

Using Dynamic Applets to Deepen Students' Statistical Understanding

2023 AP Annual Conference (Seattle, WA)

Daren Starnes (Grand Lake, CO)

Sampling distribution of a sample proportion

Suppose that 83% of students at a large high school regularly use a backpack to carry their books from class to class. Imagine taking a simple random sample (SRS) of 100 students and calculating \hat{p} = the proportion of students in the sample who regularly use a backpack. Let's use an applet to investigate the sample-to-sample variability in the values of the sample proportion \hat{p} .

1. Launch the *Simulating Sampling Distributions* applet at www.stapplet.com.
2. In the Population box, select a Population distribution that's Categorical. Enter 0.83 as the true proportion of successes.
3. In the Sample box, change the sample size to $n = 100$. Click on Select sample to simulate taking one random sample from the population. Notice that three distributions are now displayed by the applet.
 - Population distribution
 - What variable is being recorded for each individual (student) in the population? Is this variable categorical or quantitative?

 - Describe the distribution of this variable in the population.
 - Distribution of sample data
 - What variable is being recorded for each individual (student) in the sample? Is this variable categorical or quantitative?

 - Describe the distribution of this variable in the sample.
 - Simulated sampling distribution of the sample proportion \hat{p}
 - What does the graph look like after you have simulated taking one random sample?

 - In the Sampling Distribution box, click on Select samples 9 times so that you have 10 simulated samples of size $n = 100$ from a population with $p = 0.83$. Now make an educated guess: Describe how you think the distribution will look (shape, center, and variability) if we take many, many random samples of size 100 from the population.

4. Use the applet to simulate taking at least 1,000 random samples of 100 students from the population in which 83% regularly use a backpack. How do the shape, center, and variability of the distribution compare to what you guessed in Question 3?

5. What is the largest distance that we can expect the value of the sample proportion \hat{p} to be from the population proportion p ?

- In about 95% of all possible samples of size $n = 100$ from this population, the sample proportion \hat{p} will be within _____ standard deviations $\sigma_{\hat{p}}$ of the population proportion p .
- In about 95% of all possible samples of size $n = 100$ from this population, the population proportion p will be within _____ standard deviations $\sigma_{\hat{p}}$ of the sample proportion \hat{p} .

6. How would the simulated sampling distribution of the sample proportion \hat{p} change if we randomly selected 25 students each time instead of 100? Make an educated guess about the shape, center, and variability of the simulated sampling distribution.

- Click on Clear samples and change the sample size to $n = 25$ in the Sample box.
- Take at least 1,000 samples of 25 students. How do the shape, center, and variability of the approximate sampling distribution compare to the one in Question 4?

7. How would the simulated sampling distribution of the sample proportion \hat{p} change if we randomly selected 25 students at a time from another school where 60% of students regularly use a backpack? Make an educated guess about the shape, center, and variability of the simulated sampling distribution.

- Click on Clear samples. Change the true proportion of successes to 0.60 in the Population box.
- Take at least 1,000 samples of 25 students. How do the shape, center, and variability of the approximate sampling distribution compare to the one in Question 4?

8. For what combinations of n and p is the sampling distribution of \hat{p} approximately normal? Use the applet to investigate.

9. Summarize what you learned in this activity about the sampling distribution of a sample proportion.

What is a confidence interval?

1. Go to www.stapplet.com and launch the *Simulating Confidence Intervals* applet.
 - For the Population distribution, choose Categorical. Set the population proportion to 0.83.
 - Change the sample size to $n = 100$ and choose to compute a 95% confidence interval for p .

2. Click Go!

What just happened?

Confidence interval you obtained:

How was this interval calculated?

Based on the interval you obtained, is it plausible (believable) that the population proportion of students at this school who regularly use a backpack is:

0.60?

0.80?

0.90?

0.83?

What is a confidence level?

- In the Confidence Intervals box, click Go! 9 more times.

What just happened?

What proportion of your 10 confidence intervals contained (captured) $p = 0.83$?

What does a green interval represent?

What does a red interval represent?

- Change the number of intervals to 50, and then click Go!
- Make an educated guess: If you keep clicking Go!, what percentage of the confidence intervals will capture $p = 0.83$?
- Use the applet to determine whether your guess is correct.

- Click on Reset everything.
- In the Sample box, change the confidence level to 99% and click Go!

How does the width of a 99% confidence interval compare to the width of a 95% confidence interval? Explain why that makes sense.

- Make an educated guess: If you keep clicking Go!, what percentage of the confidence intervals will capture $p = 0.83$?
- Use the applet to determine whether your guess is correct.


Summarize: What does the confidence level tell us?

Does encouragement increase athletic performance? *What is a P-value?*

A researcher wondered if people would perform better in the barbell curl if they were encouraged by a coach. The researcher recruited 31 subjects to participate in an experiment and then split them into two groups using random assignment. Group 1 received positive encouragement during the exercise, while Group 2 served as a control group that received no encouragement. The researcher recorded the number of barbell curls each participant was able to complete before setting the bar down. Here are the data:

No encouragement	43	19	33	1	0	2	6	61	21	11	40	15	15	18	37	22
With encouragement	50	11	60	5	40	45	11	30	8	50	11	15	9	40	75	

Parallel dotplots and boxplots of the data are shown below, along with summary statistics.



Summary statistics:

	n	Mean	SD
None	15	21.47	17.84
Enco	15	30.67	22.36

Do these data provide *some evidence* that people like the ones in this study who receive encouragement can do more barbell curls than those who receive no encouragement? Explain your answer.

BIG QUESTION: Is it plausible (believable) that there's really no effect of encouragement on number of barbell curls, and that random chance alone produced the observed difference between these two groups?

Simulation of the random assignment

- If there is no treatment effect, then the data values will be the same as in the original study. How large a difference in group means will there be with different random assignments?
- Write each of the 31 data values on a separate card.
- Mix the cards well, then deal one group of 16 cards (No encouragement) and one group of 15 cards (Encouragement).
- Calculate the difference (Encouragement – None) in mean number of barbell curls for the two groups.
- Make a class dotplot of the simulated difference in means.
- What proportion of the class's simulation trials yielded a difference in means greater than or equal to the one in the original experiment: $30.67 - 21.47 = 9.20$?

Computer simulation:

1. Go to www.rossmanchance.com/applets/index2021.html. Choose Statistical Inference→Randomization test for quantitative response→Two means.
2. At the top left of the applet window, check the box to Enter data Unstacked and then click Clear.
3. Open the Encouragement Data [file](#). Copy the entire data set and paste into the applet. Then click Use Data.
4. Click the box to Show Shuffle Options.
5. Set Number of Shuffles to 1 and Hypothesized $\mu_2 - \mu_1$ to 0.
6. Select display as Plot rather than Data.
7. Click the Shuffle Responses button. What just happened?

8. Click the Shuffle Responses button 9 more times, so that you have a total of 10 simulated differences in sample means. Make a guess about what the distribution of shuffled differences in means will look like if you keep doing more shuffles.

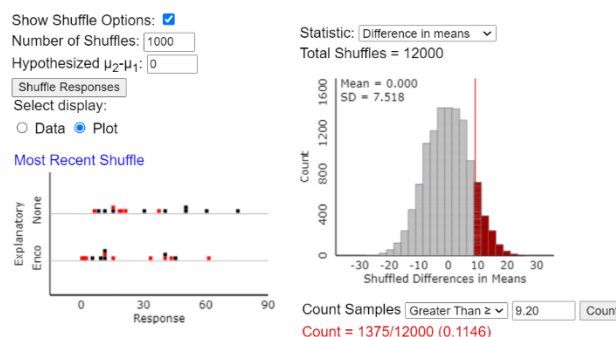
Shape: _____ Center (Mean): _____ Variability (SD): _____

9. Change the Number of Shuffles to add to 9990 and click the Shuffle Responses button. You should now have a distribution with 10,000 simulated differences in sample means. Describe the simulated sampling distribution of $\bar{x}_{\text{encouragement}} - \bar{x}_{\text{none}}$.

Shape: _____ Center (Mean): _____ Variability (SD): _____

10. What percent of dots on the graph have values greater than or equal to 9.20? _____

In other words, IF there is no effect of encouragement on the number of barbell curls that students like the ones in this study can do ($\mu_{\text{encouragement}} - \mu_{\text{none}} = 0$), the estimated probability of obtaining a difference in sample means at least as surprising as $\bar{x}_{\text{encouragement}} - \bar{x}_{\text{none}} = 9.20$ by chance alone due to the random assignment is _____.



11. Based on the estimated probability in Step 10, do these data provide *convincing evidence* that people like the ones in this study who receive encouragement can do more barbell curls, on average, than those who receive no encouragement? Explain your answer.

How does light intensity affect plant growth?

Relationships between two quantitative variables

Meadowfoam seed oil is used in making various skin care products. Researchers interested in maximizing the productivity of meadowfoam plants designed an experiment to investigate the effect of different light intensities on plant growth. The researchers planted 120 meadowfoam seedlings in individual pots, placed 10 pots on each of 12 trays, and put all the trays into a controlled enclosure. Two trays were randomly assigned to each light intensity level (micromoles per square meter per second): 150, 300, 450, 600, 750, and 900. The number of flowers produced by each plant was recorded and the average number of flowers was recorded for each tray. Here are the data:

Light intensity	150	150	300	300	450	450	600	600	750	750	900	900
Average number of flowers	62.3	77.4	55.3	54.2	49.6	61.9	39.4	45.7	31.3	44.9	36.8	41.9

(Source: M. Seddigh and G.D. Jolliff, "Light Intensity Effects on Meadowfoam Growth and Flowering," *Crop Science* 34 (1994): 497–503.)

1. Go to www.rossmanchance.com/applets/index2021.html. Choose Data Analysis → Least Squares Regression.
2. Open the Meadowfoam Experiment [data file](#).
3. Copy and paste the data into the Enter Data panel.
4. Click the Use Data button. A scatterplot of the data should appear.
5. Describe the relationship between average number of flowers produced and light intensity for the meadowfoam plants in this experiment.

6. Identify and interpret the slope of the least-squares regression line in context.

7. Use these data to estimate the average number of flowers produced, *irrespective of light intensity*.

8. How much better does the least-squares regression line with $x =$ light intensity do at predicting average number of flowers than using the mean \bar{y} to predict average number of flowers *irrespective of light intensity*?

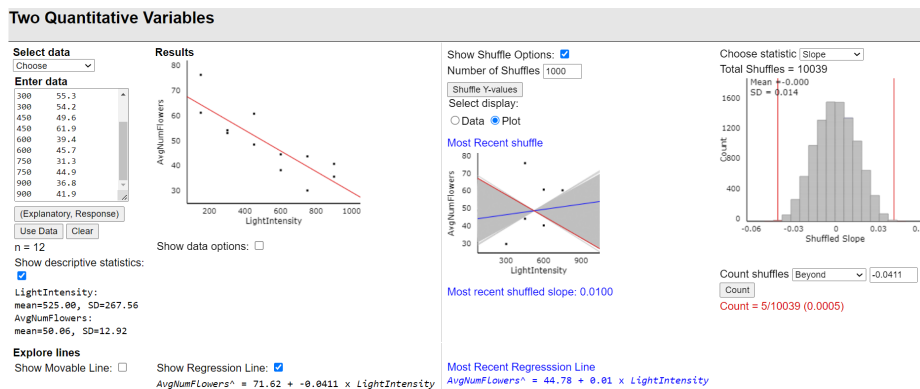
9. If we use the least-squares regression line to predict average number of flowers from light intensity, about how far off is a typical prediction?

10. Do the data give *some* evidence of a linear relationship between light intensity and average number of flowers produced for meadowfoam plants? Explain your answer.

11. On the right side of the applet screen, click the box to Show Shuffle Options. Be sure that Cards is chosen under Select display. In the Choose statistic drop-down menu at the top right, select Slope.
12. Click the Shuffle Y-values button. What just happened?

13. Change the Number of Shuffles to 1000 and click Shuffle several times. What do you notice about the simulated distribution of the sample slopes?

14. What percent of dots on the graph have values at least as extreme (in either direction) as -0.0411 ? _____
(Use the Count shuffles Beyond feature under the graph to help answer this question.)
15. In other words, IF there is actually no linear relationship between light intensity and average number of flowers produced in the population of meadowfoam plants like the ones in this study, the estimated probability of obtaining a sample regression line with a slope at least as surprising (in either direction) as -0.0411 purely by chance due to the random assignment is _____.
16. Based on the estimated probability in Step 15, do these data provide convincing evidence of a linear relationship between light intensity and average number of flowers produced in the population of meadowfoam plants like the ones in this study? Explain your answer.



17. In the left panel, click the radio button to show the Regression table. How does the estimated probability in Step 15 compare with the P -value in the computer output?

18. Now click the radio button to obtain a 95% Confidence Interval for Slope. Interpret this interval.