

SPARKING DEEPER UNDERSTANDING WITH SIMULATIONS IN STATISTICS

DOUG TYSON AND LUKE WILCOX

CMC-S 2016

PALM SPRINGS, CA

SIMULATION 1 – SMELLING PARKINSON’S DISEASE	(TEST FOR A COUNT/PROPORTION)
SIMULATION 2 – FLIPPING AND SPINNING (PENNIES)	(TEST FOR A COUNT/PROPORTION)
SIMULATION 3 – B IS FOR BOTOX (AND BACK PAIN)	(TEST FOR A DIFFERENCE IN PROPORTIONS)

COPIES OF THIS HANDOUT CAN BE FOUND AT MRTYSONSTATS.COM

SMELLING PARKINSON'S DISEASE

INTRODUCTION

As reported by the Washington Post (<http://tinyurl.com/SmellPark>), Joy Milne of Perth, UK, smelled a “subtle musky odor” on her husband Les that she had never smelled before. At first, Joy thought maybe it was just from the sweat after long hours of work. But when Les was diagnosed with Parkinson's 6 years later, Joy suspected the odor might be a result of the disease.

Scientists were intrigued by Joy's claim and designed an experiment to test her ability to “smell Parkinson's.” Joy was presented with 12 different shirts, each worn by a different person, some of whom had Parkinson's and some of whom did not. The shirts were given to Joy in a random order and she had to decide whether each shirt was worn by a Parkinson's patient or not.

1. Why would it be important to know that someone can smell Parkinson's disease?
2. How many correct decisions (out of 12) would you expect Joy make if she couldn't really smell Parkinson's and was just guessing?
3. How many correct decisions (out of 12) would it take to *convince* you that Joy really could smell Parkinson's?

SIMULATING THE EXPERIMENT

Although the researchers wanted to believe Joy, there was a chance that she may not really be able to tell Parkinson's by smell. It's logical to be skeptical of claims that are very different than our experiences. If Joy couldn't really distinguish Parkinson's by smell, then she would just have been guessing which shirt was which. The researchers were not willing to commit time and resources to a larger investigation unless they could be convinced to that Joy's wasn't just guessing. When researchers have a claim that they suspect (or hope) to find evidence against, it's called the **null hypothesis**.


4. What claim were the researchers hoping to find evidence *against*? That is, what was their prior belief (**null hypothesis**) about the ability to smell Parkinson's?
5. What claim were the researchers hoping to find evidence *for*? This is called the **alternative hypothesis** or the **research hypothesis**.

To investigate the idea that Joy was just guessing which shirt was worn by which type of person, we will assume that the null hypothesis is true.

6. Your instructor will hand you 12 cards (shirts) that have been shuffled into a random order. Don't turn them over yet! On the back of some of them is "Parkinson's" and on the back of others is "No Parkinson's." For each card, guess Parkinson's or No Parkinson's. Once you have made your guess, turn the card over and see if you were correct. Repeat this for each card and record the number of correct identifications (out of 12) below.

Tally of correct identifications	Number of correct identifications	Proportion of correct identifications

7. Create a dotplot of the number of correct identifications with the rest of the class. Record the results below.



8. In the actual experiment, Joy identified 11 of the 12 shirts correctly. Based on the very small-scale simulation by you and your classmates, what proportion of the simulations resulted in 11 or more shirts correctly identified, assuming that the person was guessing?
9. The proportion you just calculated is a crude estimate of a true probability called a **P-value**. How might we improve our estimate of the true probability?

STATISTICAL INFERENCE FROM THE SIMULATION

10. Use the SPA Applet for One Categorical Variable at stapplet.com/SPA to run this simulation 10000 times. Then use that simulation to get a (likely) better estimate of the p -value for 11 or more shirts correctly identified, assuming that this person was just guessing. Is it *possible* that Joy correctly identified 11 shirts just by random chance (guessing)? Is it *likely*?
11. An interesting side note is that Joy's one "mistake" really wasn't a mistake. The shirt was worn by a person who supposedly didn't have Parkinson's even though Joy claimed that she could smell the telltale smell on that shirt. That person called the experimenters 8 months after the experiment and reported that he had just been diagnosed with Parkinson's disease. That meant that Joy correctly identified 12 out of 12 shirts. What is the approximate P -value for 12 shirts correctly identified, assuming that this person was just guessing?

Note: A small P -value is considered strong evidence against the null hypothesis and in favor of the alternative hypothesis. But how small is small? As a rule of thumb, statisticians generally agree that P -values below 0.05 provide pretty strong evidence against the null hypothesis. Observed results with small P -values are said to be **statistically significant**.

DEEPER MATHEMATICAL CONNECTIONS

12. The true theoretical probability to get k successes in n trials when there is a true probability p of a success on each trial is given by the **binomial probability formula**: $nkpk^{1-p}n-k$. Compute the exact theoretical probability to get 11 or more successes in 12 trials when the true probability of success is 0.5. (*Hint:* calculate the probability for 11 successes and then do another calculation for 12 successes and then add these probabilities together.)

FLIPPING AND SPINNING PENNIES

FLIPPING PENNIES

1. Consider the penny given to you by your instructor. If you flip the penny, you might be interested in whether the coin is fair (balanced). Brainstorm some ways to determine whether the coin is fair (balanced).
2. In the absence of fancy equipment or techniques, we could investigate the fairness of the coin by spinning the coin repeatedly. In fact, that's what you should do right now. Your teacher will give you a penny and you should spin it 50 times. Make sure you have a hard, flat surface and that the penny spins without any obstructions. If the penny bumps into something, do not count that spin. As you spin, tally the outcome of each spin in the table below. When you are finished 50 spins, record the number of tails and heads, and the proportion of tails and heads.

	Heads	Tails	Total
Tally			50
Count			50
Proportion			1.0

3. Assume for a moment that your penny is fair. What number of tails would you expect to see in 50 flips?
4. Do you think that 50 flips will *always* produce the number of tails you answered in question number 3 above?
5. Suppose you flip a penny that you believe is fair, but lands tails up 26 times out of the 50 spins. Would this be enough evidence to convince you that the coin is unfair? How about 28 times? 30 times? 34 times? 45 times? Fill in the table below.

Number of Tails in 50 Spins	26	28	30	34	45
Do you believe the coin is unfair?					


6. Again suppose that you have a fair penny, what number of tails would you need to see before you would change your beliefs and conclude that the coin is unfair?
7. Based on the number of tails you obtained in your 50 flips, do you suspect that your coin is unfair?

SPINNING PENNIES

8. Now, spin your penny on its edge 50 times. Record your results in the table below.

	Heads	Tails	Total
Tally			50
Count			50
Proportion			1.0

9. Imagine a penny that is spun 50 times and has landed tails up 34 times. You might wonder how likely this is to happen by chance. Below is a histogram that shows the number of tails for 1000 sets of 50 spins of a fair penny from a computer simulation. In how many of the 1000 sets did 34 or more tails occur? Turn that number into a percentage.



10. From the histogram above, what is the *probability*/chance that in one set of 50 flips/spins, a fair coin lands tails up 34 times or more (just by chance alone)?

11. From the histogram above, what is the *probability*/chance that in one set of 50 flips/spins, a fair coin lands tails up 28 times or more (just by chance alone)?

STATISTICAL INFERENCE FROM THE SIMULATION

When a coin is spun and it lands tails up enough times to make you suspicious, one of two things is true:

- The coin is fair and you got lots of tails just by chance.
- The coin is not fair.

12. Use the histogram to find a probability that a fair coin lands tails up (just by chance) as many or more times as your penny did. Do you have enough evidence to reject the belief that *your* penny is fair?



Obverse (front) of the
Lincoln Memorial
penny

Statisticians have a rule of thumb for making the kind of decision you made in #12 above: when the probability for some event to occur by chance is less than 5%, we tend to discard the original belief (that the coin was fair) and choose to believe the new belief (that the coin is not fair). An event that is unlikely to happen by chance is called **statistically significant**. That is, although lots of tails from a fair coin can occur by chance *sometimes*, we don't believe that it occurred by chance *this one time* because the probability is so small.



Reverse (back) of the
Lincoln Memorial
penny

B IS FOR BOTOX (AND BACK PAIN)

THE EXPERIMENT

Botulinum toxin (Botox) injections are used to remove wrinkles, but they have other uses as well. One experiment investigated the effectiveness of Botox for treating back pain. A sample of 31 patients who suffered back pain were randomly assigned to receive injections of either Botox or saline. After 3 weeks, 11 of the 15 patients who received Botox had a large amount of pain relief. Of the 16 who received saline, only 4 had a large amount of pain relief during the same time period.

1. Organize the given counts into the following two-way table.

		Type of Injection		
		Botox	Saline	Total
Large Pain Relief	Yes			
	No			
	Total			

2. Calculate the **conditional proportion** of the Botox group that had pain relief. Do likewise for the saline group.
3. Calculate the difference in these proportions (Botox – saline).
4. Is there *some* evidence that the Botox is more effective than the saline at relieving back pain? Why or why not?
5. Is there *convincing* evidence that the Botox is more effective than the saline at relieving back pain? (Why can't you really answer this question – at least for now?)
6. State the null and alternative hypotheses for this experiment.

SIMULATING THE EXPERIMENT

One possible explanation for the observed difference in proportions is that the patients would have had the same effect for their pain no matter which type injection they received (that is, the injections were equally effective), and that the observed difference in proportions was just the result of the random shuffling of patients into groups.

7. Explain in detail how you could use cards to simulate the difference in proportions that would result from random assignment under this (null) hypothesis.

8. Use the Randomization Test for a Difference in Proportions applet on the StatKey website (<http://lock5stat.com/statkey/>) to run a simulation to estimate the p -value for an observed difference in proportions as large or larger than 0.483, assuming that the null hypothesis is true.

STATISTICAL INFERENCE FROM THE SIMULATION

9. Based on the p -value from your simulation, would you say that it is reasonable to believe that a difference of 0.483 or larger would occur *just by the chance of the random assignment*?

10. What other reason might there be for such a large difference in proportions between these two groups?

11. Is a cause-and-effect conclusion justified in this situation? Explain.