## Learning outcomes

- Explain how and why polls are done by sampling
- Explain how sampling causes error
- Explain how error is characterized by confidence intervals
- Calculate a confidence interval in simple cases
- Choose suitable sample size for desired accuracy
- Recognize situations where you need an advanced formula
- Critically evaluate Cls reported in polls
- Identify errors that are not covered by Cl


## Different kinds of polls - same mathematics



Binary opinion on an issue, 2+1 choices
U.S. Adults Show Strong Support for Plastic Straw Bans

Do you support or oppose the new policy restaurants are enacting to use recyclable paper straws, instead of plastic straws, in the coming years?

■ Support ■ Oppose ■ Don't know, No opinion


## Different kinds of polls - same mathematics

Voting intent, multiple choices


Very rough meaning:
"The numbers should not be off by more than that, at least not too often"

Binary opinion on an issue, 2+1 choices

## U.S. Adults Show Strong

 Support for Plastic Straw BansDo you support or oppose the new policy restaurants are enacting to use recyclable paper straws, instead of plastic straws, in the coming years?

■ Support ■ Oppose ■ Don't know, No opinion

## Black box formula




If you master this, you have exceeded $95 \%$ of the population!
[Statistics made up.]

## Black box formula



In the sample, proportion of supporters is $67 \%$. This we know. In the population, we can now say that the proportion of supporters is $(67 \pm 2) \%$, at $95 \%$ confidence. Is this magic?

## Back to basics: How does sampling work?

Let us first understand the direct problem.

- What is the process that creates the numbers we observe?
- Where is the randomness?
- Is there a familiar distribution involved?

We do this on the blackboard.

Then hopefully we can understand the inverse or inference problem: What do the observations tell us?

## (From the blackboard)

- From a large population ( $N$ ), we took a small sample ( $n$ ).
- In the population we had $N p$ persons of the type we are interested in (e.g. supporters of party A).
- In the sample we will have some number $K$, which has binomial distribution $\operatorname{Bin}(n, k)$.
- Naturally we would like to say that the sample proportion

$$
\hat{p}=K / n
$$

is a good estimate for $p$. Can we do that?

## Properties of the estimator

- At least the expected value is just right,

$$
\mathrm{E}(\hat{p})=\mathrm{E}(K / n)=\mathrm{E}(K) / n=n p / n=p
$$

because $K$ is binomially distributed with expectation $n p$.

- But how far will it the estimator be from its expectation? The standard deviation

$$
\mathrm{D}(\hat{p})=\mathrm{D}\left(\frac{K}{n}\right)=\frac{\mathrm{D}(K)}{n}=\frac{\sqrt{n p(1-p)}}{n}=\sqrt{\frac{p(1-p)}{n}}
$$

gives a good understanding of this.

## Simulate taking a sample

```
% Simulate an opinion poll from a huge population
p = 0.14; % True proportion of supporters (in population)
n = 1000; % Sample size
k = binornd(n,p); % Take a sample of n persons
phat = k/n; % Observed proportion (in sample)
fprintf('Sample has K=%d: ', k);
fprintf(' estimate phat=%.3f,', phat);
fprintf(' error = %+.3f\n', phat-p);
```


## It is not very complicated!

## Simulate taking a sample

```
% Simulate an opinion poll from a huge population
p = 0.14; % True proportion of supporters (in population)
n = 1000; % Sample size
k = binornd(n,p); % Take a sample of n persons
phat = k/n; % Observed proportion (in sample)
fprintf('Sample has K=%d: ', k);
fprintf(' estimate phat=%.3f,', phat);
fprintf(' error = %+.3f\n', phat-p);
>> onepoll
Sample has K=135: estimate phat=0.135, error = -0.005
```


## Simulate taking a sample

```
O.Simuate an opinion poll from a huge population
p = 0.14; % True proportion of supporters (in population)
n = 1000; % Sample size
k = binornd(n,p); % Take a sample of n persons
phat = k/n; % Observed proportion (in sample)
fprintf('Sample has K=%d: ', k);
fprintf(' estimate phat=%.3f,', phat);
fprintf(' error = %+.3f\n', phat-p);
>> onepoll
Sample has K=135: estimate phat=0.135 error = -0.005
```

Oh well.
We were "expecting" $E(K)=140$ supporters, but we got only $K=135$. We were off by 5 . Equivalently, our proportion is off by $5 / 1000$.

Given that $D(K)=\sqrt{ }(n p(1-p))=11$, we should not be too surprised. This was a simulation, but same happens when you really sample.

## Simulate 20 polls

$$
\begin{aligned}
& \text { Poll \# 1: Sample has } \mathrm{K}=126 \text { : } \\
& \text { Poll \# 2: Sample has } \mathrm{K}=142: \\
& \text { Poll \# 3: Sample has } \mathrm{K}=146: \\
& \text { Poll \# 4: Sample has } \mathrm{K}=145 \text { : } \\
& \text { Poll \# 5: Sample has K=135: } \\
& \text { Poll \# 6: Sample has K=127: } \\
& \text { Poll \# 7: Saple has K=169: } \\
& \text { Poll \# 8: Sample has K=143: } \\
& \text { Poll \# 9: Sample has K=145: } \\
& \text { Poll \#10: Sample has K=155: } \\
& \text { Poll \#11: Sample has K=138: } \\
& \text { Poll \#12: Sample has K=145: } \\
& \text { Poll \#13: Saple has K=149: } \\
& \text { Poll \#14: Sample has K=136: } \\
& \text { Poll \#15: Sample has K=142: } \\
& \text { Poll \#16: Sample has K=149: } \\
& \text { Poll \#17: Sample has K=147: } \\
& \text { Poll \#18: Sample has K=122: } \\
& \text { Poll \#19: Saple has K=168: } \\
& \text { Poll \#20: Sample has } K=138:
\end{aligned}
$$

estimate phat=0.126, error $=-0.014$ estimate phat=0.142, error $=+0.002$ estimate phat=0.146, error $=+0.006$ estimate phat=0.145, error $=+0.005$ estimate phat=0.135, error = -0.005 estimate phat=0.127, error $=-0.013$ estimate phat=0.169, error $=+0.029$ estimate phat=0.143, error $=+0.003$ estimate phat=0.145, error $=+0.005$ estimate phat=0.155, error $=+0.015$ estimate phat=0.138, error $=-0.002$ estimate phat=0.145, error $=+0.005$ estimate phat=0.149, error $=+0.009$ estimate phat=0.136, error $=-0.004$ estimate phat=0.142, error $=+0.002$ estimate phat=0.149, error $=+0.009$ estimate phat=0.147, error $=+0.007$ estimate phat=0.122, error $=-0.018$ estimate phat=0.168, error $=+0.028$ estimate phat=0.138, error $=-0.002$

## Let's be very cheap

- Approximate the binomial with a normal, with same mean and same standard deviation

$$
K \sim N(n p, n p(1-p))
$$

- The error of our estimate is

$$
\hat{p}-p=(K / n)-p
$$

- So the error is normally distributed with zero mean and standard deviation

$$
\sqrt{\frac{p(1-p)}{n}} \approx \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}
$$

## And that gives us our basic formula

$$
\text { MOE }=1.96 \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}
$$

The estimate will be within $\pm$ MOE of the true proportion $95 \%$ of the time, when we do this kind of sampling.

That is because a normally distributed variable is within $\pm 1.96 \cdot$ sd of its own mean, $95 \%$ of the time.

We will say that

$$
[\hat{p}-M O E, \hat{p}+M O E]
$$

is the $95 \%$ confidence interval for $p$.

## Improved code: Compute Cl

```
% Simulate an opinion poll from a huge population
p = 0.14; % True proportion of supporters (in population)
n = 1000; % Sample size
k = binornd(n,p); % Take a sample of n persons
phat = k/n; % Observed proportion (in sample)
```

```
moe = 1.96 * sqrt(phat*(1-phat)/n);
```

moe = 1.96 * sqrt(phat*(1-phat)/n);
ci = [phat-moe, phat+moe];
ci = [phat-moe, phat+moe];
fprintf('Sample has K=%d: ', k);
fprintf(' estimate phat=%.3f,', phat);
fprintf(' CI = [%.3f, %.3f] ', ci(1), ci(2));
if ci(1)<=p \&\& p<=ci(2)
fprintf('OK\n');
else
fprintf('FAIL\n');
end

```

\section*{Simulate 20 polls}

\section*{Recall that the true proportion in our population is 0.140 .}
\begin{tabular}{|c|c|c|}
\hline Poll \# & Sample h & \\
\hline \# & Sample has & \\
\hline Poll \# & Sample has & \\
\hline Poll \# & Sample has & \\
\hline Poll \# & Sample h & \\
\hline Poll \# & Sample has & \\
\hline \# & Sample has & \\
\hline Poll \# & Sample has & \\
\hline Poll \# & Sample has & \\
\hline Poll \#10 & Sample has & \\
\hline Poll \# & Sample has & \(\mathrm{K}=13\) \\
\hline Ooll \# & Sample has & \\
\hline Poll \#13 & Sample has & \\
\hline oll \#1 & Sample ha & = \\
\hline Poll \#1 & Sample has & \\
\hline Poll \#1 & Sample has & =149 \\
\hline Poll \#1 & Sample has & 14 \\
\hline Poll \#18 & Sample has & \\
\hline oll \#19 & Sample has & \\
\hline & Sample h & \\
\hline
\end{tabular}

estimafe phat \(=0.126, \mathrm{CI}=[0.105,0.147] \quad \mathrm{OK}\)
estimate plrat 0.112 CT \(=[0.120 .0\) 161] On
estimate phat=0.146, CI = [0.124, 0.168] OK
estimate phat \(=0.145, \mathrm{CI}=[0.123,0.167] \quad 0 \mathrm{~K}\)
estimate phat=0.135, CI = [0.114, 0.156] OK
estimate phat \(=0.127\), CI \(=[0.106,0.148]\) OK
estimate phat \(=0.169, \mathrm{CI}=[0.146,0.192]\) FAIL
estimate phat \(=0.143, \mathrm{CI}=[0.121,0.165]\) OK
estimate phat=0.145, CI = [0.123, 0.167] OK
estimate phat=0.155, \(\mathrm{CI}=[0.133,0.177]\) OK
estimate phat=0.138, CI = [0.117, 0.159] OK
estimate phat=0.145, CI = [0.123, 0.167] OK
estimate phat=0.149, CI = [0.127, 0.171] OK
estimate phat=0.136, CI = [0.115, 0.157] OK
estimate phat=0.142, CI = [0.120, 0.164] OK
estimate phat=0.149, CI = [0.127, 0.171] OK
estimate phat=0.147, CI = [0.125, 0.169] OK
estimate phat=0.122, CI = [0.102, 0.142] OK
estimate phat=0.138, CI = [0.117, 0.159] OK

\section*{We were cheap}

We approximated many times:
- sample without replacement \(\approx\) binomial sample
- binomial distribution \(\approx\) normal distribution
- standard deviation estimated from sample

All of these are reasonable under mild assumptions.
Discuss with your neighbor when they might fail.
See exercises and reading material for improved formulas.

\section*{Why not take a bigger sample?}
\[
\text { MOE }=1.96 \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}
\]

Look at the formula.
We can decrease the margin of error by increasing \(n\).
Suppose we took a sample of size \(n=1000\), but we want to decrease MOE by a factor of ten.

What is our new sample size?


\title{
Oh, and the other errors... we only covered sampling error!
}


Do now: Explain the pink "error arrows" to your neighbor, and discuss. (1 min)

\section*{Further examples}
1. The population is \(N=1000\) and the sample is \(\boldsymbol{n}=\boldsymbol{N}\). What does our basic formula say? Explain why the result does/doesn't make sense. Then read about finite sample correction and repeat using that.
2. We have a large population and a sample of \(n=100\).

We observe 0 supporters of party X . Calculate the CI ,
(a) by our basic formula,
(b) exactly, by considering the binomial distribution, and
(c) by a computer simulation.

Compare and explain the results.

\section*{Questions?}


I DON'T KNOW HOW TO PROPAGATE ERROR CORRECTLY, SO I JUST PUT ERROR BARS ON ALL MY ERROR BARS.```

