

Margin of Error

After completing this section, students should be able to:

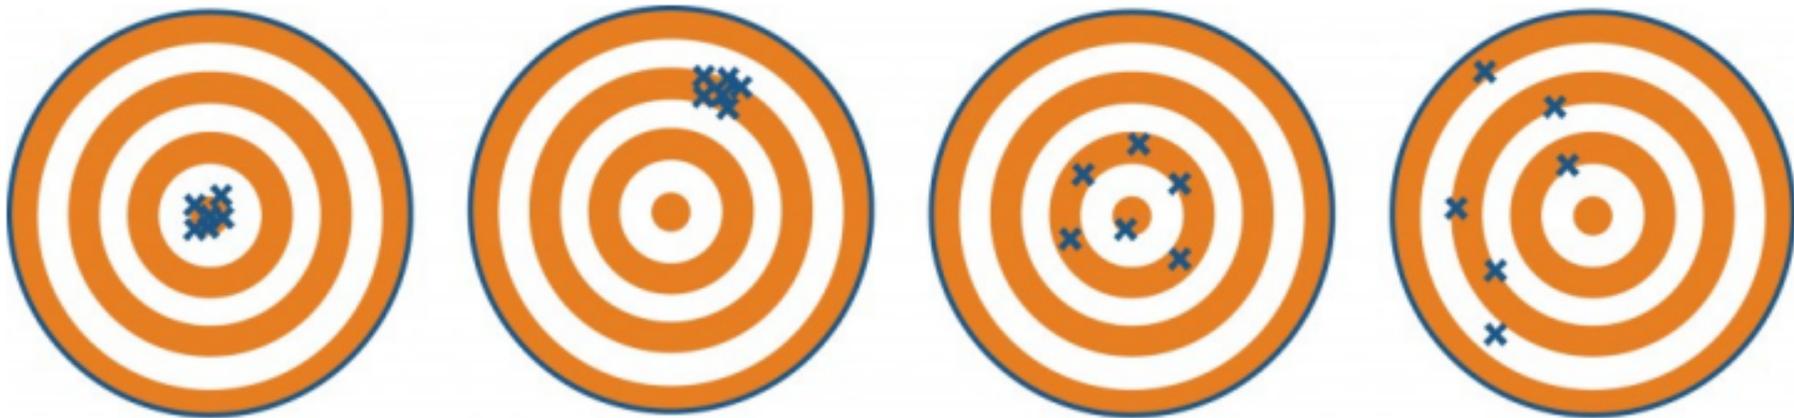
- Explain the relationship between sample size and variability in sample results.
- Estimate a 95% confidence interval for a simple random sample given the sample percentage and the sample size.

Precision vs. accuracy

When I go geocaching, I am using my phone to hone in on where a small box, called the cache, is located based on the GPS coordinates recorded by the person who hid it. I don't always find the geocache, and there are a variety of things that can go wrong, including:

- Sometimes due to tall buildings or hills, my phone is bouncing around with the GPS coordinates. First, it says the cache is 3 feet away over here, then it says, no that is 35 feet away and 3 feet away is at this other spot, etc. So I end up having to search a wide area for the geocache. This is a **LACK OF PRECISION**.
- Sometimes, the person who hid the cache gives inaccurate GPS coordinates to begin with, or the person who last found the cache didn't put it back in the same place. Then, even if my phone is precisely honing in on a tiny location without bouncing around, it is honing in on the wrong spot. This is a **LACK OF ACCURACY**.

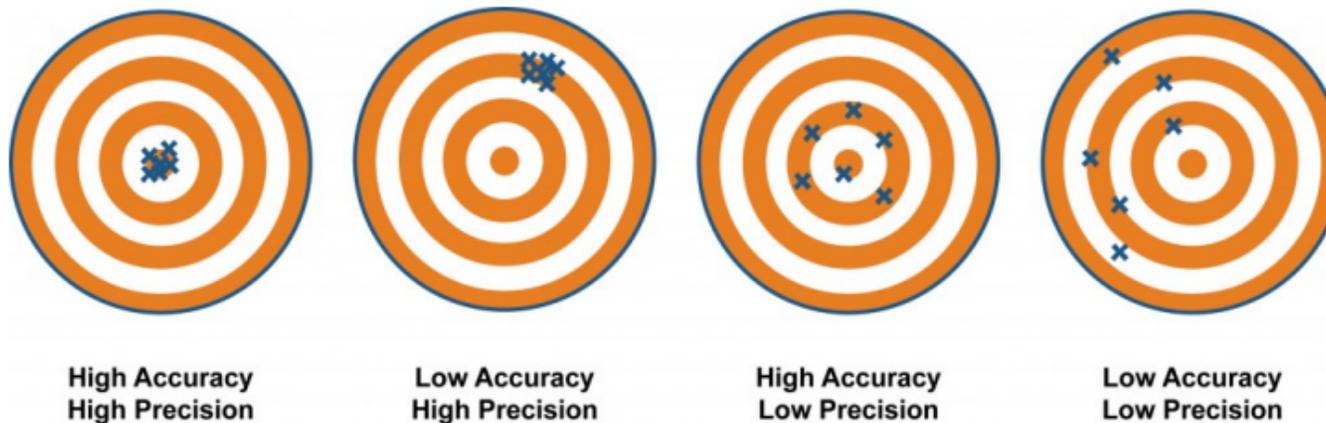
Which of these images have high precision vs. low precision? High accuracy vs. low accuracy?



Last time we talked about bias. If a poll is biased, it will likely have low (circle one) precision / accuracy.

Even if there is no bias, there is chance variability around who you happened to talk to, leading to a lack of (circle one) precision / accuracy. A bigger sample size will increase the (circle one) precision / accuracy of a poll but does not necessarily increase the (circle one) precision / accuracy if the poll is biased.

Margin of error is a measure of the (circle one) precision / accuracy of the poll.



Which of these images correspond to a poll with a large sample size but lots of bias due to inaccurate responses (e.g. a social stigma that inhibits people from answering honestly)?

A small sample size but not much bias?

Goals

Today's Goal: get a handle on how much variability there is in sample estimates

- even when using a simple random sample
- even with a 100% response rate
- even in the absence of any kind of bias

Today's Goal: make more sense of the opinion poll predictions for the 2020 presidential election (e.g. in North Carolina)

North Carolina	Poll source	wikipedia	Date(s) administered	Sample size ^[b]	Margin of error	Donald Trump Republican	Joe Biden Democratic	Jo Jorgensen Libertarian	Howie Hawkins Green	Other	Undecided
	SurveyMonkey/Axios		Oct 20 – Nov 2, 2020	5,363 (LV)	± 2%	48% ^[zt]	50%	-	-	-	-
	Change Research/CNBC		Oct 29 – Nov 1, 2020	473 (LV)	± 4.51%	47%	49%	2%	1%	-	1%
	Swayable		Oct 27 – Nov 1, 2020	690 (LV)	± 5.3%	46%	52%	1%	0%	-	-
	Ipsos/Reuters		Oct 27 – Nov 1, 2020	707 (LV)	± 4.2%	48% ^[n]	49%	1%	1%	2% ^[zu]	-
		48% ^[zv]				49%	-	-	3% ^[zw]	1%	
		48% ^[zx]				50%	-	-	2% ^[zy]	-	
	Data for Progress		Oct 27 – Nov 1, 2020	908 (LV)	± 3.3%	48%	50%	1%	1%	0% ^[zz]	-
	Frederick Polls/Compete Everywhere	^[CB]	Oct 30–31, 2020	676 (LV)	± 3.7%	49%	51%	-	-	-	-
	AtlasIntel		Oct 30–31, 2020	812 (LV)	± 3.0%	50%	48%	-	-	3%	-
	Insider Advantage/Center for American Greatness	^[H]	Oct 30–31, 2020	450 (LV)	± 4.6%	48%	44%	2%	-	-	7%
	Emerson College		Oct 29–31, 2020	855 (LV)	± 3.3%	47% ^[aaa]	47%	-	-	6% ^[aab]	-
	Morning Consult		Oct 22–31, 2020	1,982 (LV)	± 2%	48%	49%	-	-	-	-
	CNN/SSRS		Oct 23–30, 2020	901 (LV)	± 4.1%	45%	51%	2%	1%	1% ^[aac]	1%

- In the poll by Data for Progress of 908 likely voters, is 50% for Biden vs. 48% for Trump a big enough difference to trust?
- What is meant by the margin of error of ±3.3 percentage points?

Beads

In a real poll, we don't have info about the real population of likely voters. But for today, we are going to use a population of beads that we do know about, to see how well our samples give us the right answers.



Our analogy: the 4000 beads in the bin play the role of ... , and the blue beads represent ...

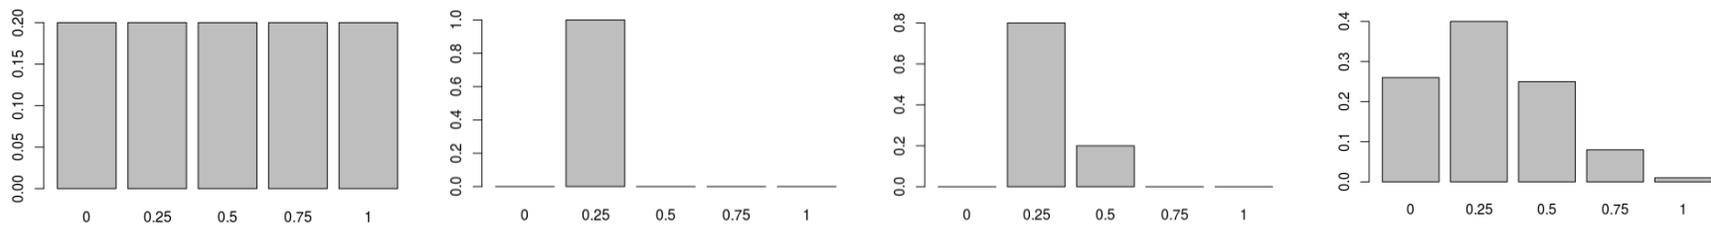
Note: In our bead box the true proportion of blue beads is 0.295. This is different from the voting situation where the true proportion of Biden voters is. ...

Overall question: If multiple samples of beads (likely voters) were taken, how different will the estimates for the proportion of blue beads (Biden voters) be?

Each student will take a sample of 4 beads with the spatula.



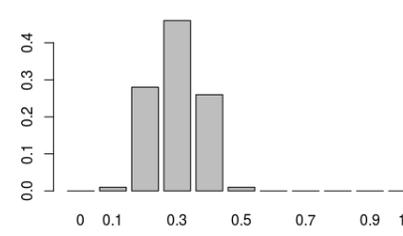
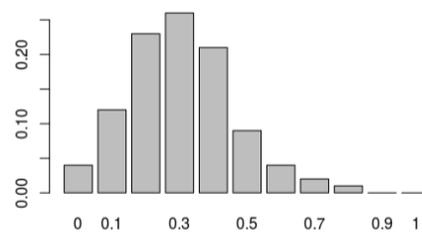
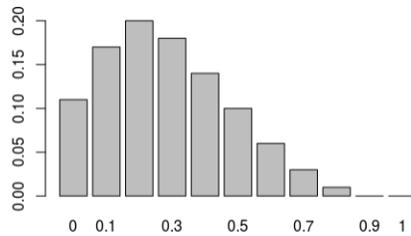
What do you think the distribution of sample proportions will look like for samples of size 4? (PollEv)



Take a sample of size 4 and enter the proportion of blue beads in the PollEv poll.



What do you think the distribution of sample proportions will look like for samples of size 10? (PollEv)



Take a sample of size 10 and enter the proportion of blue beads in the PollEv poll.



What do you think the distribution of sample proportions will look like for samples of size 20? (PollEv)

- (a) narrower than with samples of 10
- (b) about the same as with samples of 10
- (c) wider than with samples of 10

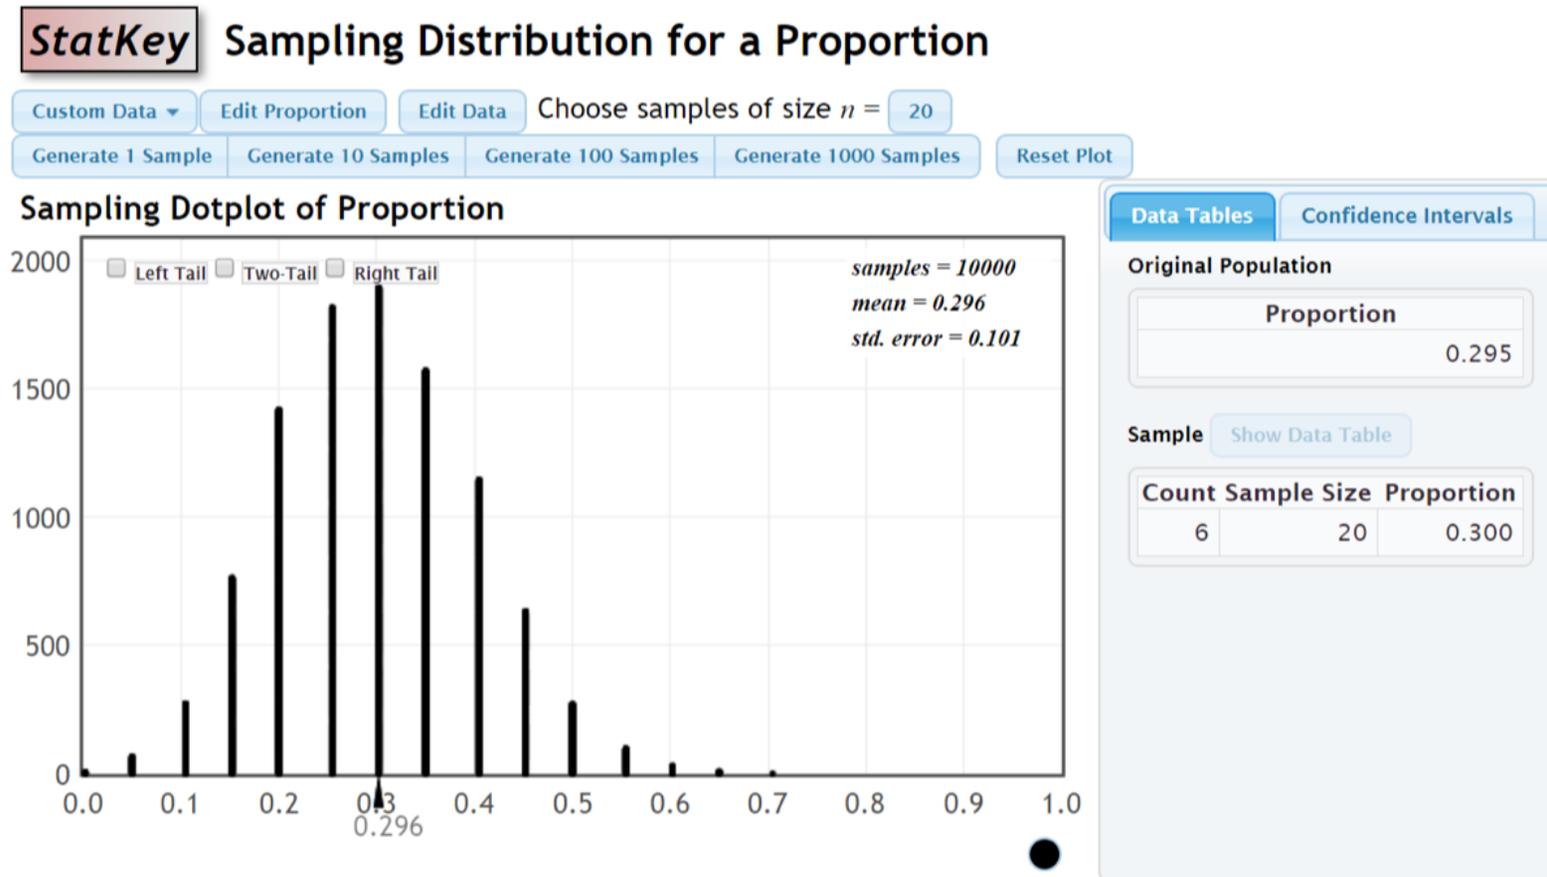
Take a sample of size 20 and enter the proportion of blue beads in the PollEv poll.

Observation: For larger sample sizes, you get (circle one) more / the same / less variability in the proportion, so (circle one) more / the same / less uncertainty in your estimates .

Conclusion: a larger sample size will give a (circle one) better / the same / worse estimate of the true proportion.

StatKey Simulations

Simulation of lots of samples from the bead box: population proportion is 0.295



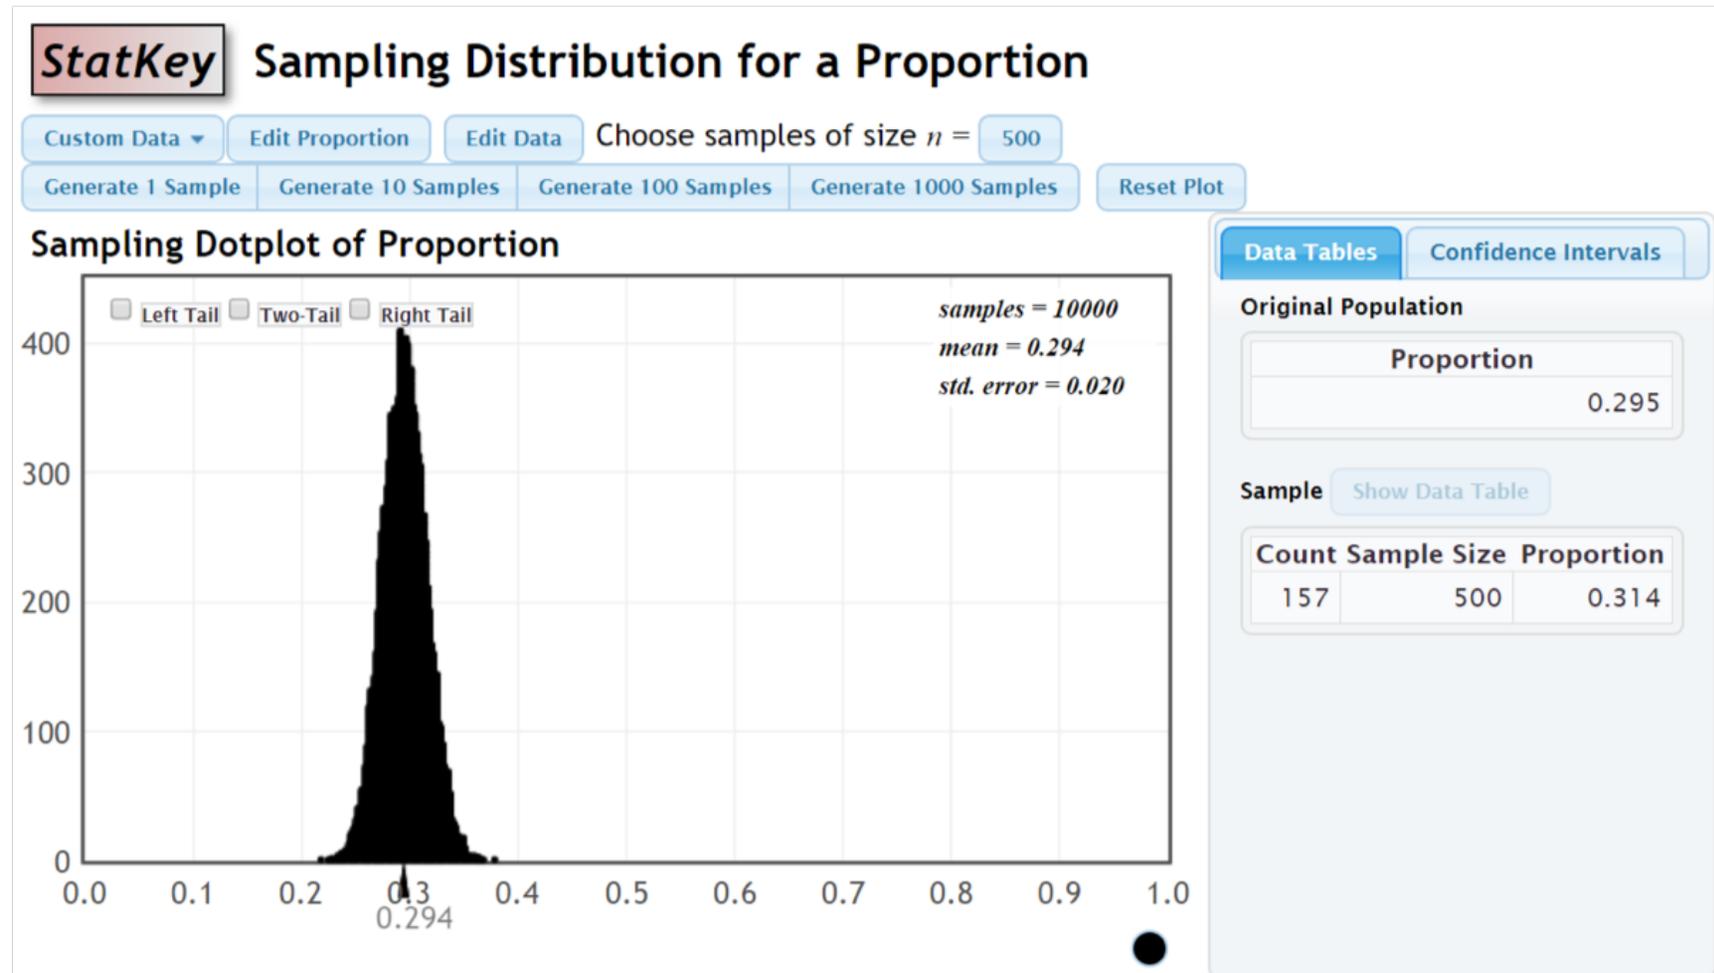
About 95% of sample proportions are in the interval _____.

About 95% of sample proportions are within _____ of the true proportion.

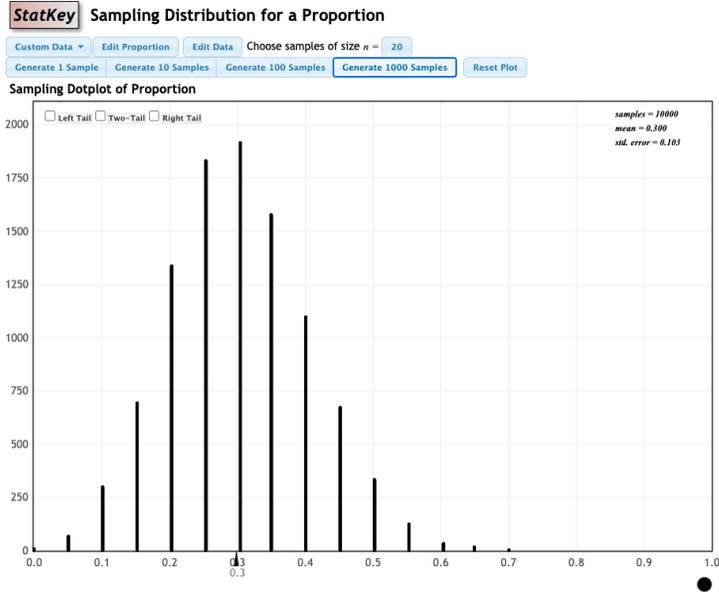
Margin of Error

- The 95% *margin of error* is the distance you have to go out from the center to capture 95% of the sampling distribution.
- In the example above, the 95% margin of error is _____
- You can also have a 90% margin of error or any other % margin of error, but 95% is most common.
- If you hear *margin of error* without the % specified, assume it is 95%.
- Would you expect a 90% margin of error to be a bigger or smaller number than a 95% margin of error, for the same sampling distribution?

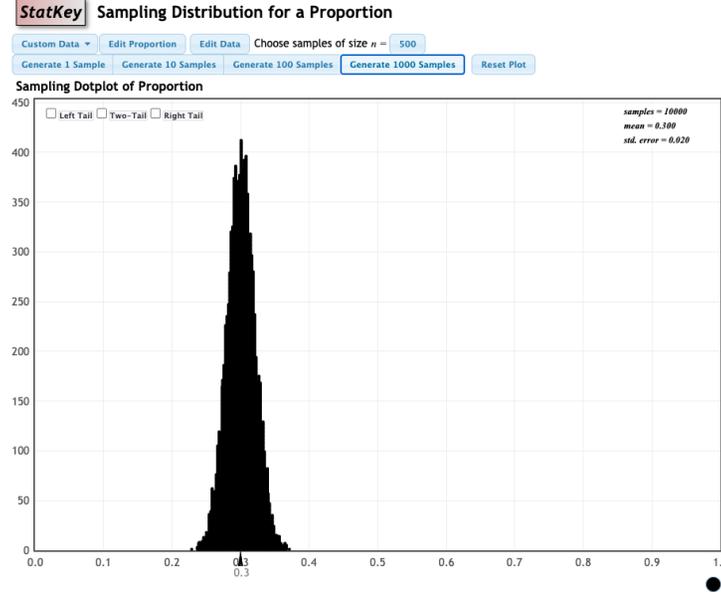
Simulation of lots of samples from the bead box: population proportion is 0.295



- About 95% of sample proportions are in the interval _____.
- About 95% of sample proportions are within _____ of the true proportion.
- What is the margin of error?



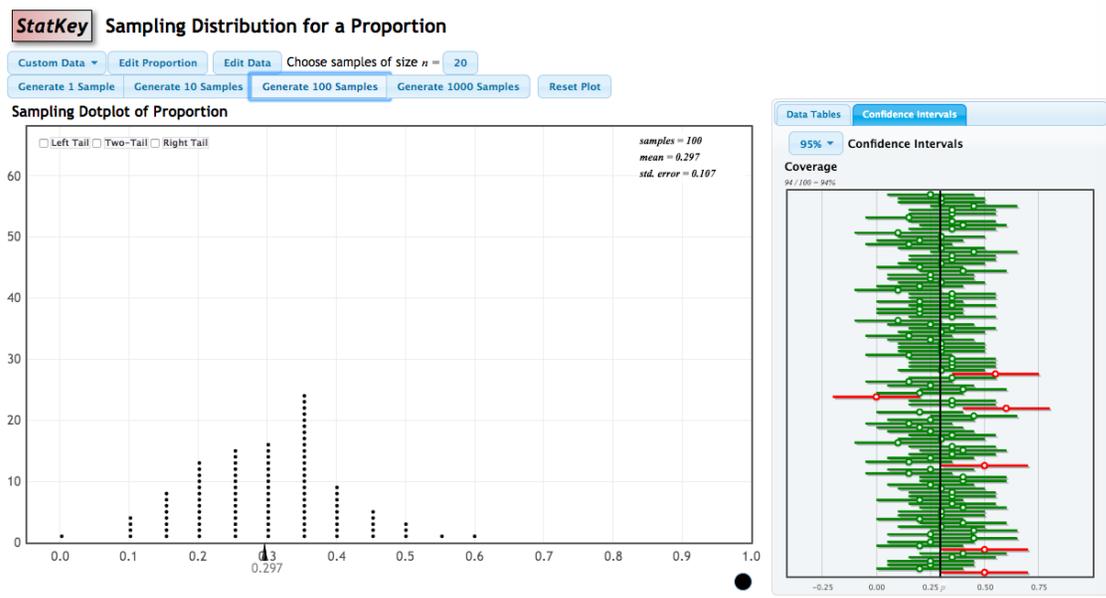
Taking lots of samples of size 20



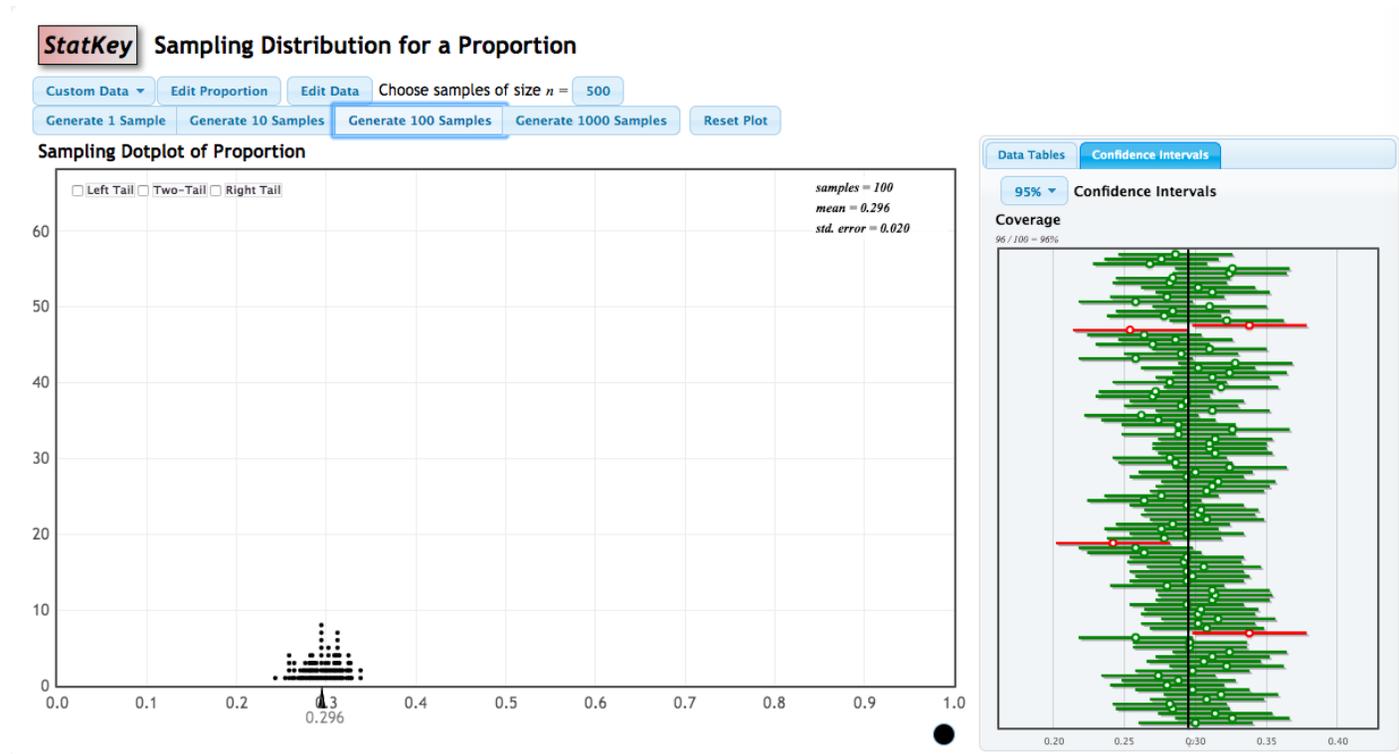
Taking lots of samples of size 500

- What is the relationship between *margin of error* and the *standard error (std. error)* reported by StatKey?

Confidence Intervals

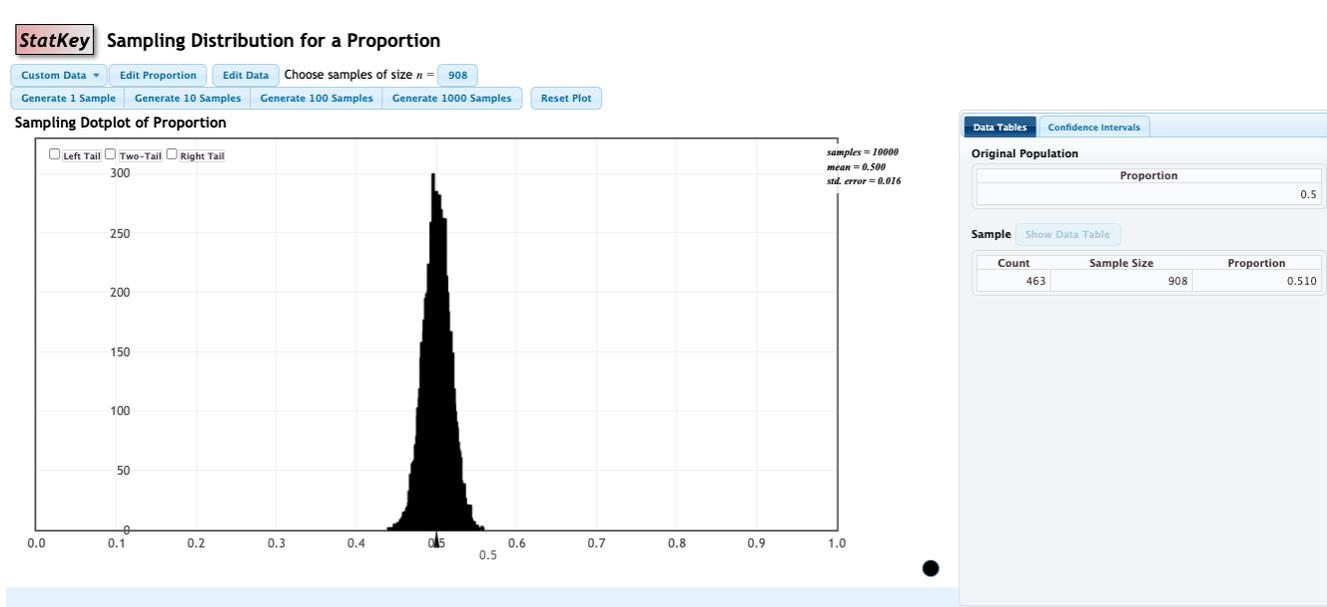


- On the right side of the figure above, the circles represent sample proportions, and the black line is the true population proportion.
- If we draw a box around the center black line, of width equal to the margin of error 0.2, then we should capture the sample proportions about 95% of the time.
- Alternatively, if we draw line segments coming out from each green circle, of length 0.2 on either side, then these line segments should capture the true proportion about 95% of the time.
- The *95% confidence interval* is the *sample proportion* \pm *margin of error*
- The confidence intervals are shown by the green (or red) lines.
- They are green if they overlap the true proportion and red if they don't.



- True or False: About 95% of the time the sample proportion will be in the interval given by population proportion \pm margin of error
- True or False: About 95% of the time the population proportion will be in the interval given by sample proportion \pm margin of error
- Which of the previous two statements are more useful in practice when we work with polls?

- In the previous examples, we estimated the margin of error from a distribution of sample proportions, with full knowledge of the true population proportion. But normally, we don't know the true population proportion. We ONLY know the proportion for ONE sample.
- For example, in the Data for Progress poll, we only know that in the sample of 908 people, 0.50% favor Biden.
- So we use the SAMPLE population instead as a proxy for the TRUE population proportion to estimate the margin of error.



Simulation of lots of samples of 908 people: suppose population proportion is really 0.48

- What is our estimate for the margin of error?

Computing margin or error

We have estimated margin of error using *simulation* – that is, an app like StatKey that simulates repeatedly physically re-sampling from a sample, with replacement. This is called *bootstrapping*.

1. we eyeballed a graph of sample proportions for repeated samples
2. we used twice the standard error for a bunch of sample proportions for repeated samples

There is a third method for finding the margin of error, that uses a formula instead of a simulation:

$$\text{margin of error} = 2 \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}}$$

where \hat{p} is the sample proportion and n is the sample size.

Once we have the margin of error, we compute the 95% confidence error from

$$\text{sample proportion} \pm \text{margin of error}$$

.

We use the sample proportion \hat{p} in the formula and not the true population proportion p because ...

Back to the polls

Back to the Data for Progress poll, 908 people were interviewed and 50% favored Trump ...

- Compute the margin of error using the formula $2 \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}}$
- Does this agree with the bootstrap estimate for the margin of error that we got from Statkey?
- Does it agree with the margin of error reported on Wikipedia?
- What is a 95% confidence interval for Biden's percent of the vote?

North Carolina	Poll source	wikipedia	Date(s) administered	Sample size ^[b]	Margin of error	Donald Trump Republican	Joe Biden Democratic	Jo Jorgensen Libertarian	Howie Hawkins Green	Other	Undecided
	SurveyMonkey/Axios		Oct 20 – Nov 2, 2020	5,363 (LV)	± 2%	48% ^[zt]	50%	-	-	-	-
	Change Research/CNBC		Oct 29 – Nov 1, 2020	473 (LV)	± 4.51%	47%	49%	2%	1%	-	1%
	Swayable		Oct 27 – Nov 1, 2020	690 (LV)	± 5.3%	46%	52%	1%	0%	-	-
	Ipsos/Reuters		Oct 27 – Nov 1, 2020	707 (LV)	± 4.2%	48% ^[n]	49%	1%	1%	2% ^[zu]	-
		48% ^[zv]				49%	-	-	3% ^[zw]	1%	
		48% ^[zx]				50%	-	-	2% ^[zy]	-	
	Data for Progress		Oct 27 – Nov 1, 2020	908 (LV)	± 3.3%	48%	50%	1%	1%	0% ^[zz]	-
	Frederick Polls/Compete Everywhere ^[CB]		Oct 30–31, 2020	676 (LV)	± 3.7%	49%	51%	-	-	-	-
	AtlasIntel		Oct 30–31, 2020	812 (LV)	± 3.0%	50%	48%	-	-	3%	-
	Insider Advantage/Center for American Greatness ^[H]		Oct 30–31, 2020	450 (LV)	± 4.6%	48%	44%	2%	-	-	7%
	Emerson College		Oct 29–31, 2020	855 (LV)	± 3.3%	47% ^[aaa]	47%	-	-	6% ^[aab]	-
	Morning Consult		Oct 22–31, 2020	1,982 (LV)	± 2%	48%	49%	-	-	-	-
	CNN/SSRS		Oct 23–30, 2020	901 (LV)	± 4.1%	45%	51%	2%	1%	1% ^[aac]	1%

What do you think of Trump’s 48% to Biden’s 50% in this poll? (PollEv)

- A. Based on this poll, we can be pretty confident Biden will win.
- B. Based on this poll, it’s too close to call.

North Carolina	Poll source	wikipedia	Date(s) administered	Sample size ^[b]	Margin of error	Donald Trump Republican	Joe Biden Democratic	Jo Jorgensen Libertarian	Howie Hawkins Green	Other	Undecided
	SurveyMonkey/Axios		Oct 20 – Nov 2, 2020	5,363 (LV)	± 2%	48% ^[zt]	50%	-	-	-	-
	Change Research/CNBC		Oct 29 – Nov 1, 2020	473 (LV)	± 4.51%	47%	49%	2%	1%	-	1%
	Swayable		Oct 27 – Nov 1, 2020	690 (LV)	± 5.3%	46%	52%	1%	0%	-	-
	Ipsos/Reuters		Oct 27 – Nov 1, 2020	707 (LV)	± 4.2%	48% ^[n]	49%	1%	1%	2% ^[zu]	-
		48% ^[zv]				49%	-	-	3% ^[zw]	1%	
		48% ^[zx]				50%	-	-	2% ^[zy]	-	
	Data for Progress		Oct 27 – Nov 1, 2020	908 (LV)	± 3.3%	48%	50%	1%	1%	0% ^[zz]	-
	Frederick Polls/Compete Everywhere	[CB]	Oct 30–31, 2020	676 (LV)	± 3.7%	49%	51%	-	-	-	-
	AtlasIntel		Oct 30–31, 2020	812 (LV)	± 3.0%	50%	48%	-	-	3%	-
	Insider Advantage/Center for American Greatness	[H]	Oct 30–31, 2020	450 (LV)	± 4.6%	48%	44%	2%	-	-	7%
	Emerson College		Oct 29–31, 2020	855 (LV)	± 3.3%	47% ^[aaa]	47%	-	-	6% ^[aab]	-
	Morning Consult		Oct 22–31, 2020	1,982 (LV)	± 2%	48%	49%	-	-	-	-
	CNN/SSRS		Oct 23–30, 2020	901 (LV)	± 4.1%	45%	51%	2%	1%	1% ^[aac]	1%

Use the $2 \sqrt{\frac{\hat{p} \cdot (1 - \hat{p})}{n}}$ estimate for the margin of error to compute the margins of errors for the first four polls on the list. Do your numbers agree with the ones reported?

1. Survey Monkey / Axios
2. Change Research / CNBC
3. Swayable
4. Ipsos / Reuters