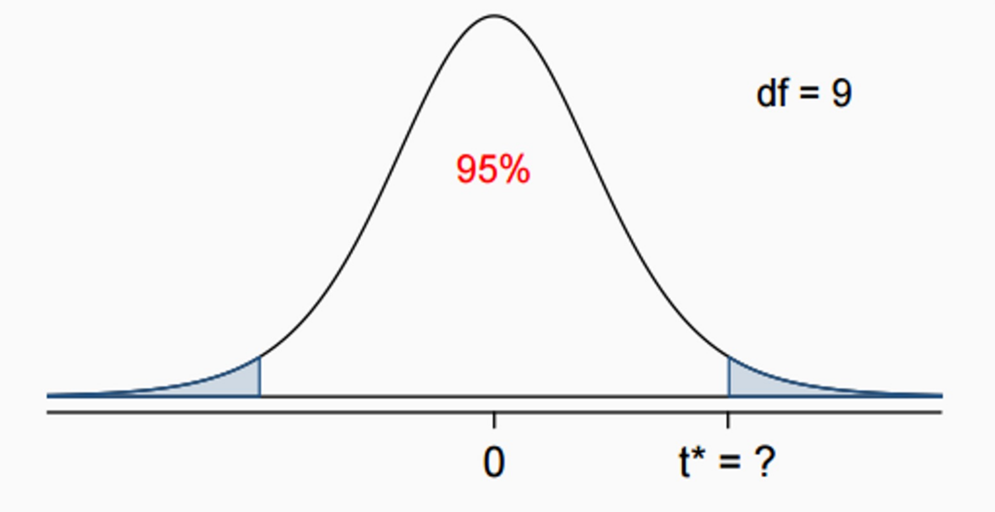

n = 10, df = 10 - 1 = 9, t^* is at the intersection of rows df = 9 and two tail probability 0.05

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.71
7	1.41	1.89	2.36	3.00	3.50
8	1.40	1.86	2.31	2.90	3.36
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

Finding the critical t (t^*)

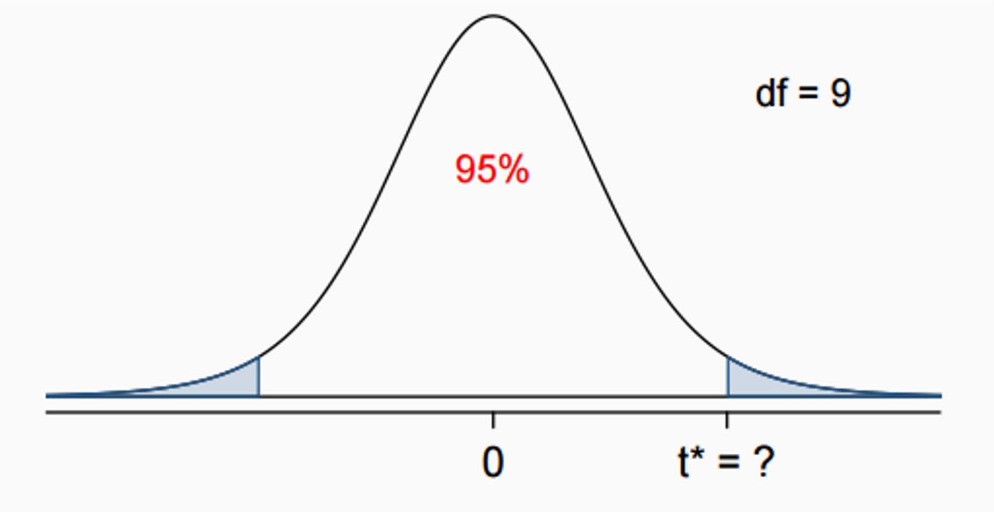

n = 10, df = 10 - 1 = 9, t^* is at the intersection of rows df = 9 and two tail probability 0.05

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	0.050	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.71
7	1.41	1.89	2.36	3.00	3.50
8	1.40	1.86	2.31	2.90	3.36
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

Finding the critical t (t^*)

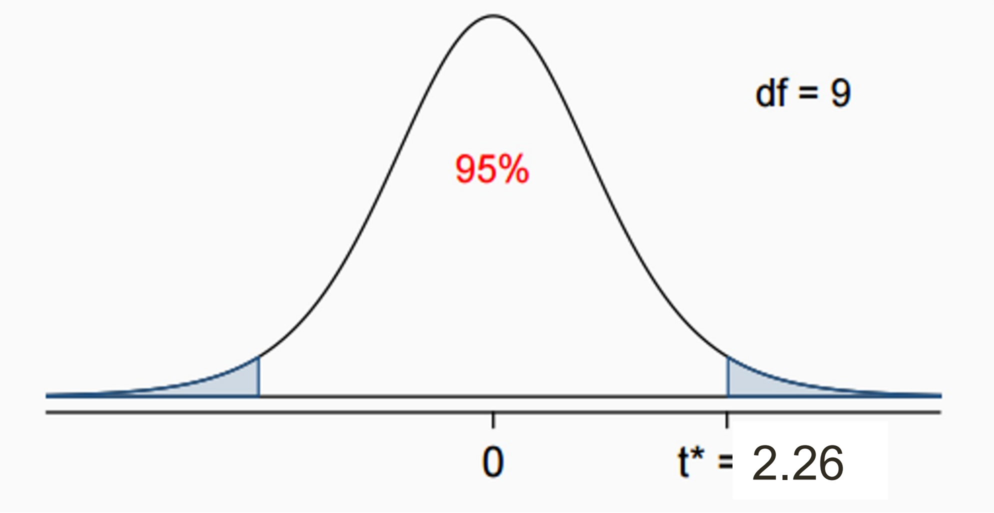

n = 10, df = 10 - 1 = 9, t^* is at the intersection of rows df = 9 and two tail probability 0.05

one tail	0.100	0.050	0.025	0.010	0.005
two tails	0.200	0.100	*0.050*	0.020	0.010
df 6	1.44	1.94	2.45	3.14	3.71
7	1.41	1.89	2.36	3.00	3.50
8	1.40	1.86	2.31	2.90	3.36
9	1.38	1.83	2.26	2.82	3.25
10	1.37	1.81	2.23	2.76	3.17

Constructing a CI for a small sample mean

Which of the following is the correct calculation of a 95% confidence interval for the difference between the traffic flow between Friday 6th and 13th?

$point\ estimate \pm t^* \ x\ SE$ t* = 2.26

$\bar{x}_{diff} = 1836$ $s_{diff} = 1176$ $n = 10$ $SE = 372$

A. 1836 ± 1.96 x 372
B. 1836 ± 2.26 x 372
C. 1836 ± -2.26 x 372
D. 1836 ± 2.26 x 1176

Constructing a CI for a small sample mean

Which of the following is the correct calculation of a 95% confidence interval for the difference between the traffic flow between Friday 6th and 13th?

$point\ estimate \pm t^* \ x\ SE$ t* = 2.26

$\bar{x}_{diff} = 1836$ $s_{diff} = 1176$ $n = 10$ $SE = 372$

A. 1836 ± 1.96 x 372

B. 1836 ± 2.26 x 372 → *(995, 2677)*

C. 1836 ± -2.26 x 372

D. 1836 ± 2.26 x 1176

Interpreting the CI

Which of the following is the *best* interpretation for the confidence interval we just calculated?

$$\mu_{diff:6th-13th} = (995, 2677)$$

We are 95% confident that…

A. the difference between the average number of cars on the road on Friday 6th and 13th is between 995 and 2,677
B. on Friday 6th there are 995 to 2,677 fewer cars on the road than on the Friday 13th, on average
C. on Friday 6th there are 995 fewer to 2,677 more cars on the road than on the Friday 13th, on average
D. on Friday 13th there are 995 to 2,677 fewer cars on the road than on the Friday 6th, on average

Interpreting the CI

Which of the following is the *best* interpretation for the confidence interval we just calculated?

$$\mu_{diff:6th-13th} = (995, 2677)$$

We are 95% confident that…

A. the difference between the average number of cars on the road on Friday 6th and 13th is between 995 and 2,677
B. on Friday 6th there are 995 to 2,677 fewer cars on the road than on the Friday 13th, on average
C. on Friday 6th there are 995 fewer to 2,677 more cars on the road than on the Friday 13th, on average
D. on Friday 13th there are 995 to 2,677 fewer cars on the road than on the Friday 6th, on average

Synthesis

Does the conclusion from the hypothesis test agree with the findings of the confidence interval?

Do you think the findings of this study suggests that people believe Friday 13th is a day of bad luck?

Synthesis

Does the conclusion from the hypothesis test agree with the findings of the confidence interval?

Yes, the hypothesis test found a significant difference, and the CI does not contain the null value of 0.

Do you think the findings of this study suggests that people believe Friday 13th is a day of bad luck?

Synthesis

Does the conclusion from the hypothesis test agree with the findings of the confidence interval?

Yes, the hypothesis test found a significant difference, and the CI does not contain the null value of 0

Do you think the findings of this study suggests that people believe Friday 13th is a day of bad luck?

No, this is an observational study. We have just observed a significant difference between the number of cars on the road on these two days. We have not tested for people's beliefs

Recap: Inference using the t-distribution

- If $\boldsymbol{\sigma}$ is unknown, use the $\boldsymbol{t}$-distribution with $SE = \frac{s}{\sqrt{n}}$

Recap: Inference using the t-distribution

- If σ is unknown, use the t-distribution with $SE = \frac{s}{\sqrt{n}}$
- Conditions:
  - independence of observations (often verified by a random sample, and if sampling without replacement, n < 10% of population)
  - no extreme skew

Recap: Inference using the t-distribution

- If $\boldsymbol{\sigma}$ is unknown, use the $\boldsymbol{t}$-distribution with $SE = \frac{s}{\sqrt{n}}$
- Conditions:
  - independence of observations (often verified by a random sample, and if sampling without replacement, n < 10% of population)
  - no extreme skew
- Hypothesis Testing:

$$T_{df} = \frac{point\ estimate\ - null\ value}{SE}, where\ df = n - 1$$

Recap: Inference using the t-distribution

- If σ is unknown, use the t-distribution with $SE = \frac{s}{\sqrt{n}}$
- Conditions:
  - independence of observations (often verified by a random sample, and if sampling without replacement, n < 10% of population)
  - no extreme skew
- Hypothesis Testing:

$$T_{df} = \frac{point\ estimate\ - null\ value}{SE}, where\ df = n - 1$$

- Confidence interval:

$$point\ estimate \pm t^*_{df} \times SE$$

Recap: Inference using the *t*-distribution

- If $\boldsymbol{\sigma}$ is unknown, use the $\boldsymbol{t}$-distribution with $SE = \frac{s}{\sqrt{n}}$
- Conditions:
  - independence of observations (often verified by a random sample, and if sampling without replacement, n < 10% of population)
  - no extreme skew
- Hypothesis Testing:

$$T_{df} = \frac{point\ estimate\ - null\ value}{SE}, where\ df = n - 1$$

- Confidence interval:

$$point\ estimate \pm t^{*}_{df} \times SE$$

Note: The example we used was for paired means (difference between dependent groups). We took the difference between the observations and used only these differences (one sample) in our analysis

Useful tool

http://gallery.shinyapps.io/dist_calc/

(a) Let X denote returns on this portfolio, then $X \sim N(\mu = 14.7, \sigma = 33)$.

Useful tool

http://gallery.shinyapps.io/dist_calc/

(a) Let X denote returns on this portfolio, then $X \sim N(\mu = 14.7, \sigma = 33)$.

$$P(X < 0) = P\left(Z < \frac{0 - 14.7}{33}\right) = P(Z < -0.45) = 0.3264 \rightarrow 32.64\%$$

Useful tool

http://gallery.shinyapps.io/dist_calc/

Let X be the number of people who have consumed alcohol in the group. Then, using a binomial distribution with $n = 5$ and $p = 0.697$:

$P(X \leq 2)$

Useful tool

http://gallery.shinyapps.io/dist_calc/

Let X be the number of people who have consumed alcohol in the group. Then, using a binomial distribution with $n = 5$ and $p = 0.697$:

$$
\begin{aligned}
P(X \leq 2) &= P(X = 0) + P(X = 1) + P(X = 2) \\
&= \binom{5}{0} \times 0.697^0 \times 0.303^5 + \binom{5}{1} \times 0.697^1 \times 0.303^4 + \binom{5}{2} \times 0.697^2 \times 0.303^3 \\
&= 0.0026 + 0.0293 + 0.1351 \\
&= 0.167
\end{aligned}
$$

Useful tool

http://gallery.shinyapps.io/dist_calc/

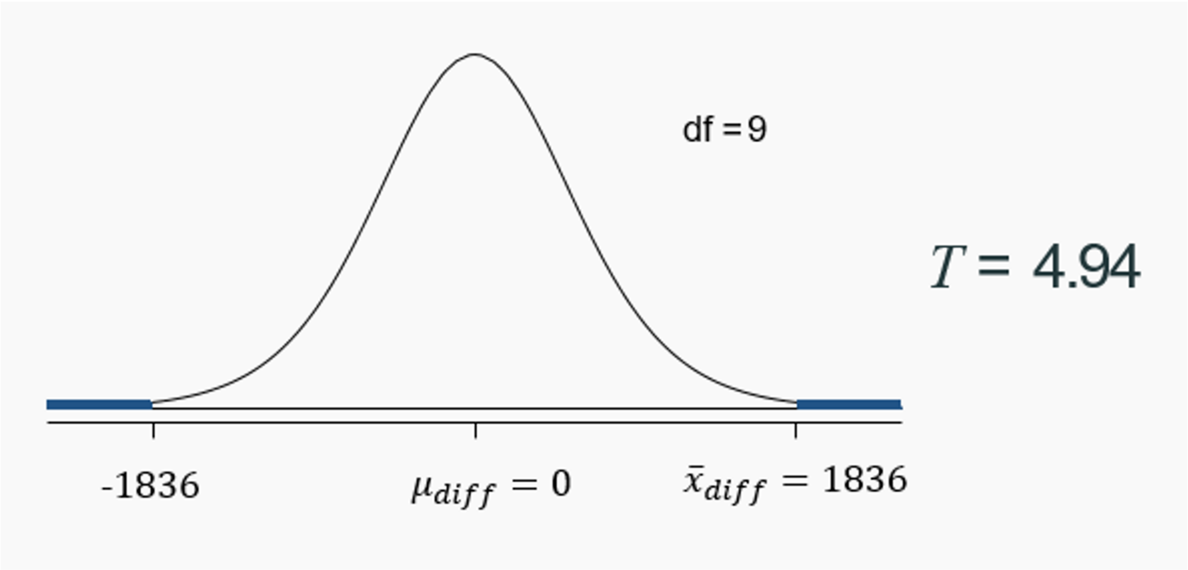

Useful tool

http://gallery.shinyapps.io/dist_calc/

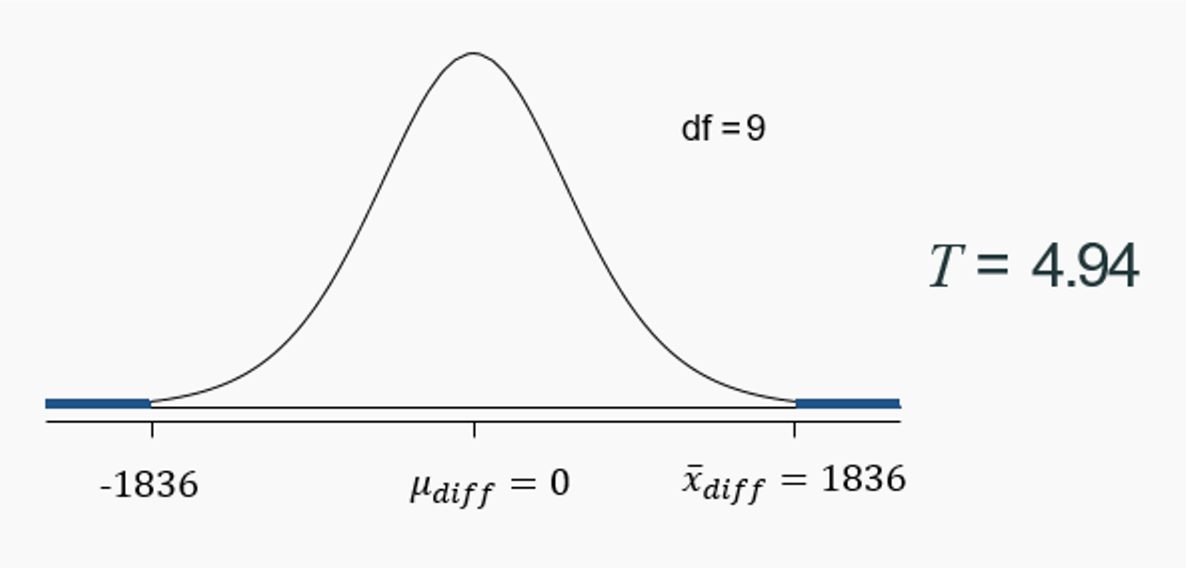

p = 0.0006