Lecture 14: Finish 8.2 Sample Mean; 9.1 Inference for Categorical Variable: Confidence Intervals

- 68-95-99.7 Rule; Checking Assumptions
- 3 Forms of Inference
- Probability vs. Confidence
- Constructing Confidence Interval
- Sample Size; Level of Confidence

Looking Back: *Review*

- **4 Stages of Statistics**
  - Data Production (discussed in Lectures 1-3)
  - Displaying and Summarizing (Lectures 3-8)
  - Probability
    - Finding Probabilities (discussed in Lectures 9-10)
    - Random Variables (discussed in Lectures 10-12)
    - Sampling Distributions
      - Proportions (discussed in Lecture 13)
      - Means (began in Lecture 13)
  - Statistical Inference

68-95-99.7 Rule for Normal R.V. *(Review)*

Sample at random from normal population; for sampled value X (a R.V.), probability is

- 68% that X is within 1 standard deviation of mean
- 95% that X is within 2 standard deviations of mean
- 99.7% that X is within 3 standard deviations of mean

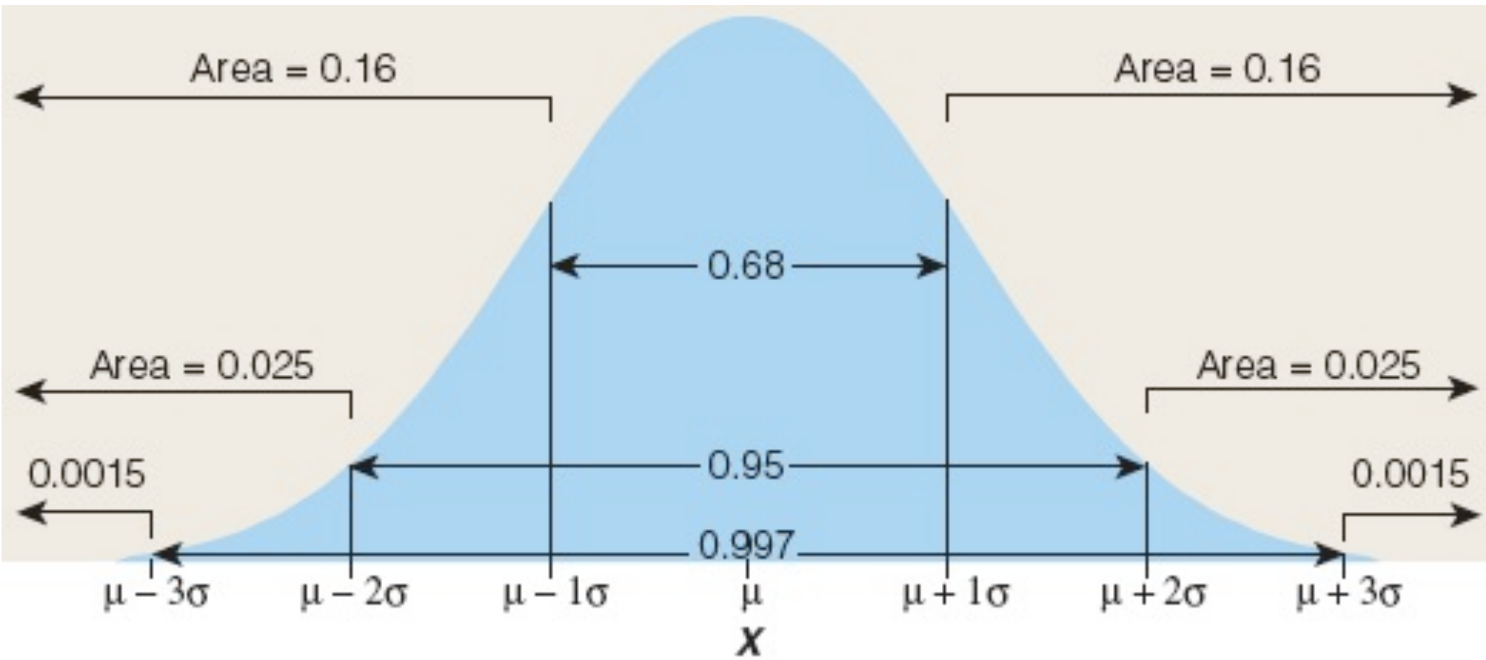

68-95-99.7 Rule for Sample Mean

For sample means $\bar{X}$ taken at random from large population with mean μ, s.d. σ, probability is

- 68% that $\bar{X}$ is within $1\frac{\sigma}{\sqrt{n}}$ of μ
- 95% that $\bar{X}$ is within $2\frac{\sigma}{\sqrt{n}}$ of μ
- 99.7% that $\bar{X}$ is within $3\frac{\sigma}{\sqrt{n}}$ of μ

These results hold only if n is large enough.

Example: *68-95-99.7 Rule for 8 Dice*

- **Background**: Population of dice rolls has $\mu = 3.5,\ \sigma = 1.7$. For random samples of size 8, sample mean roll $\bar{X}$ has mean 3.5, standard deviation 0.6, and shape fairly normal.
- **Question:** What does 68-95-99.7 Rule tell us about the behavior of $\bar{X}$?
- **Response:** The probability is approximately
  - 0.68 that $\bar{X}$ is within ______ of _____: in (2.9, 4.1)
  - 0.95 that $\bar{X}$ is within ______ of _____: in (2.3, 4.7)
  - 0.997 that $\bar{X}$ is within ______ of _____: in (1.7, 5.3)

Practice: 8.76 p.382

Typical Problem about Mean *(Review)*

The numbers 1 to 20 have mean 10.5, s.d. 5.8. If numbers picked "at random" by sample of 400 students has mean 11.6, does this suggest bias in favor of higher numbers?

Solution Method: Assume (temporarily) that population mean is 10.5, find probability of sample mean as high as 11.6. If it's too improbable, we won't believe population mean is 10.5; we'll conclude there *is* bias in favor of higher numbers.

Example: *Establishing Behavior of* $\bar{X}$

- **Background**: We asked the following: "*The numbers 1 to 20 have mean 10.5, s.d. 5.8. If numbers picked 'at random' by 400 students have mean* $\bar{x}$ *=11.6, does this suggest bias in favor of higher numbers?*"
- **Question:** What are the mean, standard deviation, and shape of the R.V. $\bar{X}$ in this situation?
- **Response:** For μ= 10.5, σ=5.8, and n=400, $\bar{X}$ has
  - mean ____
  - standard deviation ___________
  - shape ______________________________

Practice: 8.17 p.366

Example: *Testing Assumption About* μ

- **Background**: Sample mean number picked at random from 1 to 20 by 400 students should have mean 10.5, s.d. 0.3.
- **Questions:** Is 11.6 improbably high for $\bar{X}$? Does a sample mean of 11.6 convince us of bias in favor of higher numbers?
- **Responses:** 11.6 is ____ above ____, more than 3 s.d.s. The probability of being this high (or higher) is ______________. Since this is extremely improbable, we _____ believe μ =10.5. Apparently, there______ bias in favor of higher numbers.

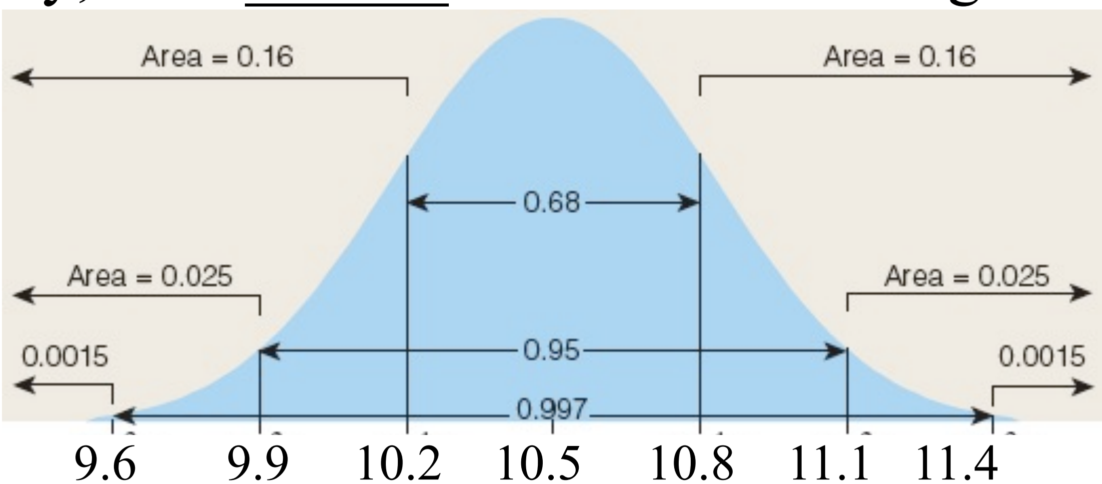

Practice: 8.32 p.369

Example: *Behavior of Individual vs. Mean*

- **Background**: IQ scores are normal with mean 100, s.d. 15.
- **Question:** Is 88 unusually low for…
  - IQ of a randomly chosen individual?
  - Mean IQ of 9 randomly chosen individuals?
- **Response:**
  - IQ X of a randomly chosen individual has mean 100, s.d. 15. For x=88, z = ________________: not even 1 s.d. below the mean→ __________________
  - Mean IQ $\bar{X}$ of 9 randomly chosen individuals has mean 100, s.d. ___________. For $\bar{x}$=88, z = ______________: unusually low (happens less than _____ of the time, since _____________).

Example: *Checking Assumptions*

- **Background**: Household size X in the U.S. has mean 2.5, s.d. 1.4.
- **Question:** Is 3 unusually high for…
  - Size of a randomly chosen household?
  - Mean size of 10 randomly chosen households?
  - Mean size of 100 randomly chosen households?
- **Response:**
  - ______________________________
  - ______________________________
  - ______________________________ n=100 large→ $\bar{X}$ normal; mean 2.5, s.d. $\frac{1.4}{\sqrt{100}} = 0.14$ so $\bar{x} = 3$ has z = (3-2.5)/0.14 = +3.57: unusually high.

Practice: 8.25b-c p.368

Looking Back: *Review*

- **4 Stages of Statistics**
  - Data Production (discussed in Lectures 1-3)
  - Displaying and Summarizing (Lectures 3-8)
  - Probability (discussed in Lectures 9-13)
  - Statistical Inference
    - 1 categorical
    - 1 quantitative
    - categorical and quantitative
    - 2 categorical
    - 2 quantitative

Four Processes of Statistics

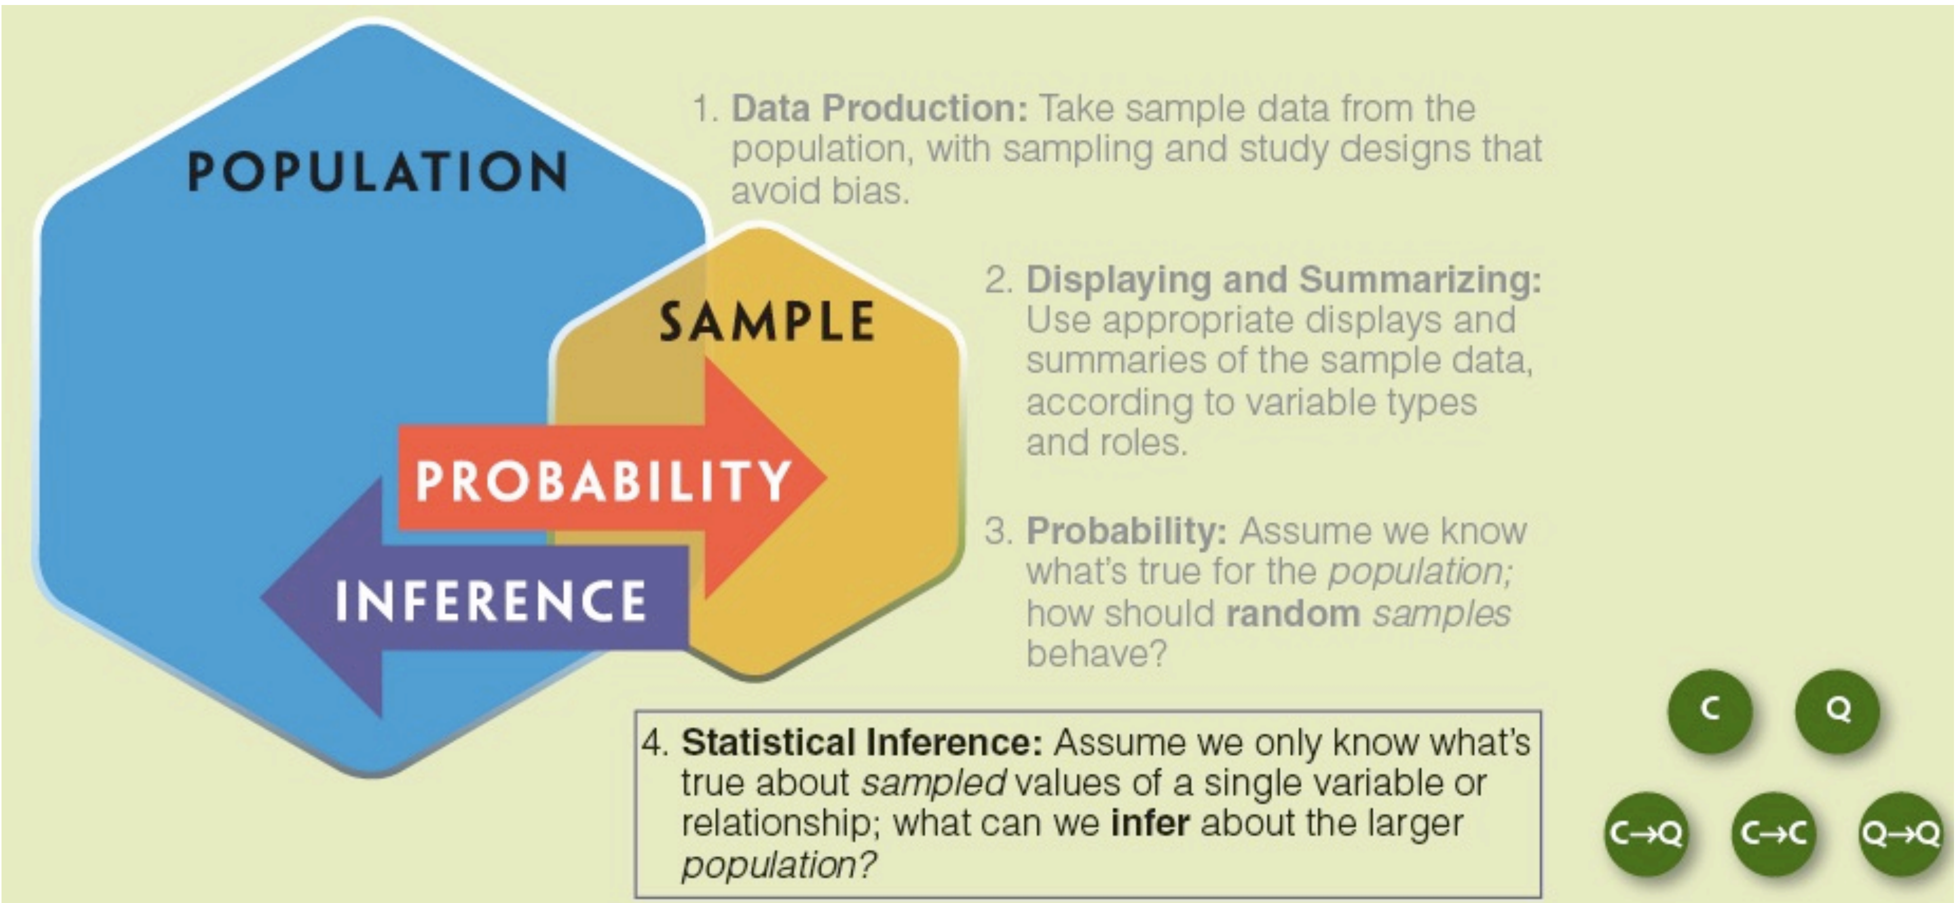

Summarizing Categorical Sample Data (*Review*)

What proportion of sampled students ate breakfast the day of the survey? $\hat{p} = \frac{X}{n} = \frac{246}{446} = 0.55$

Looking Back: *In Part 2, we summarized **sample** data for single variables or relationships.*

Looking Ahead: *In Part 4, our goal is to go beyond sample data and draw conclusions about the larger **population** from which the sample was obtained.*

Three Types of Inference Problem

In a sample of 446 students, 0.55 ate breakfast.

1. What is our best guess for the population proportion of students who eat breakfast?

   Point Estimate

2. What interval should contain the population proportion of students who eat breakfast?

   Confidence Interval

3. Is the population proportion of students who eat breakfast more than half (50%)?

   Hypothesis Test

Behavior of Sample Proportion *(Review)*

For random sample of size n from population with p in category of interest, sample proportion $\hat{p} = \frac{X}{n}$ has

- mean p

→ $\hat{p}$ is *unbiased estimator* of p

(sample must be random)

Example: *Checking if Estimator is Unbiased*

- **Background**: Survey produced sample proportion of intro stat students (various ages and times of day) at a university who'd eaten breakfast.
- **Questions:**
  - Is the sample representative of *all* college students? All students at that university?
  - Were the values of the variable (breakfast or not) recorded without bias?
- **Responses:**
  - Differences among college cafeterias, etc.→ ______________________________
  - Question not sensitive→ ______________________________

Practice: 9.6 p.407

Example: *Point Estimate for p*

- **Background**: In a representative sample of students, 0.55 ate breakfast.
- **Question:** What is our best guess for the proportion of all students at that university who eat breakfast?
- **Response:** $\hat{p}$ unbiased estimator for p→ ______ is best guess for p

Example: *Point Estimate Inadequate*

- **Background**: Our best guess for *p*, population proportion eating breakfast, is sample proportion 0.55.
- **Questions:**
  - Are we pretty sure the population proportion is 0.55?
  - By approximately what amount is our guess "off"?
  - Are we pretty sure population proportion is > 0.50?
- **Responses:**
  - ____
  - ________________________
  - ________________________

Three Types of Inference Problem

In a sample of 446 students, 0.55 ate breakfast.

1. What is our best guess for the population proportion of students who eat breakfast?

   Point Estimate

2. What interval should contain the population proportion of students who eat breakfast?

   Confidence Interval

3. Do more than half (50%) of the population of students eat breakfast?

   Hypothesis Test

Beyond a Point Estimate

Sample proportion from unbiased sample is best estimate for population proportion.

Looking Ahead: *For point estimate we don't need* ***sample size*** *or info about* ***spread****. These are required for confidence intervals and hypothesis tests, to quantify how good our point estimate is.*

Probability vs. Confidence

- **Probability:** given population proportion, how does sample proportion behave?
- **Confidence:** given sample proportion, what is a range of plausible values for population proportion?

Turning Point!

Example: *Probability Statement*

- **Background**: If students pick numbers from 1 to 20 at random, p=0.05 should pick #7. For n=400, $\hat{p}$ has
  - mean 0.05
  - standard deviation $\sqrt{\frac{0.05(1-0.05)}{400}} = 0.01$
  - shape approximately normal
- **Question:** What does the "95" part of the 68-95-99.7 Rule tell us about $\hat{p}$?
- **Response:** Probability is approximately 0.95 that $\hat{p}$ falls within ____________ of ________.

Looking Ahead: *This statement about **sample proportion** is correct but not very useful for practical purposes. In most real-life problems, we want to draw conclusions about an **unknown population proportion.***

Practice: 8.57a p.376

Example: *How Far is One from the Other?*

- **Background**: An instructor can say about his/her position in the classroom: "I'm within 10 feet of this particular student."
- **Question:** What can be said about where that student is in relation to the instructor?
- **Response:**

Definitions

Margin of Error: *Distance* around a sample statistic, within which we have reason to believe the corresponding parameter falls.

A common margin of error is 2 s.d.s.

Confidence Interval for parameter: *Interval* within which we have reason to believe the parameter falls = **range of plausible values**

A common confidence interval is sample statistic plus or minus 2 s.d.s.

> ***A Closer Look:*** *A parameter is* ***not*** *a R.V. It does* ***not*** *obey the laws of probability so we must use the word "confidence".*

Example: *Confidence Interval for p*

- **Background**: 30/400=0.075 students picked #7 "at random" from 1 to 20. Let's assume sample proportion for *n*=400 has s.d. 0.01.
- **Question:** What can we claim about population proportion *p* picking #7?
- **Response:** We're pretty sure *p* is

____

Looking Back: *In Part I, we learned about biased samples. The data suggest p>0.05: students were apparently biased in favor of #7. Their selections were* ***haphazard****, not random. If sampling individuals or assigning them to experimental treatments is not randomized, then we produce a confidence interval that is not centered at p.*

Practice: 9.7 p.407

Level of Confidence Corresponds to Multiplier

By "pretty sure", we mean "95% confident", because 95% is the probability of sample proportion within 2 s.d.s of p (for large enough n).

Looking Back: *Our probability statement claimed sample proportion should fall within 2 s.d.s of population proportion. Now, the inference statement claims population proportion should be within 2 s.d.s of sample proportion.*

Probability Interval for $\hat{p}$ Picking #7

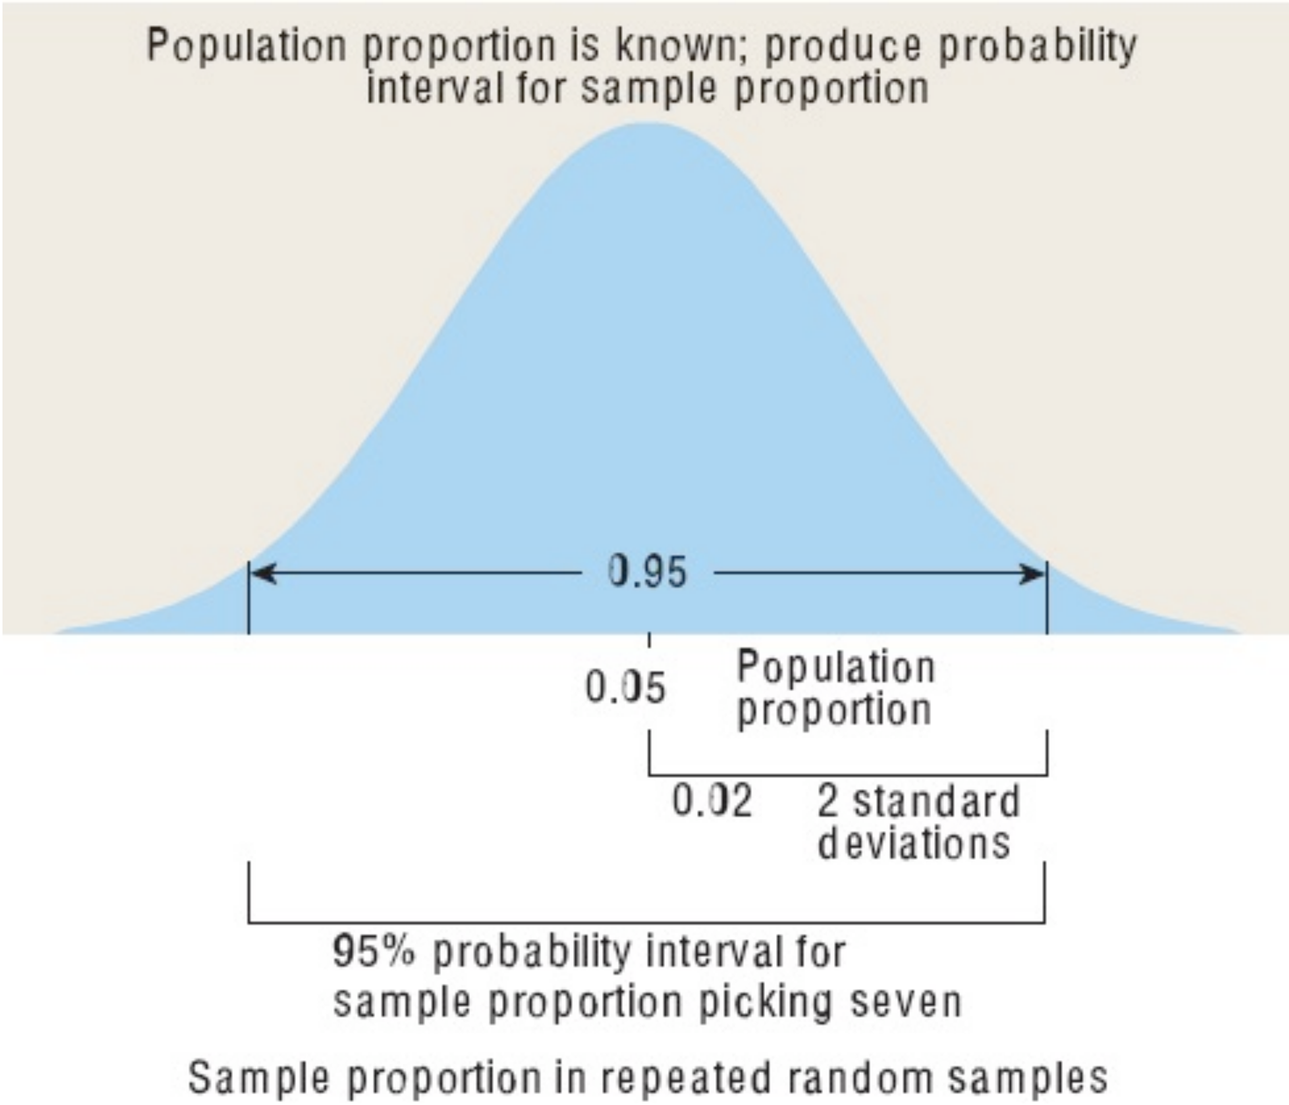

Confidence Interval for p Picking #7

We do not sketch a curve showing probabilities for population proportion because it is ***not a random variable****.*

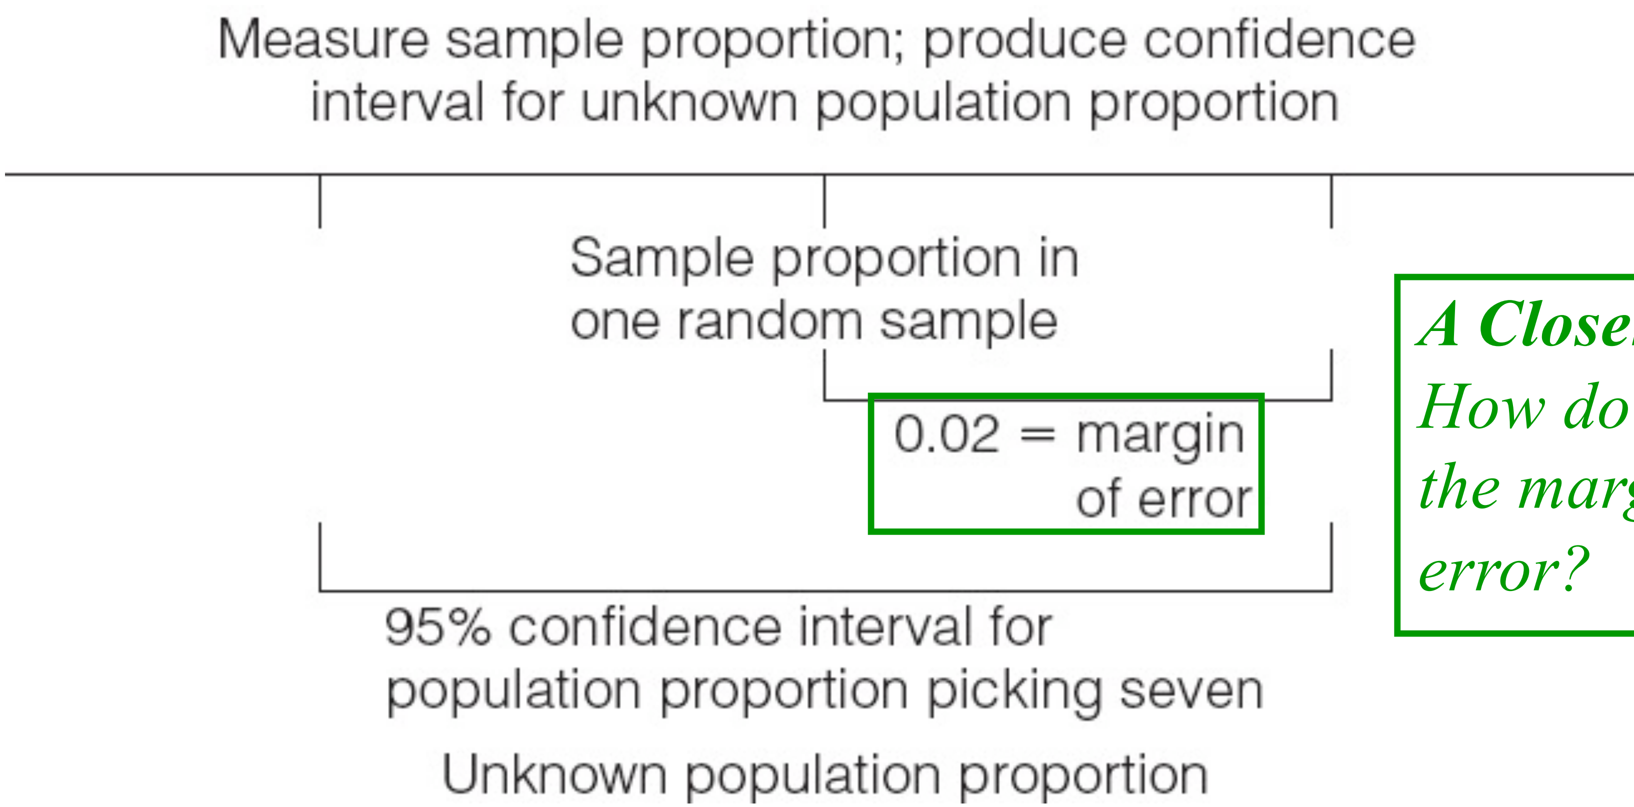

A Closer Look:
How do we know the margin of error?

Behavior of Sample Proportion *(Review)*

For random sample of size n from population with p in category of interest, sample proportion $\hat{p}$ has

- mean p
- standard deviation $\sqrt{\frac{p(1-p)}{n}}$

We do inference because p is unknown; how can we know the standard deviation, which involves p?

Definition

Standard error: estimated standard deviation of a sampling distribution.

We estimate standard deviation of $\hat{p}$ with standard error $\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$.

Looking Ahead: *In many situations throughout inference, when needed information about the* ***population*** *is unknown, we substitute known information about the* ***sample****.*

Definition

95% confidence interval for *p*: (approx.)

$$\hat{p} \pm 2\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

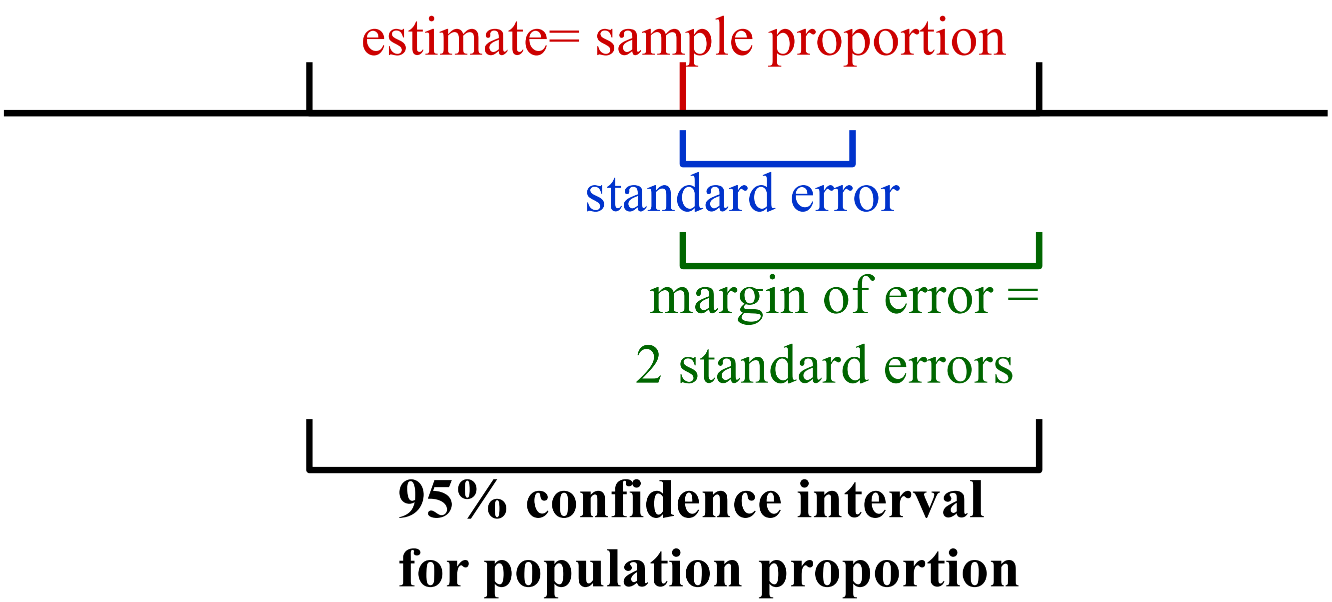

Confidence Interval Formula: Conditions

95% confidence interval for *p*: (approx.)

$$\hat{p} \pm 2\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

- Sample must be unbiased
  (otherwise interval is not really centered at $\hat{p}$)
- n must be large enough so $\hat{p}$ is approx. normal
  (otherwise multiplier 2 from 68-95-99.7 Rule is incorrect)
- Population size must be at least $10n$
  (otherwise formula for s.d., which requires independence, is incorrect)

Conditions for Normality in Confidence Interval

Multiplier 2 from normal dist. approximately correct if np and $n(1-p)$ both at least 10.

But p is unknown so substitute $\hat{p}$:

Require

$$n\hat{p} = nX/n = \boxed{X \geq 10}$$

$$n(1-\hat{p}) = n - nX/n = \boxed{n - X \geq 10}$$

Sample count in (X) and out (n-X) of category of interest should both be at least 10.

Example: *Checking Sample Size*

- **Background**: 30/400=0.075 students picked #7 "at random" from 1 to 20.
- **Question:** Do the data satisfy requirement for approximate normality of sample proportion?
- **Response:**

Example: *Checking Population Size*

- **Background**: To draw conclusions about criminal histories of a city's 750 bus drivers, a random sample of 100 drivers was used.
- **Question:** Is there approximate independence in spite of sampling without replacement, so formula for standard error is accurate?
- **Response:**

Example: *Revisiting Original Question*

- **Background**: In sample of 446 college students, 246 (proportion 0.55) ate breakfast.
- **Question:** Assuming sample is representative, what interval should contain proportion of all students at that university who eat breakfast?
- **Response:** Approx. 95% confidence interval for *p* is

$$\hat{p} \pm 2\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}$$

Looking Back:
Earlier we wondered if a majority of students eat breakfast. The interval suggests this is the case, since it is entirely above 0.50.

Practice: 9.9a,c p.407

Example: *Role of Sample Size*

- **Background**: 95% confidence intervals based on sample proportion 0.54 from various sample sizes:

Sample Size n	Standard Error of $\hat{p}$	Margin of Error	95% Confidence Interval
50	$\sqrt{\frac{0.54(1-0.54)}{50}} = 0.070$	$2(0.070) = 0.14$	(0.40, 0.68)
200	$\sqrt{\frac{0.54(1-0.54)}{200}} = 0.035$	$2(0.035) = 0.07$	(0.47, 0.61)
1,000	$\sqrt{\frac{0.54(1-0.54)}{1,000}} = 0.016$	$2(0.016) = 0.03$	(0.51, 0.57)

- **Question:** What happens as n increases?
- **Response:**

Practice: 9.17b p.409

Example: *A Common Margin of Error*

- **Background**: Pollsters most often report a 3% error margin.

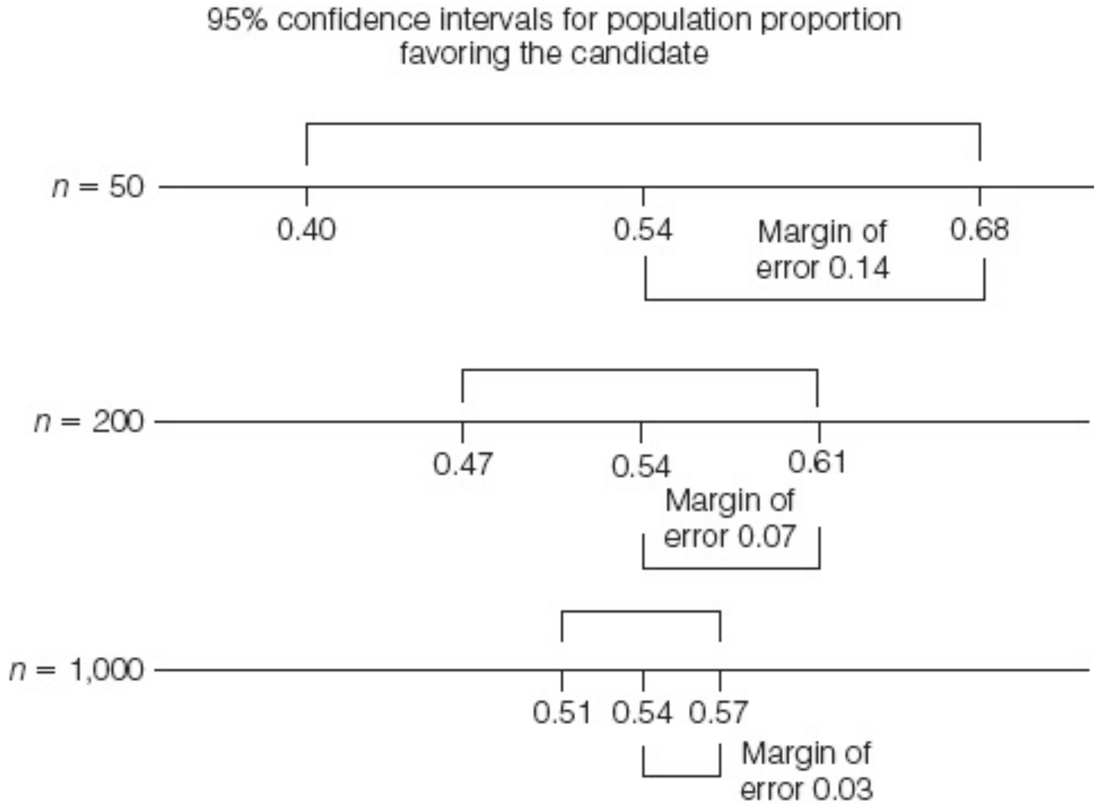

- **Question:** What is the most common sample size for polls?
- **Response:** Approximately ________.

Practice: 9.17c p.409

Other Levels of Confidence

Confidence level 95% uses multiplier 2. Other levels use other multipliers, based on normal curve:

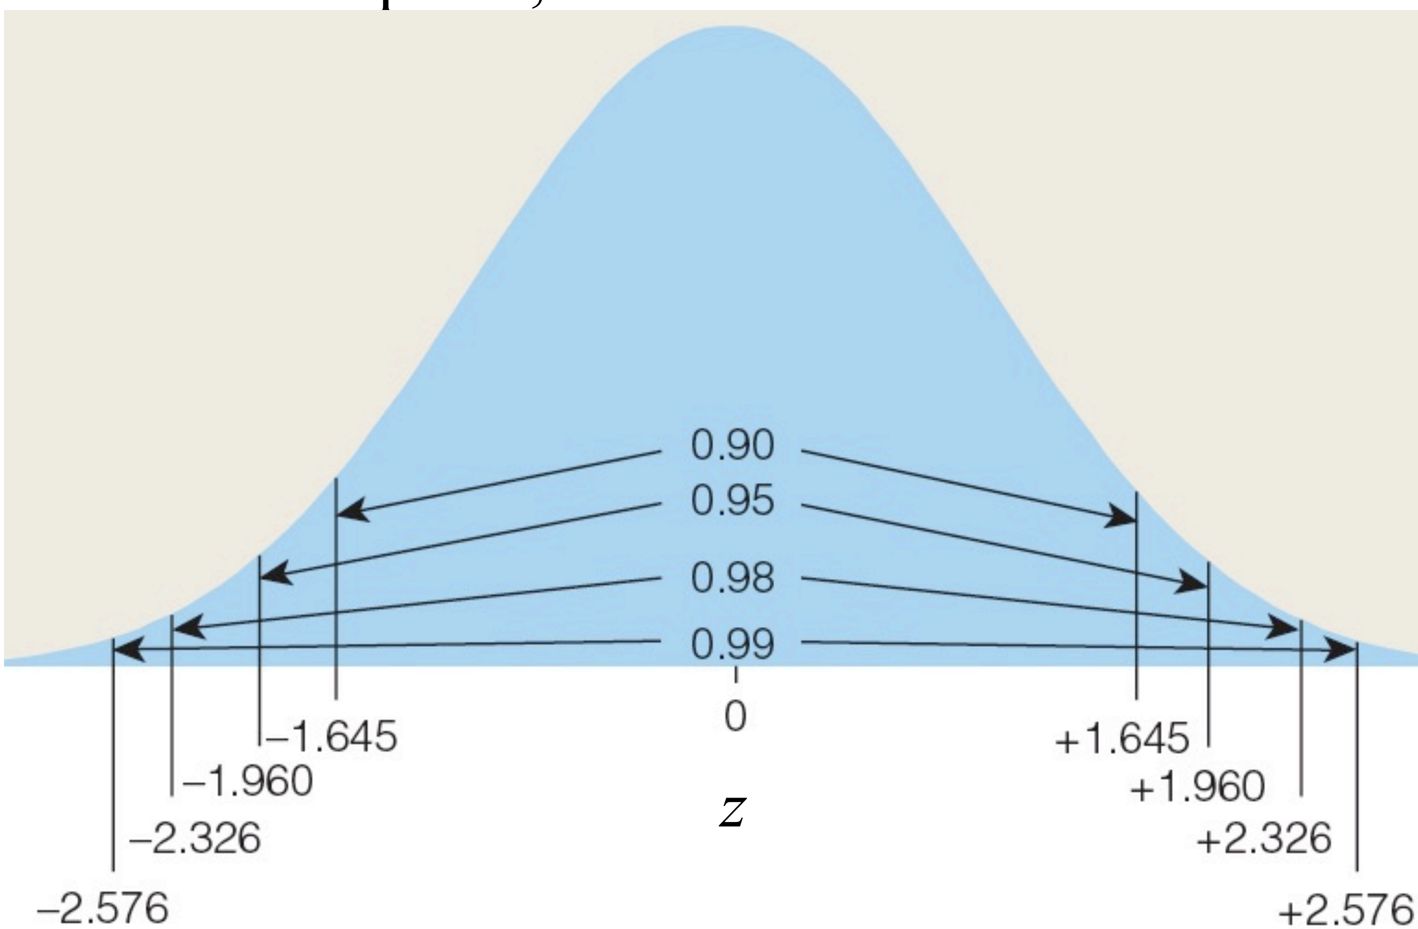

Other Levels of Confidence

Confidence level 95% uses multiplier 2. Other levels use other multipliers, based on normal curve.

More precise multiplier for 95% is 1.96 instead of 2.

Level	Multiplier
90%	1.645
95%	1.960
98%	2.326
99%	2.576

Example: *Other Levels of Confidence*

- **Background**: Of 108 students in committed relationships, 0.70 said they took comfort by sniffing out-of-town partner's clothing. Standard error can be found to be 0.04.
- **Question:** How do 90%, 95%, 98%, 99% confidence intervals compare?
- **Response:**
  - 90% C.I. is ____________________=(0.63, 0.77)
  - 95% C.I. is ____________________=(0.62, 0.78)
  - 98% C.I. is ____________________=(0.61, 0.79)
  - 99% C.I. is ____________________=(0.60, 0.80)

Practice: 9.18a p.409

Example: *Other Levels of Confidence*

Intervals get ______ as confidence level increases:

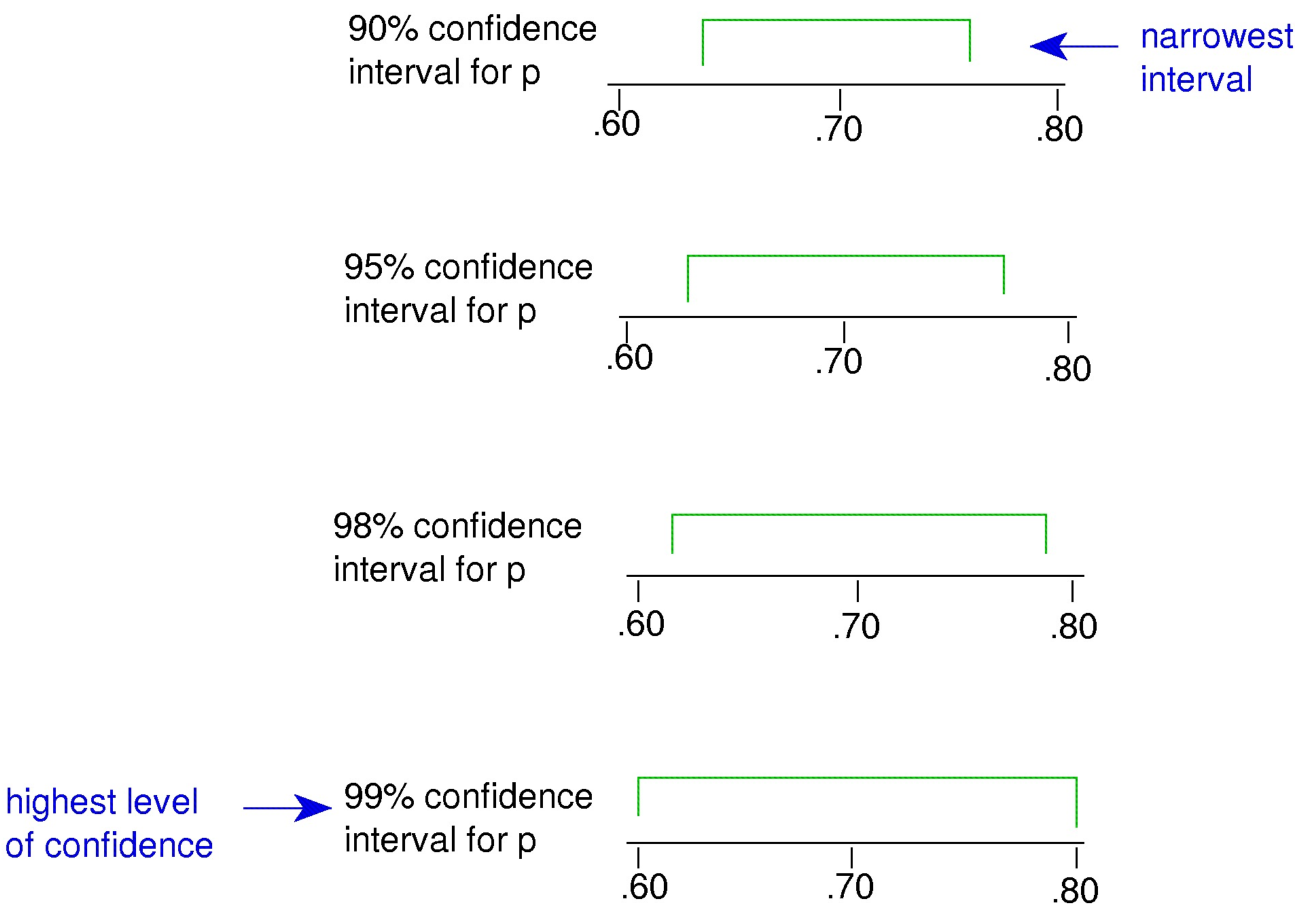

Confidence Interval and Long-Run Behavior

Repeatedly set up 95% confidence interval for proportion of heads, based on 20 coinflips.

In the long run, 95% of the intervals should contain population proportion of heads, 0.5.

Confidence Interval and Long-Run Behavior

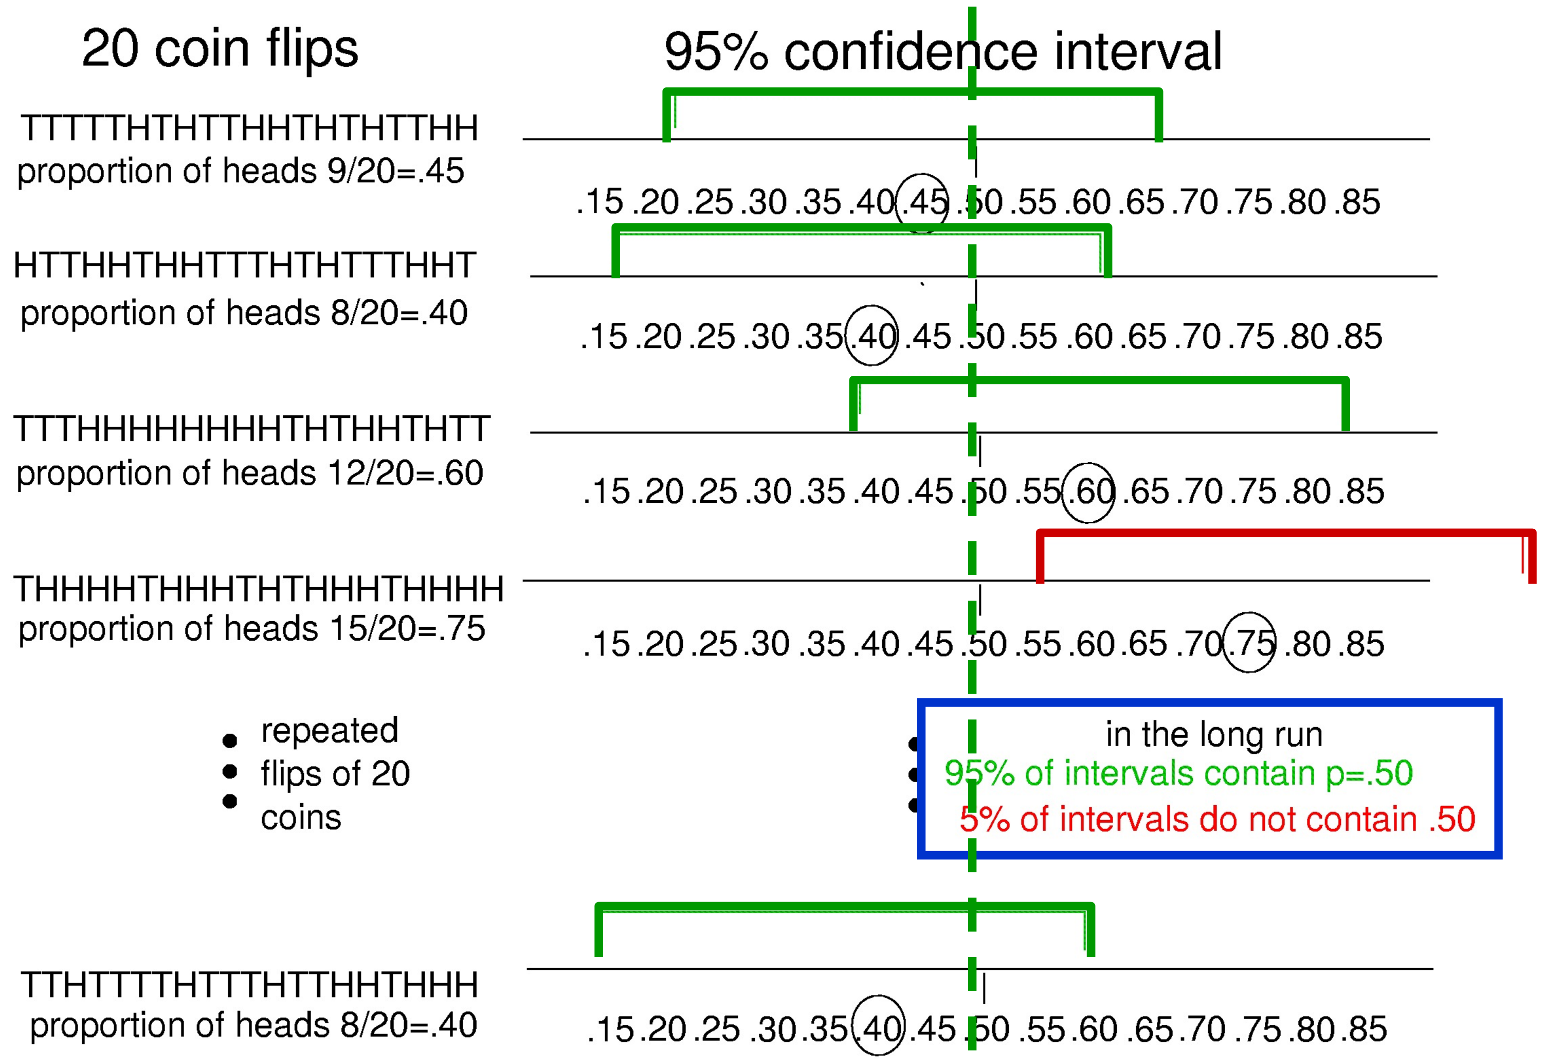

Example: *Confidence in the Long Run*

- **Background**: "President-elect Barack Obama's campaign strategists weren't the only ones vindicated Tuesday. Pollsters came out looking pretty good, too. Of 27 polls of Pennsylvania voters released in the campaign's final two weeks, only seven missed Obama's 10.3-point victory by more than their margins of error. Obama's national victory of about 6 points was within the error margins of 16 of the 21 national polls released in the final week."
- **Question:** Should pollsters be pleased with success rates of 20/27=16/21= 75% ?
- **Response:**

pittsburghlive.com/x/pittsburghtrib/news/cityregion/s_597288.html

Practice: 9.35 p.413

Lecture Summary
(Sampling Distributions; Means)

- 68-95-99.7 Rule for sample mean
  - Revisit typical problem
  - Checking assumptions for use of Rule

Lecture Summary
(Inference for Proportions: Confidence Interval)

- 3 forms of inference; focus on confidence interval
- Probability vs. confidence
- Constructing confidence interval
  - Margin of error based on standard error
  - Conditions
- Role of sample size
- Confidence at other levels
- Confidence interval in the long run