Goals for the Week

Inference For Numerical Data

t-Distribution

One-Sample t-Interval

One-Sample t-Test

Two-Sample Inference

Statistical Toolbox, part 2

STAT 140: Week 10

Scheduling 1

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- Week 10 (11/01-11/05): Numerical Inference
  - Homework 9, Journal 9, due Monday, 11/01
  - Progress Check 10 due Wednesday, 11/03
- Week 11 (11/08-11/12): Practice and Review for Numerical Inference
  - Homework 10, Journal 10, due Monday, 11/08
  - Progress Check 11 due Wednesday, 11/10
  - Test 2 due Sunday, 11/14!!
- Week 12 (11/15-11/20): More on Linear Regression
  - Progress Check 12 due Wednesday, 11/17

Scheduling 2

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- Week 13 (11/21 - 11/25): Multiple Linear Regression
  - Video Lecture rather than in-person–there is NO NEED to be in class on Monday, November 21.
  - Progress Check 12 due Monday, 11/29
  - Homework 12, Journal 12, due Wednesday, 12/01
- Week 14 (11/29 - 12/03)
  - Monday: Practice for Week 13 material
  - Wednesday: Week 13 Catch Up
  - Friday: Project Work Day
- Week 15 (12/06): TBD, but either final review or project work day.

Statistical Toolbox

STAT 140: Week 10

Goals for the Week

Inference For Numerical Data

t-Distribution

One-Sample t-Interval

One-Sample t-Test

Two-Sample Inference

Statistical Toolbox, part 2

Analysis methods we have learned so far...

- One-proportion test, confidence interval

My hope is that after Week 10 (11/01-11/05), you will be able to:

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

My hope is that after this week, you will be able to:

- Understand the differences between the normal distribution and the t-distribution.
- Understand when and how to use a one-sample t-test, including preparing for a test, checking conditions, calculating test-statistics and p-values, and interpreting results.
- Discuss tests for two samples–the two-proportion z-test and the two-sample t-test.
- Analyze data using the tests mentioned above,

Inference for Numerical Data

STAT 140: Week 10

Goals for the Week

Inference For Numerical Data

t-Distribution

One-Sample t-Interval

One-Sample t-Test

Two-Sample Inference

Statistical Toolbox, part 2

- Up until this week, the analysis tools we have learned have been for categorical data or proportions.
- This week, we will focus on analysis tools for numerical data–typically means.
- In both cases, the inference ideas remain the same:
  - Determine which point estimate or test statistic is useful,
  - Identify an appropriate sampling distribution for the point estimate or test statistic, and
  - Apply the ideas of inference.

Central Limit Theorem for $\bar{x}$ 1

When we collect a sufficiently large sample of n independent observations from a population with mean μ and standard deviation σ, the sampling distribution of the sample mean, $\bar{x}$, is nearly normal with

$$\text{mean, or } E(X)\text{: } \mu\text{, and}$$

$$\text{standard error, or } SE\text{: } \frac{\sigma}{\sqrt{n}}.$$

In other words,

$$\bar{x} \sim N(\mu, \frac{\sigma}{\sqrt{n}}).$$

Conditions for Modeling $\bar{x}$

- **Independence**: The sample observations must be independent. The most common way to satisfy this condition is when the sample is a simple random sample from the population. If the data come from a random process, this would also satisfy the independence condition.
- **Normality**
  - When a sample is small, we also require that the sample observations come from a normally distributed population. We can relax this condition more and more for larger sample sizes.
  - What is too small?

Rules of Thumb

- $n \geq 30$: If the sample size is at least 30 and there are no *particularly extreme* outliers, then we typically assume the sampling distribution is $\bar{x}$ is nearly normal, even if the underlying distribution of individual observations is not.
- $n < 30$ If the sample size is less than 30 and there are no clear outliers in the data, then we typically assume the data come from a nearly normal distribution to satisfy the condition.
- There is no *perfect* way to check the normality condition–BUT! You can always **plot the data!!**
  - `geom_histogram()`
  - `geom_density()`

Central Limit Theorem for $\bar{x}$ 2

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- Once we know that $\bar{x} \sim N(\mu, \frac{\sigma}{\sqrt{n}})$, we might think the best strategy is to use a Z-score:

$$z = \frac{\bar{x} - \mu}{\frac{\sigma}{\sqrt{n}}}$$

- Why can't we adopt this strategy?

Introducing the t-distribution 1

- In practice, we cannot directly calculate the standard error for $\bar{x}$ since we do not know the population standard deviation, σ.
- We encountered a similar issue when computing the standard error for a sample proportion ($\sqrt{\frac{p(1-p)}{n}}$), but as a workaround, we substituted $\hat{p}$ for p.
- We can use a similar strategy for computing the standard error of $\bar{x}$, substituting s for σ.

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

Introducing the t-distribution 2

- This strategy tends to work well when we have a lot of data and can estimate σ using s accurately, but it is less precise with smaller samples.
- Because it is less precise with small samples, we encounter problems when trying to use the normal distribution to model $\bar{x}$.
- Instead of using the normal distribution, we will use a new continuous distribution called the t-**distribution**.

Introducing the t-distribution 3

STAT 140: Week 10

Goals for the Week

Inference For Numerical Data

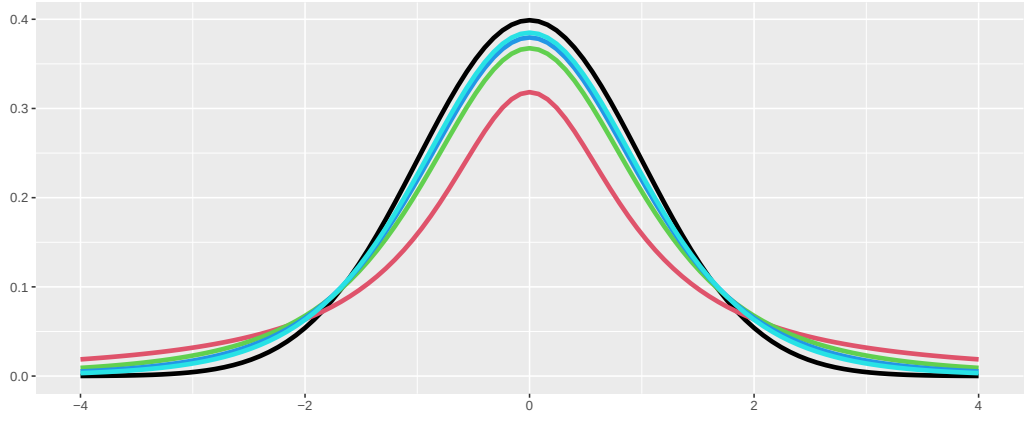
t-Distribution

One-Sample t-Interval

One-Sample t-Test

Two-Sample Inference

Statistical Toolbox, part 2



Introducing the t-distribution 4

- Like the standard normal curve, the t-distribution is always centered at zero.
- Where the normal curve has two parameters, μ and σ, the t-distribution has a single parameter: **degrees of freedom**.
  - The degrees of freedom *df* control the shape of the curve.
- When we increase the degrees of freedom, the t distribution looks closer and closer to a normal curve.
  - *When modeling $\bar{x}$ using the t-distribution, use $df = n - 1$.*

Introducing the t-distribution 5

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- Regardless of the value of df, the tails of any t-distribution are always thicker than the tails of the standard normal distribution.
- This means that observations considered to be unlikely in the normal distribution are somewhat less unlikely in the t-distribution.
- The extra thick tails of the t-distribution are the perfect solution for addressing the extra uncertainty that gets introduced when we use s instead of σ.

Introducing the *t*-distribution 6

You can see that as the degrees of freedom (`df`) increase, the probabilities get closer to the probabilities associated with the normal distribution (`pnorm(2, lower.tail = F) = 0.02275013`).

```
pt(2, df = 10, lower.tail = F)
```

```
## [1] 0.03669402
```

```
pt(2, df = 20, lower.tail = F)
```

```
## [1] 0.02963277
```

```
pt(2, df = 30, lower.tail = F)
```

```
## [1] 0.02731252
```

Introducing the *t*-distribution 7

Also, as the degrees of freedom increase, the percentiles get closer to the percentiles associated with the normal distribution (`qnorm(0.975, mean = 0, sd = 1) = 1.959964`).

```
qt(0.975, df = 10)
```

```
## [1] 2.228139
```

```
qt(0.975, df = 20)
```

```
## [1] 2.085963
```

```
qt(0.975, df = 30)
```

```
## [1] 2.042272
```

One-Sample t-Interval 1

- Recall that confidence intervals represent a range of reasonable values for a parameter.
- So far, we have only built confidence intervals for a proportion (p), but it is easy to see why you would need to build a confidence interval for a mean (μ).
- Proportion:

$$\hat{p} \pm z^* \times \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

- Mean:

$$\bar{x} \pm t^*_{df} \times \frac{s}{\sqrt{n}}$$

One-Sample t-Interval 2

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

1. **Prepare**
   - Identify $\bar{x}$, s, n, and determine what confidence level you wish to use.
2. **Check**
   - Verify the conditions to ensure $\bar{x}$ is nearly normal.
3. **Calculate**
   - If the conditions hold, compute SE, find t^*_{df}, and construct the interval.
4. **Conclude**
   - Interpret the confidence interval in the context of the problem.

Practicing t-intervals 1

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- In April 2020, McDonald's transitioned to a limited menu to simplify tasks for restaurant workers and safely serve customers during the coronavirus public health crisis.
- By transitioning to a limited menu and simplifying tasks, the company hoped that interaction times between customers and staff would be shorter, giving the virus less time to move between hosts.

Practicing t-intervals 2

- Assume that a manager at a local McDonald's needs a 95% confidence interval, has provided data from a random sample of 64 trips. A histogram and summary statistics appear below and on the next slide.
- What are the assumptions for a t-test? Do the data appear to fit these assumptions?

```
(xbar_drivethrough <- mean(waittime$Wait))
```

```
## [1] 236.6991
```

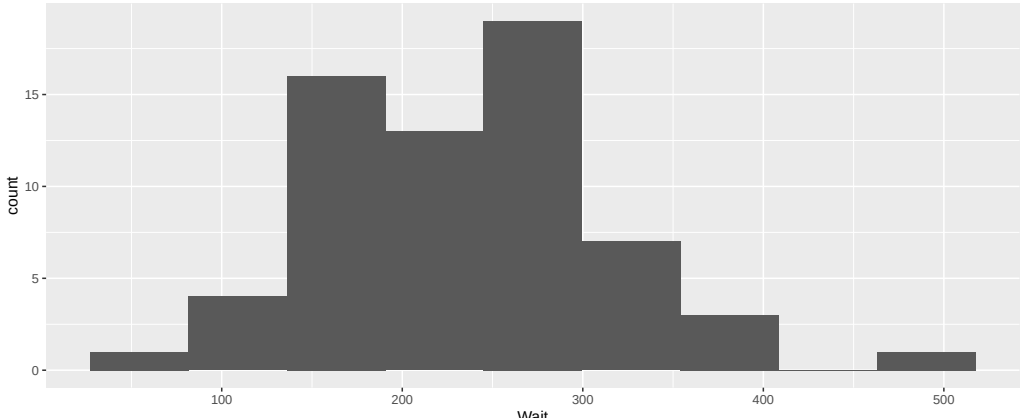
```
(s_drivethrough <- sd(waittime$Wait))
```

```
## [1] 82.27895
```

```
n_drivethrough <- 64
```

Practicing t-intervals 3

- **Check**: Verify the conditions to ensure $\bar{x}$ is nearly normal.
  - There is one potential outlier, but the sample size is pretty large, so we will proceed with caution.



Practicing *t*-intervals 4

- **Calculate**

```
(se_drivethrough <- s_drivethrough/sqrt(n_drivethrough))
```

```
## [1] 10.28487
```

```
(t_star <- qt(0.975, df = 63))
```

```
## [1] 1.998341
```

```
(MoE_drivethrough <- t_star*se_drivethrough)
```

```
## [1] 20.55267
```

Practicing t-intervals 5

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

```
xbar_drivethrough - MoE_drivethrough
```

```
## [1] 216.1465
```

```
xbar_drivethrough + MoE_drivethrough
```

```
## [1] 257.2518
```

- **Interpret**: We are 95% confident that the true mean drive through wait time falls between 216.1 and 257.3 seconds.

Practicing *t*-Intervals 6

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample *t*-Interval
One-Sample *t*-Test
Two-Sample Inference
Statistical Toolbox, part 2

```
t.test(waittime$Wait, conf.level = 0.95)$conf.int
```

```
## [1] 216.1465 257.2518
## attr(,"conf.level")
## [1] 0.95
```

One-Sample t-Test 1

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- We know that you can build both confidence intervals and hypothesis tests for a proportion, p.
- Hopefully, it feels natural to use both confidence intervals and hypothesis tests for a mean, μ, as well!
- Hypotheses:
  - $H_0 : \mu = \mu_0$
  - $H_A : \mu \neq \mu_0$
- Assumptions
  - The sample is randomly selected and therefore independent.
  - The sample is nearly normally distributed.

One-Sample t-Test 2

- Z-score:

$$Z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0 \times (1 - p_0)}{n}}}$$

- t-score:

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

A Z-score tells you how many standard errors an observation is from the mean–the t-score is almost identical! It tells you how far away your observed result is from the proposed mean in terms of sample standard deviations.

One-Sample t-Test 3

- p-value:
  - One-Sided, Lower-Tailed ($H_A : \mu < \mu_0$): `pt(t, df = n - 1)`
  - One-Sided, Upper-Tailed ($H_A : \mu > \mu_0$): `pt(t, df = n - 1, lower.tail = F)`
  - Two-Sided ($H_A : \mu \neq \mu_0$): `2*pt(abs(t), df = n - 1, lower.tail = F)`
- Once you have the p-value, the hypothesis test proceeds according to our previously learned framework.
  - If the p-value is greater than α, fail to reject.
  - If the p-value is less than α, reject.
- Interpretations
  - Fail to Reject: There is not sufficient evidence to say that the true mean is significantly different from the null value.
  - Reject: There is evidence saying that the true mean is significantly different from the null value.

One-Sample t-Test 4

- **Prepare**
  - Identify the parameter of interest, list out hypotheses, identify the significance level, and identify $\bar{x}$, s, n.
- **Check**
  - Verify the conditions to ensure $\bar{x}$ is nearly normal.
- **Calculate**
  - If the conditions hold, compute *SE*, compute the *T-score*, and identify the p-value.
- **Conclude**
  - Evaluate the hypothesis test by comparing the p-value to α, and provide a conclusion in the context of the problem.

Practicing t-tests 1

- The drive through wait time from 2019 was 284 seconds.
- The McDonald's manager noticed that drive through wait times appear to be much lower–is there evidence of a significant decrease in average wait time due to the switch to the limited menu?

1. Prepare

- Write down the hypotheses for a t-test to investigate the manager's hypothesis.
  - What parameter are we investigating?
  - What is the null value?
  - Is the test one- or two-sided?
- State the significance level.

Practicing t-tests 2

$$H_0 : \mu \geq 284$$

$$H_A : \mu < 284$$

- We are investigating the *mean* wait time, denoted with μ, because we want to see if the *average* wait time has decreased.
- The null value we are using is the value reported from 2019, 284 seconds.
- The test is one-sided–we don't care if the average wait time has increased, we are only interested in showing that it has decreased.
- We will use $\alpha = 0.05$.

Practicing t-tests 3

2. Check

- The assumptions for a t-test are the same as the assumptions for the t-interval.
- The sample is randomly selected and therefore independent.
- The sample is nearly normally distributed.

We saw a histogram already that verifies the conditions.

Practicing t-tests 4

3. Calculate

- Given the null value, the mean, sample standard deviation, and sample size, compute the t-score.
- Intuitively, does this t-score seem large or small?

Practicing *t*-tests 5


```
t_drivethrough <- (xbar_drivethrough - 284)/(se_drivethrough)
t_drivethrough
```

```
## [1] -4.599073
```

```
pt(t_drivethrough, df = 63)
```

```
## [1] 1.049771e-05
```

- The *t*-score is -4.60. We know that *t*-scores are close to *Z*-scores, and a *Z*-score of -4.60 would be considered far away from the mean.
- The p-value is 1.05×10^{-5}, or 0.0000105. What decision would we make about the null hypothesis?

Practicing t-tests 6

- Our p-value is less than $\alpha = 0.05$, so we reject the null hypothesis!
- There is significant evidence that the mean drive through wait time at this McDonald's has decreased from the 2019 average.

Practicing *t*-tests 7


```
t.test(waittime$Wait, mu = 284, alternative = "less")
```

```
##
##  One Sample t-test
##
## data:  waittime$Wait
## t = -4.5991, df = 63, p-value = 1.05e-05
## alternative hypothesis: true mean is less than 284
## 95 percent confidence interval:
##      -Inf 253.8687
## sample estimates:
## mean of x
##  236.6991
```

Hypothesis Tests and Confidence Intervals

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

The relationships between a confidence interval and hypothesis test still hold!

- If the null value for a hypothesis test at significance level α is included in the $100 \times (1 - \alpha)$ confidence level, it is considered reasonable and thus we would fail to reject the null hypothesis.
- If the null value for a hypothesis test at significance level α is *NOT* included in the $100 \times (1 - \alpha)$ confidence level, it is *NOT* considered reasonable and thus we would reject the null hypothesis.

Hypothesis Tests and p-values

- Just like in a test for one or two proportions, we can use a one- or two-sided test for a mean.
- Remember, the choice of test (proportions versus means) depends on what type of data you are using.
- The choice of hypothesis depends on what question you want to ask.

Two-Sample Inference

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- Just as we can build intervals or carry out hypothesis tests for single samples, there are methods for comparing the parameters in two different groups!
- This works for either comparing two proportions or two means.
- In *all* cases, the inference ideas remain the same:
  - Determine which point estimate or test statistic is useful,
  - Identify an appropriate sampling distribution for the point estimate or test statistic, and
  - Apply the ideas of inference.

Two-Sample t-Interval 1

1. **Prepare**
   - Identify $\bar{x}_1$, $\bar{x}_2$, s_1, s_2, n_1, n_2, and determine what confidence level you wish to use.
2. **Check**
   - Verify the conditions to ensure $\bar{x}_1 - \bar{x}_2$ is nearly normal.
3. **Calculate**
   - If the conditions hold, compute SE, find t^*_{df}, and construct the interval.
4. **Conclude**
   - Interpret the confidence interval in the context of the problem.

Two-sample t-Interval 2

The formulas for the interval are slightly different.

- One-Sample:

$$\bar{x} \pm t^*_{df} \times \frac{s}{\sqrt{n}}$$

- Two-Sample

$$\bar{x}_1 - \bar{x}_2 \pm t^*_{df} \times \sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}$$

The degrees of freedom calculation is quite tricky! Use the value given to you by software, or the smaller of $n_1 - 1$ and $n_2 - 1$.

Practicing Two-Sample t-Intervals 1

"The U.S. Census Bureau, in collaboration with five federal agencies, is in a unique position to produce data on the social and economic effects of COVID-19 on American households.

The Household Pulse Survey is designed to deploy quickly and efficiently, collecting data to measure household experiences during the Coronavirus (COVID-19) pandemic. Data will be disseminated in near real-time to inform federal and state response and recovery planning."

Practicing Two-Sample t-Intervals 2

- Are there differences in consumer habits in two different cities in the United States during the covid-19 public health crisis?
- Specifically, are there differences in the average amounts of money households spent on prepared food items, including eating out, fast food, carry out, delivered meals, money spent in cafeterias in work or at school, and vending machines during from Week 17 of the HPS collection period (October 14 to October 26) from DFW to Chicago?
- We will focus on the people who are spending greater than \$0 per week.

Practicing Two-Sample t-Intervals 3

1. Prepare:

Let μ_C is the mean log amount of money spent on prepared food items per week in Chicago, and μ_{DFW} is the mean log amount of money spent on prepared food items per week in DFW.

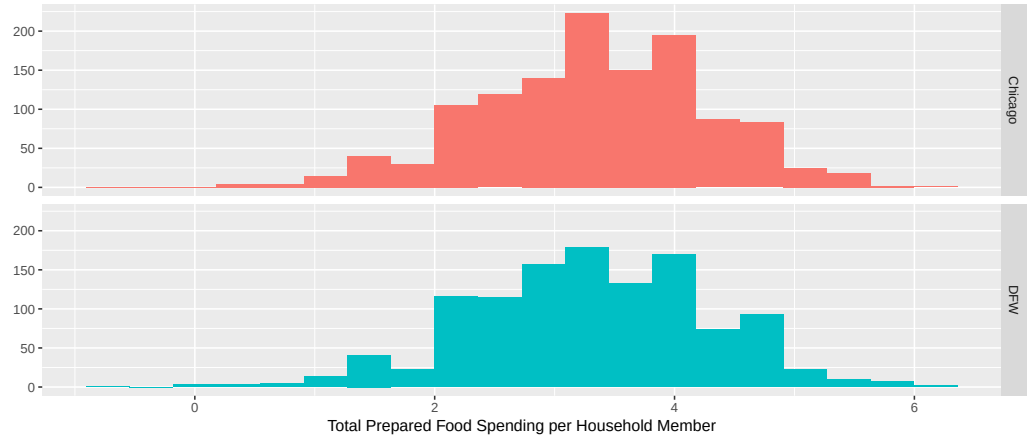
- We will use a confidence level of 99%.
- $\bar{x}_C = \$3.32$, $s_C = \$0.914$, $n_C = 1241$
- $\bar{x}_{DFW} = \$3.30$, $s_{DFW} = \$0.956$, $n_{DFW} = 1165$

2. Check

- The samples were randomly selected and therefore independent (within and between).
- The samples are nearly normally distributed.

Practicing Two-Sample t-Intervals 4

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2



Practicing Two-Sample t-intervals 5


```
t.test(data = prepfood_cities, log(TSPNDPRPD) ~ MSA,
       conf.level = 0.99)$conf.int
```

```
## [1] -0.0715672  0.1253295
## attr(,"conf.level")
## [1] 0.99
```

We are 99% confident that the true mean difference in log amount of money spent on prepared foods per person from Chicago to DFW is between -0.072 and 0.125.

Two-Sample t-Test 1

1. **Prepare**
   - Identify the parameters of interest, list out hypotheses, identify the significance level, and identify $\bar{x}_1,\ \bar{x}_2,\ s_1,\ s_2,\ n_1,\ n_2$.
2. **Check**
   - Verify the conditions to ensure $\bar{x}_1 - \bar{x}_2$ is nearly normal.
3. **Calculate**
   - If the conditions hold, compute SE, compute the *t-score*, and identify the p-value.
4. **Conclude**
   - Evaluate the hypothesis test by comparing the p-value to α, and provide a conclusion in the context of the problem.

Two-Sample t-Test 2

- To test for significant differences between two groups, the hypotheses are written
  - $H_0 : \mu_1 = \mu_2$
  - $H_A : \mu_1 \neq \mu_2$
- Assumptions:
  - Independence—the data are independent within and between groups.
  - Normality—Check outliers and rules of thumb for each group separately.

Two-Sample t-Test 3

The formulas for the tests are slightly different.

- One-Sample

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

- Two-sample

$$t = \frac{(\bar{x}_1 - \bar{x}_2) - 0}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$$

- However, the remainder of the problem is identical to the one-sample t-test!
  - Calculate the p-value, reject or fail to reject as usual, interpret in the context of the problem.

Practicing Two-Sample t-tests 1

We will use the same dataset as the confidence interval. The hypotheses are:

$$H_0 : \mu_C = \mu_{DFW}$$
$$H_A : \mu_C \neq \mu_{DFW}$$

To match with our 99% confidence level, we will use a 1% significance level.

Practicing Two-Sample *t*-tests 2


```
t.test(data = prepfood_cities, log(TSPNDPRPD) ~ MSA,
       conf.level = 0.99)
```

```
##
##  Welch Two Sample t-test
##
## data:  log(TSPNDPRPD) by MSA
## t = 0.70389, df = 2376.4, p-value = 0.4816
## alternative hypothesis: true difference in means between group Chi
## 99 percent confidence interval:
##  -0.0715672  0.1253295
## sample estimates:
## mean in group Chicago     mean in group DFW
##              3.323315              3.296434
```

Practicing Two-Sample t-tests 3

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- The p-value is 0.4816, which is greater than 1%.
- We fail to reject our null hypothesis!
- There is not enough evidence to say that the mean log amounts of money spent on prepared food per person is different in Dallas and Chicago.

Two-Proportion z-Interval 1

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

1. **Prepare**
   - Identify $\hat{p}_1$, $\hat{p}_2$, n_1, n_2, and determine what confidence level you wish to use.
2. **Check**
   - Verify the conditions for $\hat{p}_1$ and $\hat{p}_2$.
3. **Calculate**
   - If the conditions hold, compute SE, find z^*, and construct the interval.
4. **Conclude**
   - Interpret the confidence interval in the context of the problem.

Two-Proportion z-Interval 2

The formulas for the interval are slightly different.

- One-Sample:

$$\hat{p} \pm z^* \times \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

- Two-Sample

$$p_1 - p_2 \pm z^* \times \sqrt{\frac{p_1(1-p_1)}{n1} + \frac{p_2(1-p_2)}{n2}}$$

Practicing Two-Proportion z-Intervals 1

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- Is the proportion of people in Chicago spending any money on prepared foods at all different from the proportion of people in DFW?

Practicing Two-Proportion z-Intervals 2

1. Prepare:

Let p_C be the proportion of people spending money on prepared food items in Chicago, and p_{DFW} be the proportion of people spending money on prepared food items in DFW.

- We will use a confidence level of 99%.
- $x_C = 1,241$, $n_C = 1875$, $\hat{p}_C = 0.662$
- $x_{DFW} = 1,165$, $n_{DFW} = 1,715$, $\hat{p}_{DFW} = 0.679$

Practicing Two-Proportion z-Intervals 3

2. Check

- The samples were randomly selected and therefore independent (within and between).
- The success failure conditions need to be checked for each group separately!

$$n\hat{p}_C = 1,241 > 10, \; n(1-\hat{p}_C) = 634 > 10$$

$$n\hat{p}_{DFW} = 1,165 > 10, \; n(1-\hat{p}_{DFW}) = 550 > 10$$

Practicing Two-Proportion z-intervals 4


```
prop.test(x = c(1241, 1165), n = c(1875, 1715),
        conf.level = 0.99, correct = FALSE)$conf.int
```

```
## [1] -0.05786566  0.02299841
## attr(,"conf.level")
## [1] 0.99
```

We are 99% confident that the true difference in proportions of people spending any money at all in Chicago and DFW is between -5.8% and 2.3%.

Two-Proportion z-Test 1

1. **Prepare**
   - Identify the parameters of interest, list out hypotheses, identify the significance level, and identify $\hat{p}_1$, $\hat{p}_2$, n_1, and n_2.
2. **Check**
   - Verify the conditions for $\hat{p}_1$ and $\hat{p}_2$.
3. **Calculate**
   - If the conditions hold, compute SE, compute the z-*score*, and identify the p-value.
4. **Conclude**
   - Evaluate the hypothesis test by comparing the p-value to α, and provide a conclusion in the context of the problem.

Two-Proportion z-Test 2

Goals for the Week
Inference For Numerical Data
t-Distribution
One-Sample t-Interval
One-Sample t-Test
Two-Sample Inference
Statistical Toolbox, part 2

- To test for significant differences between two groups, the hypotheses are written
  - $H_0 : p_1 = p_2$
  - $H_A : p_1 \neq p_2$
- Assumptions:
  - Independence—the data are independent within and between groups.
  - Normality—Check conditions using $\hat{p}_p ooled$.

$$\hat{p}_{pooled} = \frac{x_1 + x_2}{n_1 + n_2}$$

We do this because we assume that the proportions are equal in the null, so we can use information from both groups to estimate the common proportion.

Two-Proportion z-Test 3

The formulas for the tests are slightly different.

- One-Sample

$$z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0 \times (1 - p_0)}{n}}}$$

- Two-sample

$$z = \frac{(\hat{p}_1 - \hat{p}_2) - 0}{\sqrt{\frac{\hat{p}_{pooled} \times (1 - \hat{p}_{pooled})}{n_1} + \frac{\hat{p}_{pooled} \times (1 - \hat{p}_{pooled})}{n_2}}}$$

- The remainder of the problem is identical to the one-proportion z-test!

Practicing Two-Proportion z-tests 1

We will use the same dataset as the confidence interval. The hypotheses are:

$$H_0 : p_C = p_{DFW}$$
$$H_A : p_C \neq p_{DFW}$$

To match with our 99% confidence level, we will use a 1% significance level.

Practicing Two-Proportion z-tests 2


```
prop.test(x = c(1241, 1165), n = c(1875, 1715),
          conf.level = 0.99, correct = FALSE)
```

```
##
##  2-sample test for equality of proportions without continuity
##  correction
##
## data:  c(1241, 1165) out of c(1875, 1715)
## X-squared = 1.2317, df = 1, p-value = 0.2671
## alternative hypothesis: two.sided
## 99 percent confidence interval:
##  -0.05786566  0.02299841
## sample estimates:
##    prop 1    prop 2
## 0.6618667 0.6793003
```

Practicing Two-Sample t-tests 3

- The p-value is 0.2671, which is greater than 1%.
- We fail to reject our null hypothesis!
- There is not enough evidence to say that proportions of people spending any money on prepared foods is different in Dallas and Chicago.

Statistical Toolbox

Analysis methods we have learned so far...

- One-proportion z-test, confidence interval
- One-sample t-test, confidence interval
- Two-sample t-test, confidence interval
- Two-proportion z-test, confidence interval