Can Vitamin D Help Prevent COVID-19?

Overview and student objectives

In-Class Activity Length: 25 minutes

Overview

This in-class activity revisits hypothesis testing for a difference in proportions, but rather than calculating a P-value using a traditional two-sample hypothesis test, students use a tool to simulate a null distribution of sample differences in proportions. Then they compare the observed difference in proportions to this distribution to obtain an approximate P-value. This procedure is called a randomization test. Technically, this test procedure is only called a randomization test when the data arise from a randomized experiment, since the simulation mimics the random assignment in the study. When data arise from an observational study, the procedure is called a permutation test.

Objectives

Students will understand:

- Hypothesis tests can be conducted by simulating sample data under the assumption of a null hypothesis and then comparing observed data to the results of the simulation.
- One can estimate a P-value by calculating the proportion of simulated statistics (under the assumption of the null hypothesis) that is as or more extreme than the one observed.

Students will be able to:

- Conduct a randomization test involving a difference in proportions.

Suggested resources and preparation

Materials and technology

- Computer, projector, document camera
- Student Pages for In-Class Activity
- Practice Assignment
- Access to the *DCMP Data Analysis Tools*

Prerequisite assumptions

Students should be able to:

- Identify observational units and types of variables.
- Describe how to simulate the random assignment of treatments to experimental groups.
- Determine whether data provide evidence against a null hypothesis based on a simulated null distribution of sample statistics.

Making connections

This activity connects back to two-sample tests for proportions and simulation-based inference.

Background context

The activity uses data from a clinical trial that examines if taking vitamin D can reduce the severity of the COVID-19 illness among patients who have been diagnosed with COVID-19. The instructor should glance through the news article and published study article that are referenced in the activity:

- Curley, B. (2020, September 13). *Vitamin D can help reduce COVID-19 risks: Here's how*. Healthline. https://www.healthline.com/health-news/vitamin-d-can-help-reduce-covid19-risks
- Castillo, M. E., Costa, L. M. E., Barrios, J. M. V., Díaz, J. F. A., Miranda, J. L., Bouillon, R., & Gomez, J. M. Q. (2020, October). Effect of calcifediol treatment and best available therapy versus best available therapy on intensive care unit admission and mortality among patients hospitalized for COVID-19: A pilot randomized clinical study. *Journal of Steroid Biochemistry and Molecular Biology*, *203*. https://doi.org/10.1016/j.jsbmb.2020.105751

Suggested instructional plan

Frame the activity (5 minutes)

Resources and Structure	Instructor Suggestions
Think-Pair-Share *Brief Discussion*	Question 1 • Have students answer Question 1 prior to describing the study referred to in the opening paragraph. • After students have had a chance to answer Question 1 on their own or in groups, describe how the study was carried out and have a class discussion on what study design characteristics in this study match what students suggested in their answers to Question 1. • Students should have access to the *DCMP Association Between Two Categorical Variables* tool at https://dcmpdatatools.utdanacenter.org/association_categorical/. • Students could either use the tool individually or share a computer to access the tool within groups. If students do not have access to the tool, Questions 12 and 13 should be conducted together as a class, with the instructor demonstrating the tool on the projector. • Transition to the in-class activity by briefly discussing the **Objectives for the activity**.

Activity flow (17 minutes)

Resources and Structure	Instructor Suggestions
Small Groups	• Have students work in small groups to complete this activity. Questions 2–4 • Assist students with Questions 2 and 3 as needed. • In Question 4, have students refer back to Question 3 (the two groups) and ask how group membership was assigned—randomly or naturally occurring?
Guiding Question	• "How would the study be different if it were an observational study?" [Answer: Patients would have decided on their own to take vitamin D or not.]
	Questions 5–14 • Assist students with Questions 5–11 as needed.
Direct Instruction	• The description after Question 11 is very important. The entire premise of the randomization test is that if values of the variable of interest would have occurred regardless of the group, we can label cards with the values of that variable and then shuffle and re-randomize the cards into the two groups. It may help to demonstrate how one could do this using cards through the following questions: ○ "What would each card represent?" (Answer: One patient in the study) ○ "How many cards would we need to simulate the random assignment in this study?" (Answer: 76) ○ "What would we write on each card?" (Answer: Whether the patient was admitted to the ICU or not) ○ "What would we do with the cards to replicate the random assignment process?" (Answer: Shuffle the cards and then deal them into a "calcifediol" group of 50 cards and an "untreated" group of 26 cards. Then calculate the proportion of ICU admissions that ended up in each group.) • Question 12, Part D asks students to combine their simulated results with their classmates. This will require the instructor to facilitate drawing a dotplot of these results on the board. Draw the x-axis so that it ranges from -0.50 to 0.50 and label the x-axis "Difference in proportions admitted to the ICU (calcifediol – untreated)." Then have students come up to the board and draw their simulated difference in proportions on the plot. • If running low on time, the instructor can demonstrate the use of the tool required in Question 13 and then answer the parts of this question together as a class.

Extension	• If time permits, compare the results of this test to the results if we had used a two-sample z-test for a difference in proportions. Ask if conditions would have been met to use the theory-based test (Answer: No).

Wrap-up/transition (3 minutes)

Resources and Structure	Instructor Suggestions
Wrap-up	• Ask students to generalize the process of what they just did, as simulation-based hypothesis tests all follow the same general steps: 1. Set up the null and alternative hypotheses based on the research question. 2. Simulate a large number of samples (usually 1,000 or more) under the assumption of the null hypothesis, calculating a sample statistic for each simulated sample. 3. Plot the simulated sample statistics with a histogram and compare the original observed statistic to the plot. 4. The proportion of simulated statistics as or more extreme than observed is the estimated P-value. • The method described in this activity provides an alternative to the two-sample z-test in situations where the conditions are not met. Ask students whether the conditions for a two-sample z-test would have been met for this study. (The conditions are not met since there is only a single admission in the treated group.) • Give students a preview of Practice Assignment. In this activity, students conducted a randomization test for a difference in two proportions. Students will conduct a randomization test for a difference in two means in Practice Assignment. The process is identical: to carry out a single randomization of the data, write the values of the variable of interest on cards, shuffle the cards, and then deal the cards into two groups. For a difference in proportions, the variable of interest is categorical (the cards contain category names); for a difference in means, the variable of interest is quantitative (the cards contain numerical values). • Have students refer back to the **Objectives for the activity** and check the ones they recognize. Alternatively, they may check the objectives throughout the activity.

Suggested assessment, assignments, and reflections

- Give Practice Assignment

Can Vitamin D Help Prevent COVID-19?

During the COVID-19 pandemic, clinical trials were conducted to determine which factors may help reduce COVID-19 risk. One vitamin in particular, vitamin D, looked promising for reducing this risk.

News headlines such as "Vitamin D Can Help Reduce COVID-19 Risks: Here's How"[1] started appearing in Fall 2020. But were these headlines valid? In this in-class activity, we will examine data from one study designed to determine if vitamin D may reduce the severity of illness if one does test positive for COVID-19.[2]

Credit: iStock/tungphoto

1) How would you design a study to determine if vitamin D reduces the risk of intensive care unit (ICU) admission due to COVID-19?

Objectives for the activity

You will understand:

☐ Hypothesis tests can be conducted by simulating sample data under the assumption of a null hypothesis and then comparing observed data to the results of the simulation.

☐ One can estimate a P-value by calculating the proportion of simulated statistics (under the assumption of the null hypothesis) that is as or more extreme than the one observed.

You will be able to:

☐ Conduct a randomization test involving a difference in proportions.

In a small, randomized clinical trial conducted in Spain, 76 patients hospitalized with COVID-19 were randomized to either receive a calcifediol (vitamin D) treatment or not. All the patients were treated with the standard intervention for COVID-19, which at the time was a combination of hydroxychloroquine and azithromycin. Researchers recorded whether each patient was admitted to the ICU or not.

[1] Curley, B. (2020, September 13). *Vitamin D can help reduce COVID-19 risks: Here's how*. Healthline. https://www.healthline.com/health-news/vitamin-d-can-help-reduce-covid19-risks

[2] Castillo, M. E., Costa, L. M. E., Barrios, J. M. V., Díaz, J. F. A., Miranda, J. L., Bouillon, R., & Gomez, J. M. Q. (2020, October). Effect of calcifediol treatment and best available therapy versus best available therapy on intensive care unit admission and mortality among patients hospitalized for COVID-19: A pilot randomized clinical study. *Journal of Steroid Biochemistry and Molecular Biology*, *203*. https://doi.org/10.1016/j.jsbmb.2020.105751

2) What variable distinguishes the two groups compared in this study?
   a) Whether the patient was hospitalized with COVID-19 or not
   b) Whether the patient received the calcifediol treatment or not
   c) Whether the patient was admitted to the ICU or not
   d) Whether the patient was treated with the standard intervention or not

3) What is the variable of interest in this study?
   a) Whether the patient was hospitalized with COVID-19 or not
   b) Whether the patient received the calcifediol treatment or not
   c) Whether the patient was admitted to the ICU or not
   d) Whether the patient was treated with the standard intervention or not

4) Is this study a randomized experiment or an observational study?
   a) Randomized experiment, since patients were randomly assigned to receive calcifediol or not
   b) Randomized experiment, since there were two different treatments
   c) Observational study, since whether the patient was admitted to the ICU was not randomly assigned
   d) Observational study, since patients were not randomly selected from the population of all COVID-19 infected individuals

5) Before the data are collected, you should state the research hypothesis. What were the researchers hoping to show in this study?

6) Let p_{Calc} be the probability of admission to the ICU for COVID-19 patients on the calcifediol treatment and p_{Untr} be the probability of admission to the ICU for COVID-19 patients not treated with calcifediol. State the null and alternative hypotheses in terms of these parameters.

7) Of the 50 patients randomly assigned to the calcifediol treatment, only one was admitted to the ICU, whereas 13 of the 26 patients who did not receive the calcifediol treatment were admitted to the ICU. Organize these results into the following 2×2 table.

	Treated with calcifediol	**Untreated with calcifediol**	**Total**
Admitted to ICU			
Not admitted to ICU			
Total			

8) What proportion of the 76 patients in the study were admitted to the ICU?

9) What proportion of the patients on the calcifediol treatment were admitted to the ICU? What proportion of the untreated patients were admitted to the ICU? What is the appropriate statistical notation for each of these values?

10) Calculate the difference in proportions of ICU admissions between the two groups (use *calcifediol-treated – calcifediol-untreated* as the order of subtraction).

The following is a segmented bar graph of the data.

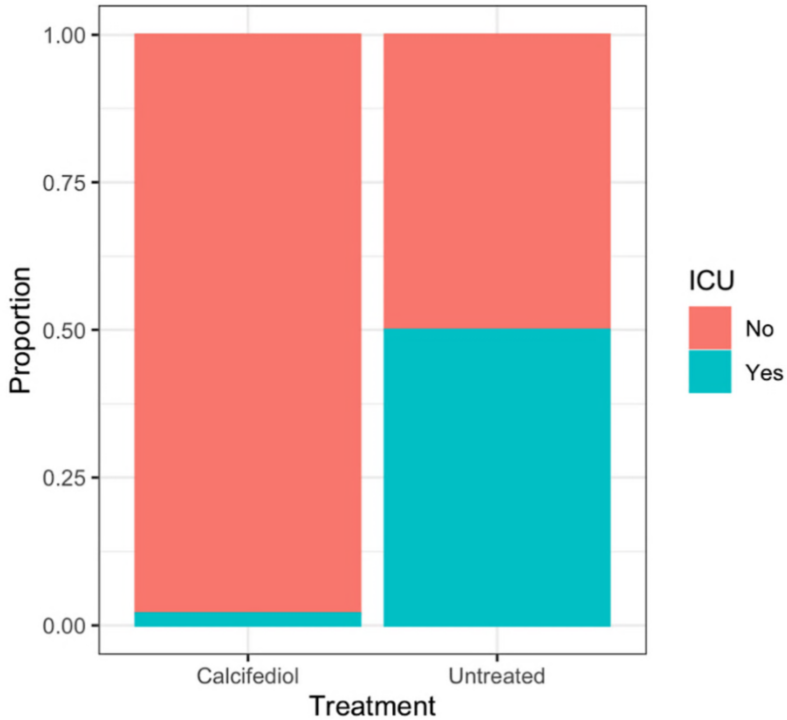

11) If the calcifediol treatment had no effect on the severity of COVID-19 symptoms, how would you expect this graph to appear? Explain.

If the calcifediol treatment had no effect on the severity of COVID-19 symptoms, each patient would have either been admitted to the ICU or not regardless of the treatment to which they were assigned. In other words, the 13 patients admitted to the ICU who were in the group that did not receive calcifediol would have ended up being admitted to the ICU even if they had received calcifediol. While it is *possible* that calcifediol does not help lower the risk of admission to the ICU for COVID-19 patients and the researchers were unlucky and just happened to "draw" more of the subjects who were going to be admitted to the ICU into the untreated group, we would like to determine whether this outcome is *probable*. If 14 out of the 76 patients were going to be admitted to the ICU no matter what, we would have expected around the same proportion of those patients to end up in each group. The key question is how unlikely the observed difference in proportions of patients admitted to the ICU is by the random assignment process alone.

We will answer this question by replicating the random assignment process all over again, under the assumption that calcifediol does *not* decrease the risk of ICU admission. We'll start with 14 ICU admissions and 62 ICU non-admissions and then randomly assign 50 of these 76 subjects to the calcifediol-treated group and the other 26 to the calcifediol-untreated group.

12) Go to the *DCMP Association Between Two Categorical Variables* tool at https://dcmpdatatools.utdanacenter.org/association_categorical/.

- Under "Data Entry & Descriptive Statistics:"
  - Select "Contingency Table" under "Enter Data."
  - Type "ICU" for the row variable, with "Admitted" and "Not admitted" for the category labels.
  - Type "Group" for the column variable, with "Treated" and "Untreated" as the category labels.
  - Enter the contingency table completed in Question 7.
- Now select "Permutation Distribution" in the top right. You should see the contingency table you entered as the "Observed Contingency Table." Check "ICU" under "Permutate Labels of" and then generate a single permutation of the data.

Part A: The shuffled counts in each group are shown under "Dataset from last permutation" on the right. Copy these results into the following 2×2 table.

	Treated with calcifediol	**Untreated with calcifediol**	**Total**
Admitted to ICU			
Not admitted to ICU			
Total			

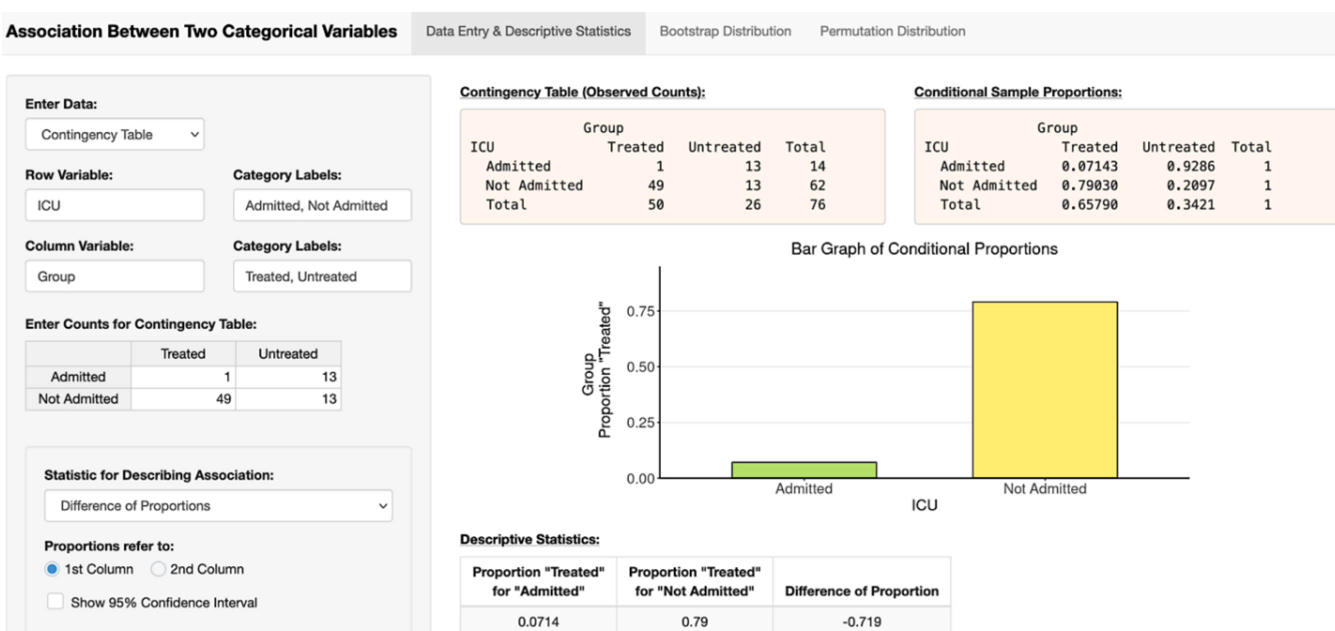

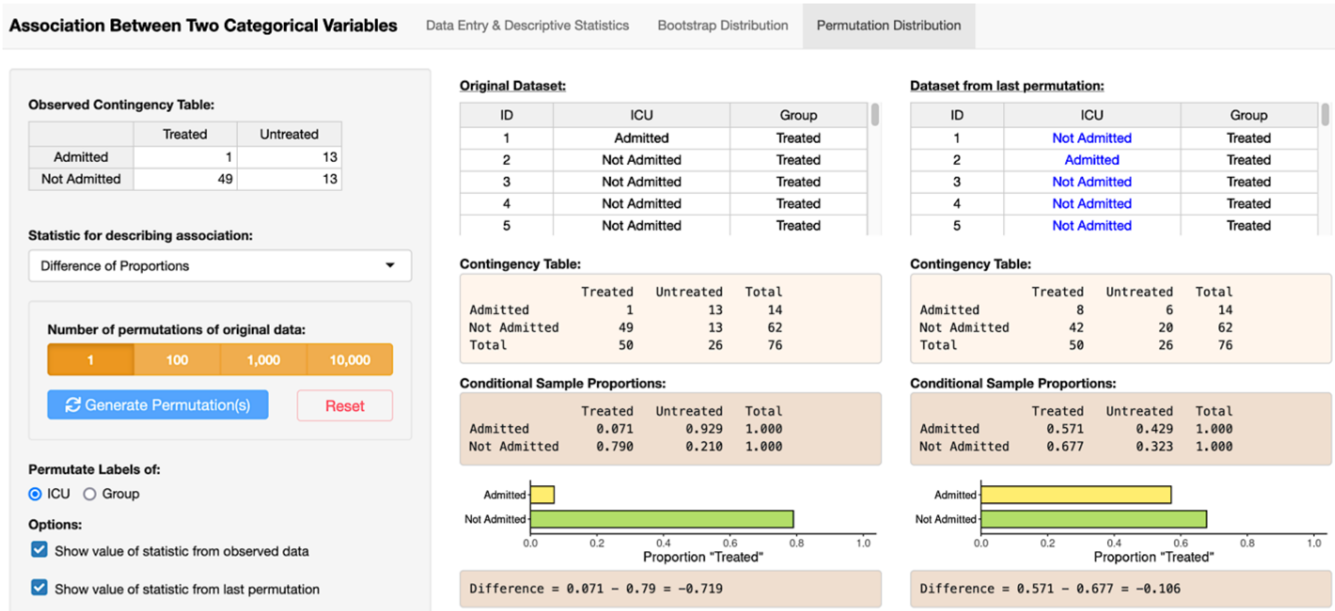

Part B: Calculate the shuffled difference in proportions of ICU admissions between the two groups (use *calcifediol-treated – calcifediol-untreated* as the order of subtraction) in your simulated table from Part A.

Hint: This should match the difference in proportions shown in the tool under "Dataset from last permutation."

Part C: Is the result of this simulated random assignment as "extreme" as the actual results that the researchers obtained? That is, did one or fewer of the ICU admissions end up in the calcifediol-treated group?

Part D: Combine your results with those from your classmates, producing a well-labeled dotplot. In what proportion of the simulated-random assignments were one or fewer of the ICU admissions assigned to the calcifediol-treated group?

13) We can use the tool to simulate 1,000s of random assignments. Click "Reset," select "1,000" for "Number of permutations of the original data," and then click "Generate Permutation(s)."

Part A: Examine the histogram of shuffled differences in proportions. Where is this plot centered? Explain.

Part B: Based on the histogram of shuffled differences in proportions, does it seem like the actual experimental result (only one ICU admission in the calcifediol-treated group) would be surprising to arise solely from the random assignment process under the assumption that calcifediol has no effect on the severity of COVID-19 symptoms? Explain.

Part C: Select "Show statistical summary of permutation distribution" under "Options." The tool will display the number and percentage of shuffled differences in proportions that are equal to or less than the one observed in the actual data at the bottom of the page. What percentage of the simulated randomizations resulted in a difference in proportions less than or equal to the one observed in the actual data?

You have just conducted what is called a **randomization test** (also sometimes called a permutation test) of the hypotheses you stated in Question 6, and your answer to Question 13, Part C is an approximate P-value for this test. A P-value is the probability of obtaining the actual results or something more extreme, under the assumption of the null hypothesis. This can be approximated by simulating many, many randomizations under the null hypothesis and calculating the proportion of randomizations that produce results like ours—or something more extreme.

14) Consider your approximate P-value from Question 13, Part C.

Part A: Is this P-value small enough so that you would consider the actual outcome surprising (or more extreme) under the null model that calcifediol has no effect on the severity of COVID-19 symptoms? Explain.

Part B: Would you say that the researchers obtained strong evidence that the risk of ICU admission decreases when treated with calcifediol? Explain your reasoning based on your simulation results, including a discussion of the purpose of the simulation process and what information it revealed to help you answer this research question.

Part C: Are you willing to draw a cause-and-effect conclusion about calcifediol-treatment and ICU admission based on these results? Justify your answer based on the design of the study.

Part D: Are you willing to generalize these conclusions to all COVID-19 patients? Justify your answer based on the design of the study.

Practice Assignment

It is common practice in the United States for dining parties to tip their servers after a meal at a restaurant. When the server brings the bill to the table, often the server includes an unexpected gift in the form of candy. But does including candy with the bill increase the server's tip?

A team of researchers set out to answer this question. Participants in the study were 92 dining parties at a restaurant in downtown Ithaca, New York. Two seasoned male servers participated. Just prior to delivering a check, the server selected a playing card from a shuffled deck of cards. If the card was red, the server gave each person in the dining party a fancy, foil-wrapped piece of chocolate when they delivered the check. When the card was black, the server delivered the check without the chocolate. Upon delivery of the check, the server thanked the customers and recorded the tip percentage (the amount of the tip divided by the size of the bill before taxes and multiplied by 100). It was hypothesized that the servers would receive larger tip percentages, on average, from those parties who received candy.[3]

1) What are the observational units in this dataset?
   a) Servers
   b) Restaurants
   c) Candy
   d) Dining parties

2) Is this a randomized experiment or an observational study? Explain in context of the study.
   Hint: Was the type of check (with or without candy) randomly assigned?

[3] Strohmetz, D. B., Rind, B., Fisher, R., & Lynn, M. (2002). Sweetening the till: The use of candy to increase restaurant tipping. *Journal of Applied Social Psychology*, 32(2), 300–309. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1559-1816.2002.tb00216.x

3) We will summarize the data both graphically and numerically. Go to the *DCMP Exploring Quantitative Data* tool at https://dcmpdatatools.utdanacenter.org/eda_quantitative/.

Enter the following inputs.

- Click on the **Several Groups** tab at the top.
- Enter Data: Your Own
- # of Groups: 2
- Variable Name: Tip Percentage
- Group Name: Treatment
- Group Labels: Candy, No candy
- Candy: Copy and paste the tip percentages from the group of parties who received candy with their bills. (The data are given in the following table, continued on the next page.)
- No candy: Copy and paste the tip percentages from the group of parties who did not receive candy with their bills. (The data are given in the following table, continued on the next page.)
- Choose Type of Plot: Dotplot

Tip percentage Candy	Tip percentage No candy
15.12	15.79
13.58	12.85
16.25	13.88
15.29	13.21
15.58	13.07
15.58	15
13.04	13.61
17.07	13.64
14.29	12.38
18.38	14.33
15.16	14.99
20.46	15.03
21.99	13.2
15.61	15.45
15.01	14.72
19.18	16.64
15.84	14.12
18.49	14.23
23.81	14.61
16.47	14.51
15.84	13.42
19.01	13.79
16.6	17.52

Tip percentage Candy	Tip percentage No candy
17.19	14.42
19.82	14.2
16.69	9.89
17.87	16.79
15.17	16.57
17.48	16.03
16.02	17.36
15.35	13.7
14.85	14.78
20.82	14.99
16.1	17.27
19.17	15.16
29.21	18.09
18.21	16.46
20.72	17.88
19.38	16.6
19.86	15.7
17.67	11.66
20.46	13.78
22.8	19.22
18.1	16.47
22.77	17.52
17.32	18.15

Part A: What is the mean tip percentage for the Candy group? For the No candy group?

Part B: Examine the dotplots of the data. (You may want to increase the dot size in order to visualize the data more clearly.) Based on these plots, write a few sentences comparing the distributions of tip percentages in each group. Be sure to address the center, shape, variability, and outliers.

Part C: Based on these plots alone, does it appear that including candy with the bill increases a server's tip percentage? Explain.

As in any hypothesis test, there are two possibilities:

A. Whether the dining party received candy or not has no effect on the tip percentage, and the difference in mean tip percentages seen in the data was due to random chance.

B. Dining parties that receive candy tend to leave higher tip percentages, and the difference in mean tip percentages seen in the data reflects this effect.

4) Which of the two previous possibilities is the null hypothesis for this study? Which is the alternative hypothesis?

To determine which of the two previous possibilities is more plausible, we will simulate data from the study under the assumption that candy has no effect on tip percentage and examine the distribution of the difference in sample means across thousands of simulated samples. Then, we can compare the observed difference in sample means to this simulated distribution to determine how often the observed data would be expected to occur under the assumption that candy has not effect.

To simulate one sample, imagine a deck of 92 cards. Write the tip percentage of each dining party on the cards. Then, shuffle the cards and deal the cards into two groups of 46 cards. One of these groups represents the dining parties who received candy with their bills and the other represents the dining parties who did not receive candy with their bills. Since the cards were dealt to the two groups randomly, this simulation assumes that whether the dining party was given candy had no effect on the tip percentage.

5) Go to https://istats.shinyapps.io/PermDist_2samples/. Enter the following inputs.

- Enter Data: Provide Own
- Name of Response Variable: Tip Percentage
- Group 1 Label: Candy
- Group 2 Label: No candy
- Group 1 Data: Copy and paste the tip percentages from the group of parties who received candy with their bills. (The data were given in Question 3.)
- Group 2 Data: Copy and paste the tip percentages from the group of parties who did not receive candy with their bills. (The data were given in Question 3.)

Part A: What is the observed difference in sample mean tip percentages between the two groups (using *Candy group – No candy group* as the order of subtraction)?

Part B: Enter "1" for "Select how many permutations you want to generate" and then click "Generate Permutation(s)." This will perform the process of dealing hypothetical "cards," as described previously. The shuffled, or "permuted," sample is shown in the plot on the left, where the colors represent the original group to which the dining parties belonged.

What is the value of the shuffled difference in means? Is this simulated difference in mean tip percentages as "extreme" as the actual results shown in the original sample? Explain.

Part C: Click "Reset." Enter "1,000" for "Select how many permutations you want to generate" and then click "Generate Permutation(s)." Take a screenshot or sketch the distribution of simulated differences in sample means.

Part D: Where is the distribution of simulated differences in mean tip percentages centered? Why does this make sense?

Part E: Click the **Permutation Test** tab at the top of the page. Select "Greater" for the alternative hypothesis. Find the proportion of simulated differences in means that is greater than or equal to the observed difference you found in Part A (using *Candy – No candy* as the order of subtraction).

Hint: The desired proportion is the *P*-value of the permutation test.

6) Consider your results from the simulation in Question 5.

Part A: Would you say that the results that the researchers obtained provide strong evidence that providing candy with the bill increases tip percentage? Explain your reasoning based on your simulation results, including a discussion of the purpose of the simulation process and what information it revealed to help you answer this research question.

Part B: Are you willing to draw a cause-and-effect conclusion about whether a dining party receives candy and tip percentage based on these results? Justify your answer based on the design of the study.

a) Yes, since the 92 dining parties were a random sample from a larger population.

b) Yes, since whether a dining party received candy was randomly assigned.

c) Yes, since the results were statistically significant.

d) No, since this was an observational study.

Part C: Are you willing to generalize these conclusions to all dining parties in all restaurants in Ithaca, New York? Justify your answer based on the design of the study.

a) Yes, since the 92 dining parties were a random sample from all restaurants in Ithaca, New York.

b) Yes, since whether a dining party received candy was randomly assigned.

c) No, since the 92 dining parties were taken from only one restaurant in Ithaca, New York; they cannot be assumed to be representative of the population of all restaurants in the city.

d) No, since this was an observational study.