Introduction

The Nov 18, 2020 NYTimes article "New Pfizer Results" says

> Pfizer and BioNTech's trial included nearly 44,000 volunteers, half of whom received the vaccine. ... out of 170 cases of Covid-19, 162 were in the placebo group, and eight were in the vaccine group. Out of 10 cases of severe Covid-19, nine had received a placebo.

Based on these numbers, which fraction tells you how effective is the vaccine at preventing *severe* Covid-19?

A. $\frac{9}{10}$

B. $\frac{8}{9}$

C. $\frac{1}{9}$

D. $\frac{1}{10}$

E. $\frac{22000-1}{22000}$

F. $\frac{22000-9}{22000}$

Could this difference (9 vs. 1) reasonably be attributed to chance?

Coin Flipping

If you flip a coin 10 times, would you be surprised to get 9 heads and 1 tail?

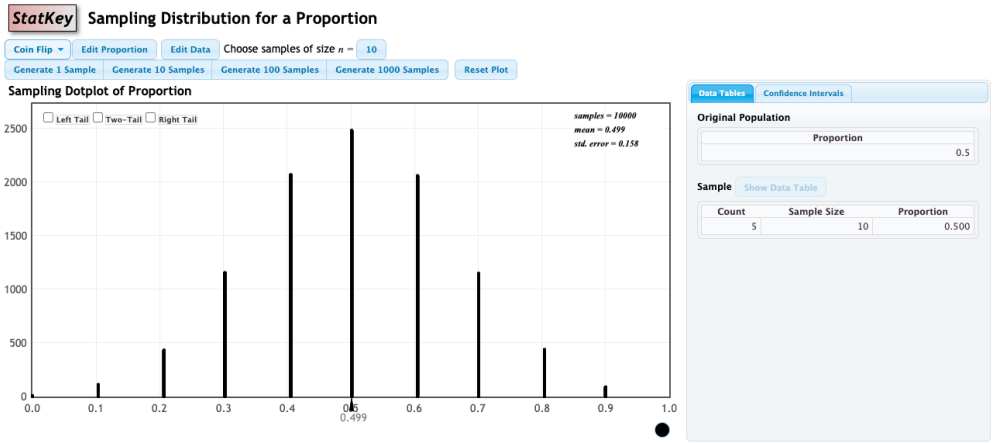

Use the statkey plot to estimate the probability of getting ≤ 1 tail out of 10 throws, by chance.

Use the google sheets function BINOMDIST to compute the probability of ≤ 1 tail out of 10 coin flips, in theory.

```
= BINOMDIST(num_successes, num_trials, prob_success, cumulative)
```

Why is it more appropriate to use the cumulative probability (probability of $\leq$ 1 tail) instead of the non-cumulative probability (probability of exactly 1 tail)?

Hint: for 10 flips, the difference between the probability of exactly 1 tail and $\leq$ 1 tail is not important. But for 100 flips, the difference between, say, the probability of exactly 45 tails and $\leq$ 45 tails is important.

- How likely is it to get exactly 45 tails?

- How likely is it to get $\leq$ 45 tails?

- Would you be surprised if you flipped a coin 100 times and got 45 tails?

Going back to the problem of the chance of getting ≤ 1 tail out of 10 coin flips, we said that the probability was ....

This number is called a *p-value*. The p-value of ..... means that

If we are gauging how odd it is to get ≤ 1 tail out of 10 throws, it would be equlvalently odd to get ≤ 1 head out of the 10 throws. So it would be reasonable to quote a "2-sided p-value" which is ...

P-values for coin flips can also be computed in google colab using the commands

```
%load_ext rpy2.ipython.
```

to allow the use of R, and then

```
%%R
pbinom(1, 10, 0.5, TRUE)
```

What does this have to do with the vaccine trial?

Johnson and Johnson Trial

In the Johnson & Johnson covid-19 vaccine trial, there were 44,325 volunteers, split approximately evenly between the vaccine arm and the placebo arm of the trial. A total of 468 cases of symptomatic covid-19 were reported as of Jan 29, 2021.

I could not find information on how many cases were in each arm of the trial, or how many people were in each arm, but let's pretend that 350 cases were in the placebo arm and 118 were in the vaccine trial. Let's also assume that the volunteers were split very closely to half and half between the arms, with 1 extra volunteer in the vaccine arm.

Based on these numbers, how effective is the J&J vaccine at preventing symptomatic covid-19?

What proportion of the cases were in the placebo arm?

Could this difference (350 cases vs. 118 cases) reasonably be attributed to chance? Justify your answer by finding the probability of seeing this extreme a result by chance, if the vaccine has no effect.

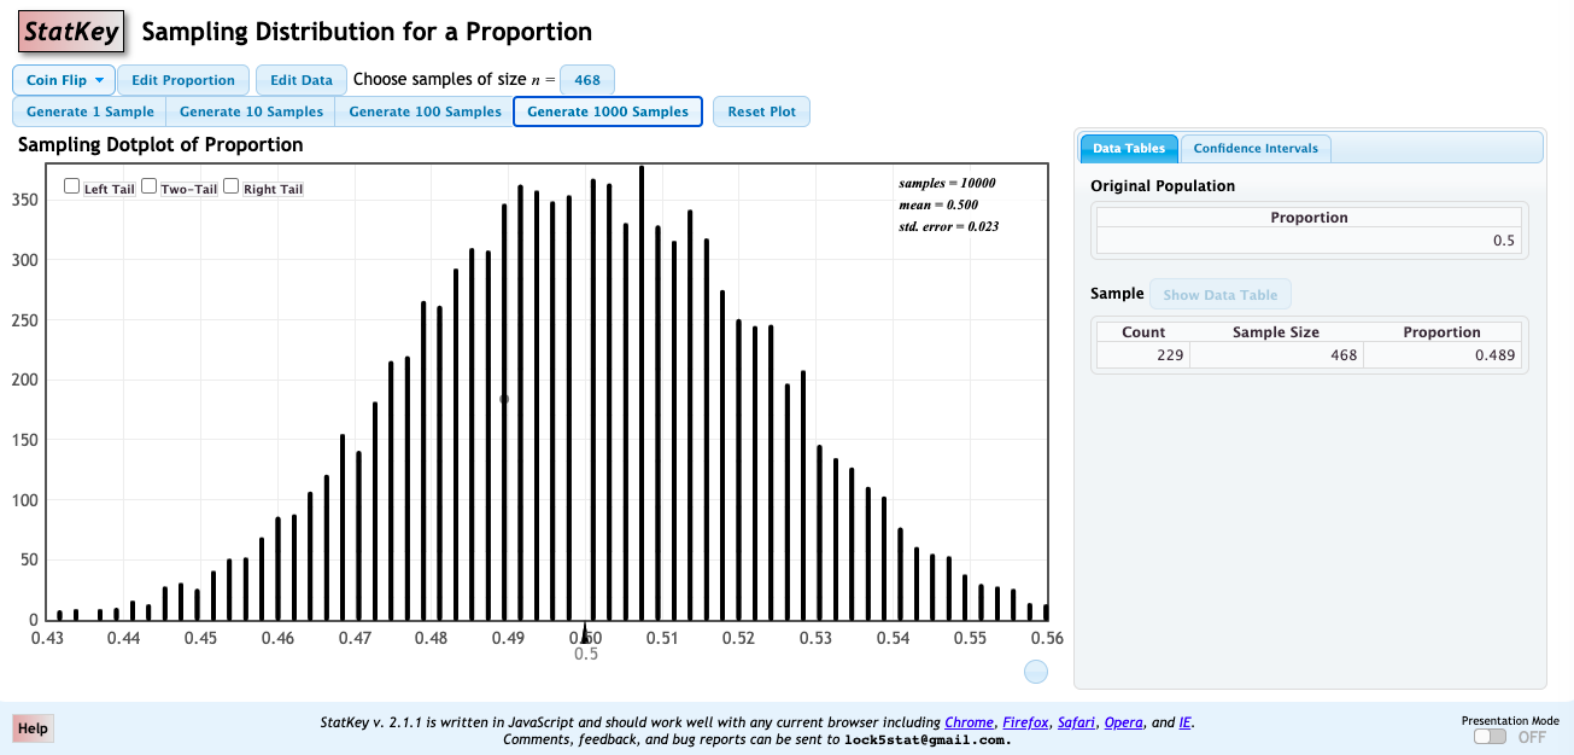
StatKey
Sampling Distribution for a Proportion
Coin Flip
Edit Proportion
Edit Data
Choose samples of size n =
468
Generate 1 Sample
Generate 10 Samples
Generate 100 Samples
Generate 1000 Samples
Reset Plot
Sampling Dotplot of Proportion
Left Tail
Two-Tail
Right Tail
samples = 10000
mean = 0.500
std. error = 0.023
350
300
250
200
150
100
50
0
0.43
0.44
0.45
0.46
0.47
0.48
0.49
0.5
0.51
0.52
0.53
0.54
0.55
0.56
Data Tables
Confidence Intervals
Original Population
Proportion
0.5
Sample
Show Data Table
Count
Sample Size
Proportion
229
468
0.489
Help
StatKey v. 2.1.1 is written in JavaScript and should work well with any current browser including Chrome, Firefox, Safari, Opera, and IE.
Comments, feedback, and bug reports can be sent to lock5stat@gmail.com.
Presentation Mode
OFF

Difference of Proportions

Here is another, more standard approach to analyzing the probability of getting the J& J vaccine result by chance.

Recall, we are assuming that there were 350 covid cases among $\tilde{2}2163$ participants who were given the placebo, and 118 new infections among the 22162 participants who received the vaccine.

What proportion of people got covid infections in the placebo arm of the trial?

What proportion of people got covid in the vaccination arm of the trial?

What is the difference in proportions?

What proportion of people got covid among all of the trial participants?

If the vaccine has absolutely no effect, then the trial is like picking people at random out of a group of people, some of whom are destined to get covid, and putting them in the two groups.

Suppose we have a population in which _____ percent overall got covid, and we take a sample of 22163 people and a sample of 22162 people at random, and look at the difference in proportions of people who got covid in those two samples. The plot below shows the distribution of differences in proportion we get if we do this over and over again.

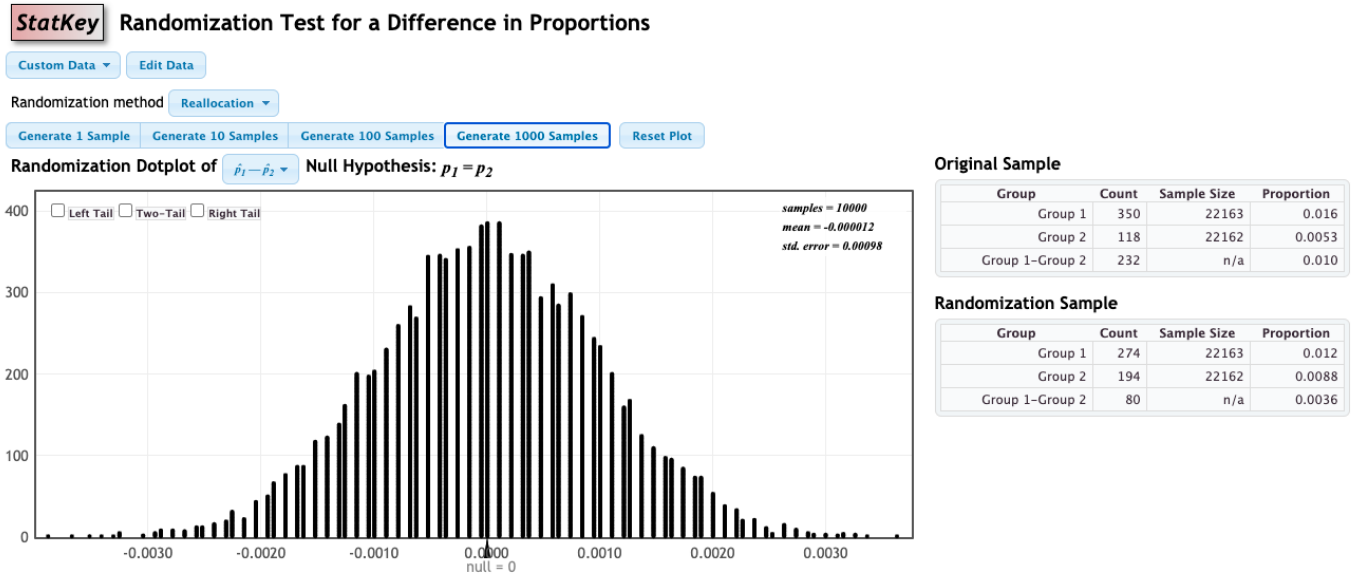

Use the statkey plot to estimate the probably of seeing this extreme a result by chance, if the vaccine actually has no effect.

Use google colab to find the exact probability.

```
[1] %load_ext rpy2.ipython
```

```
%%R
prop.test(c(350, 118), c(22163, 22162))
```

```
        2-sample test for equality of proportions with continuity correction

data:  c(350, 118) out of c(22163, 22162)
X-squared = 115.23, df = 1, p-value < 2.2e-16
alternative hypothesis: two.sided
95 percent confidence interval:
 0.008522016 0.012413298
sample estimates:
     prop 1      prop 2
0.015792086 0.005324429
```

The output says

p-value < 2.2e-16

What does this mean?

Note: there are also calculators online for finding the p-value for a test of proportions, e.g. https://www.socscistatistics.com/tests/chisquare/default2.aspx

HIV Vaccine

Extra Example. In 2009 there was an HIV vaccine trial in Thailand in which 8,197 volunteers were given the vaccine and 8,198 were given the placebo. New HIV infections occurred in 51 people who received the vaccine and 74 people who received the placebo.

- Find the efficacy of the vaccine in this trial.

- Could the difference in number of cases between the two arms be attributed to chance? Find a p-value to support your answer.

A second trial of the same HIV vaccine, called the Uhambo trial, was conducted in 2019. A Feb 3 National Institute of Health (NIH) report states:

> In the January 23, 2020 interim analysis, the DSMB [Data and Safety Monitoring Board] examined data from 2,694 volunteers who received the investigational vaccine regimen and 2,689 volunteers who received the placebo injection. ... In this analysis, 129 HIV infections occurred among the vaccine recipients, and 123 HIV infections occurred among the placebo recipients.

What do you think is the best interpretation of this result?

A. The vaccine works.

B. The vaccine makes it more likely to get HIV.

C. The vaccine doesn't have any effect on contracting HIV.

How do you reconcile the results of this second trial with the results of the first, which had a p-value < 0.05?