Large Samples (n > 29)

- The sampl**ing** distribution of $\overline{Y}$ is normally distributed if the sample is random and has at least 30 cases, even if the population distribution is not.
- Thus, we can rely on the Central Limit Theorem, which tells us six things about sampl**ing** distributions
- Therefore we can use Table A to determine the likelihood of an event (or, the probability of finding a greater difference from H_0 than the one observed in the sample).

SOC424 – Statistics w/ Ellis Godard

Student's t: Tests w/ Small Samples

Small Samples (n < 30)

- The sampl**ing** distribution of $\overline{Y}$ is **not** normally distributed when the sample size is small (**<30**) and the population distribution is not normal. *(see Fig. 4.15, p.104)*
- Thus, if n is small we cannot rely on the Central Limit Theorem, and therefore can't use Table A to determine the likelihood of an event (or, the probability of finding a greater difference from H_0 than the one observed in the sample)

OUTLINE

- **Catchup / Questions / Breathe**
  - Short lecture… lots of time… let's use it… ***now*** :)
- **Student's "t" Distribution**
  - **Small Samples (n < 30) can't use "z"**
  - **Degrees of Freedom**
  - Calculated vs. **Critical Values**
  - Table B
- **Examples**
- **Connection to Confidence Intervals**

The t-distribution

If we assume that the *population* distribution (not sampling distribution) of a variable is normal, then for a random sample of *any* size n (even 30 or more),

$$t = \frac{\overline{Y} - \hat{\mu}}{\hat{\sigma}_{\overline{Y}}}$$

is called the (Student) t distribution, with (n - 1) degrees of freedom.

"Student" = pseudonym of Guinness statistician/chemist; see text

Examining the Distributions

- z-score distribution
- z-test distribution
- one-sample t-test
- two-sample t-test

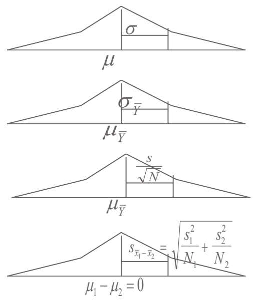

Properties of "t" distribution

(Also, see Figure 6.9, p. 181)

- The t distribution is symmetric about 0;
- The t distribution is more dispersed with fewer degrees of freedom (e.g. smaller samples);
- As the d.f. (e.g. sample size) increases, the t distribution approximates a standard normal (z) distribution.

- Otherwise, very much like z
  - Measures/represents a distance between scores
  - The bigger that difference, the smaller the tails

Degrees of Freedom - *idea*

- How much do we have to go on?

  How much do we have to assume?

- For t-test of a mean, df = n-1
- For a difference of means, df = n1 + n2 - 2

Degrees of Freedom - *effect*

- DF Determines *shape* of t distribution
- The higher the degrees of freedom, the lower the standard deviation of t
  - Mean of t-distribution is still zero, just like z
  - When the degrees of freedom are high enough (n=30), the standard deviation = 1 (so the t-distribution is "standard normal distribution" – that is, the t distribution is the z distribution if n>29)
  - With lower degrees of freedom, the standard deviation is higher than 1, and the t-distribution is flatter than the z distribution (higher kurtosis)

Small Sample Hypotest: Same Steps

Step 1 Assumptions ARE NOT HYPOTHESES!!

1. **Assumptions**
   - Still must assume the measurements are some LOM (interval or categorical)
   - Still must assume that you have a random sample
   - Now n<30 so can't assume sampling distrib is normal, so we need t distribution, so have to assume that the population distribution is normal
2. **Hypotheses**
3. **Calculate test statistic**
   - Different formulas, depending on whether you're testing for a mean or proportion (last lecture), difference of means or difference of proportions (next two), or something else (F, r, chi-square, et al.)
4. **Find p-value**
   - Now use t instead of z
   - Use t table to find "critical" value
     - If the observed t exceeds the critical t, p is less than alpha
     - If the observed t is less than the critical t, p is more than alpha
5. **Conclusions**
   - As before, in terms of *both* the Ho and Ha

Calculated vs Critical "t"

- There's some "critical" t that would be far enough to leave the "alpha" we want

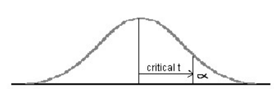

- Must compare that to the "calculated" t
  - Won't get a specific p value
  - Can only say whether p > or < some alpha

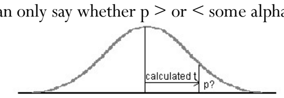

Tabular vs. Calculated "t" (illustrated)

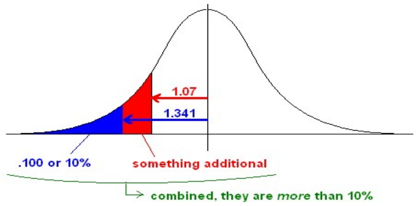

- If t = -1.07, then p > 0.100
  - We would need to observe a difference of 1.341 standard errors to leave only 10% in the tail.
  - The difference we observed isn't large enough (the observed difference isn't sufficient) to reject the null, since the error is more than 10% in *each* direction.

Finding a "critical t" value

- t_a refers to "t" value corresponding to an area of "a" in one tail, for a given d.f.
  - a ~ p; = the probability of getting a bigger t
  - area depends on the degrees of freedom
- Use Table B to find the p
  - In table A, found z for given p (e.g. 5%)
  - Now, what's "t" for 10% in each tail (t.100) with n=16 (df = 16 -1 = 15)?
  - Here's how…

Reading Table B (t's instead of z's; p.669)

- A had z's on outside (row/col headers), p's in body
- In B, Row = df, Column = p, t's inside the table
- Sample size (n) for California was 16
- So, df = n – 1 = 16 –1 = 15

If alpha = .100 (10% in tails), confidence = 90%

df	$t_{.100}$	$t_{.050}$	$t_{.025}$	$t_{.010}$	$t_{.005}$
1	3.078	6.314	12.7063	1.821	63.657
2	1.886	2.920	4.303	6.965	9.925
15	1.341	1.753	2.131	2.602	2.947

This is the "tabular t value", the t value (a number of standard errors) high enough to leave only 10% in *one* tail

Table A vs Table B

P smaller when Z bigger

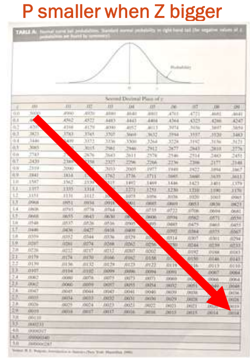

T smaller when p *or* n bigger

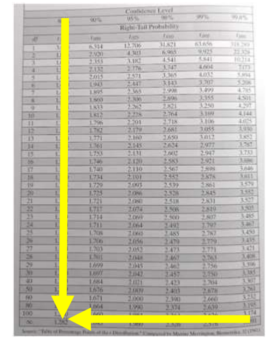

Table B

Suppose I want to be 90% confident but only have 4 cases (not 30)

Pick the column that has the right tail p

For 90% confidence, p is 0.050 so $t_{.050}$

Pick the row based on sample size (df=n-1), so with 4 cases, df=3

The body is t's, not p's

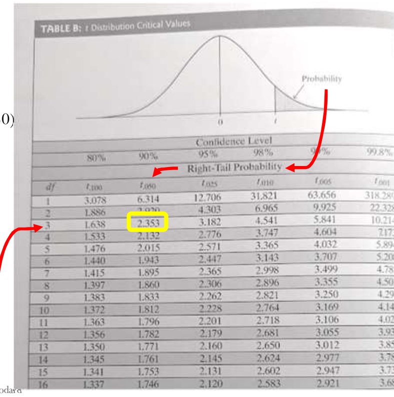

TABLE B: t Distribution Critical Values

Probability

	Confidence Level 80%	90%	95%	98%	99%	99.8%
	Right-Tail Probability					
df	$t_{.100}$	$t_{.050}$	$t_{.025}$	$t_{.010}$	$t_{.005}$	$t_{.001}$
1	3.078	6.314	12.706	31.821	63.656	318.289
2	1.886	2.920	4.303	6.965	9.925	22.328
3	1.638	2.353	3.182	4.541	5.841	10.214
4	1.533	2.132	2.776	3.747	4.604	7.173
5	1.476	2.015	2.571	3.365	4.032	5.894
6	1.440	1.943	2.447	3.143	3.707	5.208
7	1.415	1.895	2.365	2.998	3.499	4.785
8	1.397	1.860	2.306	2.896	3.355	4.501
9	1.383	1.833	2.262	2.821	3.250	4.297
10	1.372	1.812	2.228	2.764	3.169	4.144
11	1.363	1.796	2.201	2.718	3.106	4.025
12	1.356	1.782	2.179	2.681	3.055	3.930
13	1.350	1.771	2.160	2.650	3.012	3.852
14	1.345	1.761	2.145	2.624	2.977	3.787
15	1.341	1.753	2.131	2.602	2.947	3.733
16	1.337	1.746	2.120	2.583	2.921	3.686

Table B

If d.f. = 5:

- What is $t_{.100}$?
- What is $t_{.010}$?

If n=16 (so df=15)

- what t gives 90% confidence ($t_{.050}$)?
- what t gives 99% confidence ($t_{.005}$)?

Example above:

- 95% w/ n=16

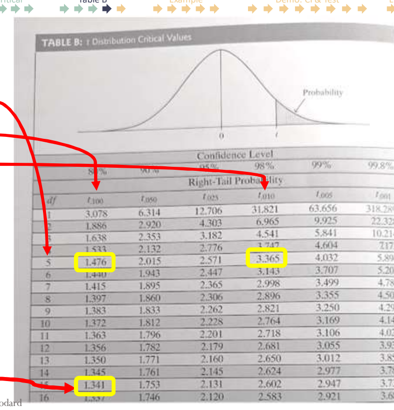

TABLE B: t Distribution Critical Values

Probability

	Confidence Level 80%	90%	95%	98%	99%	99.8%
	Right-Tail Probability					
df	$t_{.100}$	$t_{.050}$	$t_{.025}$	$t_{.010}$	$t_{.005}$	$t_{.001}$
1	3.078	6.314	12.706	31.821	63.656	318.289
2	1.886	2.920	4.303	6.965	9.925	22.328
3	1.638	2.353	3.182	4.541	5.841	10.214
4	1.533	2.132	2.776	3.747	4.604	7.173
5	1.476	2.015	2.571	3.365	4.032	5.894
6	1.440	1.943	2.447	3.143	3.707	5.208
7	1.415	1.895	2.365	2.998	3.499	4.785
8	1.397	1.860	2.306	2.896	3.355	4.501
9	1.383	1.833	2.262	2.821	3.250	4.297
10	1.372	1.812	2.228	2.764	3.169	4.144
11	1.363	1.796	2.201	2.718	3.106	4.025
12	1.356	1.782	2.179	2.681	3.055	3.930
13	1.350	1.771	2.160	2.650	3.012	3.852
14	1.345	1.761	2.145	2.624	2.977	3.787
15	1.341	1.753	2.131	2.602	2.947	3.733
16	1.337	1.746	2.120	2.583	2.921	3.686

Table B – bottom of the table…

What if we want to be 95% confident (5% risk) and have more than 100 cases?

The bigger the sample, the closer t is to z

At *some* point, t & z are the same

This table says >100. Earlier editions said >30 *Any* number is arbitrary

	Confidence Level					
	80%	90%	95%	98%	99%	99.8%
	Right-Tail Probability					
df	$t_{.100}$	$t_{.050}$	$t_{.025}$	$t_{.010}$	$t_{.005}$	$t_{.001}$
1	3.078	6.314	12.706	31.821	63.656	318.289
19	1.328	1.729	2.093	2.539	2.861	3.579
20	1.325	1.725	2.086	2.528	2.845	3.552
21	1.323	1.721	2.080	2.518	2.831	3.527
22	1.321	1.717	2.074	2.508	2.819	3.505
23	1.319	1.714	2.069	2.500	2.807	3.485
24	1.318	1.711	2.064	2.492	2.797	3.467
25	1.316	1.708	2.060	2.485	2.787	3.450
26	1.315	1.706	2.056	2.479	2.779	3.435
27	1.314	1.703	2.052	2.473	2.771	3.421
28	1.313	1.701	2.048	2.467	2.763	3.408
29	1.311	1.699	2.045	2.462	2.756	3.396
30	1.310	1.697	2.042	2.457	2.750	3.385
40	1.303	1.684	2.021	2.423	2.704	3.307
50	1.299	1.676	2.009	2.403	2.678	3.261
60	1.296	1.671	2.000	2.390	2.660	3.232
80	1.292	1.664	1.990	2.374	2.639	3.195
100	1.290	1.660	1.984	2.364	2.626	3.174
∞	1.282	1.645	1.960	2.326	2.576	3.091

But use *this* table, so we're on the same page!

An Example Test w/ Table B…

Last week, did a hypothesis test that the mean years of education in the South is less than the U.S. as a whole

There were 411 cases (>29) so we could assume the sampling distribution was normal.

But what if we had a smaller group, with < 30 cases?

Example Small Sample Hypotest

Education - Highest Year of School Completed

Region	Mean	Std Dev	Std. Error
U.S. (n=1510)	12.884	2.984	0.077
South (n=411)	12.460	3.352	0.165
Non-South (n=1099)	13.043	2.819	0.085
California (n=16)	**12.084**	**2.987**	**0.747**

Note: The California results are based on pure speculation. ☺

1. Assumptions

- We assume **interval data**, since we're using a mean and calculating the standard error based on a mean & standard deviation (which only intervals have!).
- But we don't have a large sample (since n <30) so we *cannot* assume that the sampling distribution is normal, and must assume that the **population distribution is normal.**
- We always assume a **random sample** (though it sometimes isn't).

2. Null and Alternative Hypotheses:

H_o: $\mu_{CA} = 12.88$ (No difference in schooling between CA & US)

H_a: $\mu_{CA} < 12.88$ (Mean years of schooling lower in CA than US)

3. Test Statistic

$$t = \frac{\bar{Y}_{CA} - 12.884}{s_{\bar{Y}_{CA}}} = \frac{12.084 - 12.884}{0.747}$$

$$t = -1.07$$

Interpretation: The mean number of years of education of California respondents was 1.07 standard errors away from the mean of respondents nationally.

P-value:

If t = -1.07, then p > 0.100. . Since that p-value is very high (>10% so >5%), there is too much probability of getting this difference just due to chance, which means too much risk that we would be wrong if we rejectedthe null hypothesis.

Conclusion: Since our p-value exceeds conventional levels of significance, we cannot reject the null hypothesis. We cannot conclude that educational attainment in California is any different than in the U.S. as a whole.

Small Sample Conf. Interval

- Just like with large sample – sample formula
- Only difference: Use t table instead of z

$$\bar{Y} \pm (t^* \hat{\sigma}_{\bar{Y}}) \approx \bar{Y} \pm \left(t^* \frac{s}{\sqrt{n}}\right)$$

- Still a mean plus/minus some # of standard deviations; just get a *different* # of them since n<30.

In-class demo…

- You're going to these things (***by hand***) in lab:
  1. Construct a set of *sample* data (*on paper, not in SPSS*)
  2. Construct (*& interpret!*) a 95% confidence interval
  3. Conduct a hypothesis test (*all five steps*; α =.05)
  4. Compare the results – they should be consistent!
- I'll do a quick demo in Excel – you shouldn't ☺
  - I'll paste a screenshot of that here afterwards.
  - You can mimic the layout, but don't have to do
  - But you do ~~have~~ get to *do the math* – I want to *see the work*
- First, I need small sample of students' units…

Example: Small Sample CI vs Hypotest

1. Let's make a dataset of units that one row is taking.
2. Let's estimate, with 95% confidence, how many units students in the full population take on average.
3. Let's conduct a hypothesis test (all five steps; α =.05) for whether the average student of the population takes more than 10 units.
4. Let's compare the results of the CI and hypothesis test.
   - Are they consistent?

In-class demo (10/23/25)…

A. My sample		**C. Hypothesis Test**	
a	13	*All five steps!*	
b	16	**1. Assumptions**	
c	17		Random sample
d	15		Interval variable
e	17		***Population distribution is normal (because n<30)***
		2. Hypotheses	
Univariate Statistics			Ha: M>10
n:	5		Ho: M=10
Mean:	15.6	**3. Test Statistic**	
StdDev	1.67332		t = (10 - 15.6) / 0.74833
S.E. Mean:	1.67332 / Sqrt(5) =		-7.483314774
	0.74833	**4. P Value**	
Crit.t (table b)	2.776		Table B gives a critical 't' of 2.776, for a sample size of 5 (df=4) and 2.5% tail.
B. Confidence Interval:			Since the calc t of -7.483 is larger than the critical t of 2.776
Margin:	2.776 * 0.74833		then the p value (the likelihood of being wrong if I reject the null)
	= 2.07737		is smaller than 2.5%
Conf.Int.	15.6 +/- 2.776 * 0.74833	**5. Conclusion**	
	= 15.6 +/- 2.07737		Since the observed (sample) mean is -7.48 standard errors from the hypothesized
	= 13.52263		mean, and we needed it to be 2.776 standard errors away, our sample is different
	to 17.67737		enough from the null to be confident that the population probably differs from it.
Interp:	We're 95% confident about an		We can reject the null, that 'Mu'=10, and can support the alternative, that 'Mu'>10.
	interval that doesn't contain 10.		

D. Consistency: Both the CI and the Hypotest suggest that the population parameter might not be 10.

In-class demo (3/27/25)…

A. My sample		**C. Hypothesis Test**	
a	12	*All five steps!*	
b	15	**1. Assumptions**	
c	12		Random sample
d	19		Interval variable
e	14		***Population distribution is normal (because n<30)***
		2. Hypotheses	
Univariate Statistics			Ha: M>12
n:	5		Ho: M=12
Mean:	14.4	**3. Test Statistic**	
StdDev	2.88097		t = (12 - 14.4) / 1.28841
S.E. Mean:	2.88097 / Sqrt(5) =		-1.862761262
	1.28841	**4. P Value**	
Crit.t (table b)	2.776		Table B gives a critical 't' of 2.776, for a sample size of 5 (df=4) and 2.5% tail.
B. Confidence Interval:			Since the calc t of -1.863 is smaller than the critical t of 2.776
Margin:	2.776 * 1.28841		then the p value (the likelihood of being wrong if I reject the null)
	= 3.57663		is larger than 2.5%
Conf.Int.	14.4 +/- 2.776 * 1.28841	**5. Conclusion**	
	= 14.4 +/- 3.57663		Since the observed (sample) mean is -1.86 standard errors from the hypothesized
	= 10.82337		mean, and we needed it to be 2.776 standard errors away, our sample isn't different
	to 17.97663		enough from the null to be confident that the population probably differs from it.
Interp:	We're 95% confident about an		We can't reject the null, that 'Mu'=12, and can't support the alternative, that 'Mu'>12.
	interval that does contain 12.		

D. Consistency: Both the CI and the Hypotest suggest that the population parameter might be 12.

In-class demo (4/8/24)…

A. My sample		**C. Hypothesis Test**	
a	16	*All five steps!*	
b	10	**1. Assumptions**	
c	13		Random sample
d	13		Interval variable
e			***Population distribution is normal (because n<30)***
		2. Hypotheses	
Univariate Statistics			Ha: M>12
n:	4		Ho: M=12
Mean:	13	**3. Test Statistic**	
StdDev	2.44949		t = (12 - 13) / 1.22474
S.E. Mean:	2.44949 / Sqrt(4) =		-0.816496581
	1.22474	**4. P Value**	
Crit.t (table b)	3.182		Table B gives a critical 't' of 3.182, for a sample size of 4 (df=3) and 2.5% tail.
B. Confidence Interval:			Since the calc t of -0.816 is smaller than the critical t of 3.182
Margin:	3.182 * 1.22474		then the p value (the likelihood of being wrong if I reject the null)
	= 3.89714		is larger than 2.5%
Conf.Int.	13 +/- 3.182 * 1.22474	**5. Conclusion**	
	= 13 +/- 3.89714		Since the observed (sample) mean is -0.82 standard errors from the hypothesized
	= 9.10286		mean, and we needed it to be 3.182 standard errors away, our sample isn't different
	to 16.89714		enough from the null to be confident that the population probably differs from it.
Interp:	We're 95% confident about an		We can't reject the null, that 'Mu'=12, and can't support the alternative, that 'Mu'>12.
	interval that does contain 12.		

D. Consistency: Both the CI and the Hypotest suggest that the population parameter might be 12.

In-class demo (10/26/23)…

A. My sample		**C. Hypothesis Test**	
a	17	*All five steps!*	
b	16	**1. Assumptions**	
c	10		Random sample
d	17		Interval variable
e	14		***Population distribution is normal (because n<30)***
		2. Hypotheses	
Univariate Statistics			Ha: M>12
n:	5		Ho: M=12
Mean:	14.8	**3. Test Statistic**	
StdDev	2.94958		t = (12 - 14.8) / 1.31909
S.E. Mean:	2.94958 / Sqrt(5) =		-2.122674522
	1.31909	**4. P Value**	
Crit.t (table b)	2.776		Table B gives a critical 't' of 2.776, for a sample size of 5 (df=4) and 2.5% tail.
B. Confidence Interval:			Since the calc t of -2.123 is smaller than the critical t of 2.776
Margin:	2.776 * 1.31909		then the p value (the likelihood of being wrong if I reject the null)
	= 3.6618		is larger than 2.5%
Conf.Int.	14.8 +/- 2.776 * 1.31909	**5. Conclusion**	
	= 14.8 +/- 3.6618		Since the observed (sample) mean is -2.12 standard errors from the hypothesized
	= 11.1382		mean, and we needed it to be 2.776 standard errors away, our sample isn't different
	to 18.4618		enough from the null to be confident that the population probably differs from it.
Interp:	We're 95% confident about an		We can't reject the null, that 'Mu'=12, and can't support the alternative, that 'Mu'>1
	interval that does contain 12.		
D. Consistency: Both the CI and the Hypotest suggest that the population parameter might be 12.			

In-class demo (11/1/21)…

A. My sample		**C. Hypothesis Test**	
Allison	13	*All five steps!*	
Araksya	11	**1. Assumptions**	
Edrik	10		Random sample
Aaron	17		Interval variable
			Population distribution is normal (because n<30)
Univariate Statistics		**2. Hypotheses**	
n:	4		Ha: M>12
Mean:	12.75		Ho: M=12
StdDev	3.09570	**3. Test Statistic**	
S.E. Mean:	3.0957 / Sqrt(4) =		t = (12 - 12.75) / 1.54785
	1.54785		-0.484543076
Crit.t (table b)	3.182	**4. P Value**	
B. Confidence Interval:			Table B gives a critical 't' of 3.182, for a sample size of 4 (df=3) and 2.5% tail.
Margin:	3.182 * 1.54785		Since the calc t of -0.485 is smaller than the critical t of 3.182
	= 4.92526		then the p value (the likelihood of being wrong if I reject the null)
Conf.Int.	12.75 +/- 3.182 * 1.54785		is larger than 2.5%
	= 12.75 +/- 4.92526	**5. Conclusion**	
	= 7.82474		Since the observed (sample) mean is -0.48 standard errors from the hypothesized
	to 17.67526		mean, and we needed it to be 3.182 standard errors away, our sample isn't different
Interp:	We're 95% confident about an		enough from the null to be confident that the population probably differs from it.
	interval that does contain 12.		We can't reject the null, that 'Mu'=12, and can't support the alternative, that 'Mu'>12.
D. Consistency: Both the CI and the Hypotest suggest that the population parameter might be 12.			

Lab: Small Sample CI ***and*** Hypotest

1. Construct a set of *sample* data from the ages of your group's members (*on paper, not in SPSS* – & not units; that was my demo)
2. Construct (*& interpret!*) a 95% confidence interval to estimate the average age in the *population*
3. Conduct a hypothesis test (all five steps; α =*.05)* for whether the average age of the *population* is 21.
4. Compare the results of your CI interpretation and hypothesis test.
   - Are they consistent? (Hint: they should be!)