Section 5.2

Confidence Intervals and P-values using Normal Distributions

Outline

- Central limit theorem
- Confidence interval using a normal distribution
- Hypothesis test using a normal distribution

Review

A bootstrap distribution is approximated by the normal distribution N(0.15, 0.03). What is the standard error of the statistic?

a) 0.15
b) 0.3
c) 0.03
d) 0.06

N(mean, sd)

The sd of a bootstrap distribution is the standard error of the statistic.

Central Limit Theorem

For random samples with a sufficiently large sample size, the distribution of sample statistics for a mean or a proportion is normally distributed

Bootstrap and Randomization Distributions

Slope :Restaurant tips

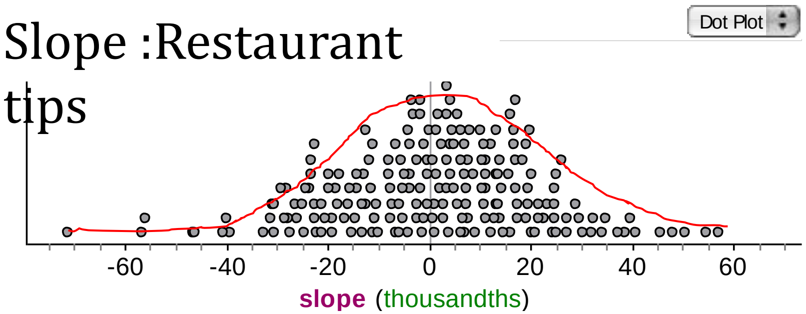

Correlation: Malevolent uniforms

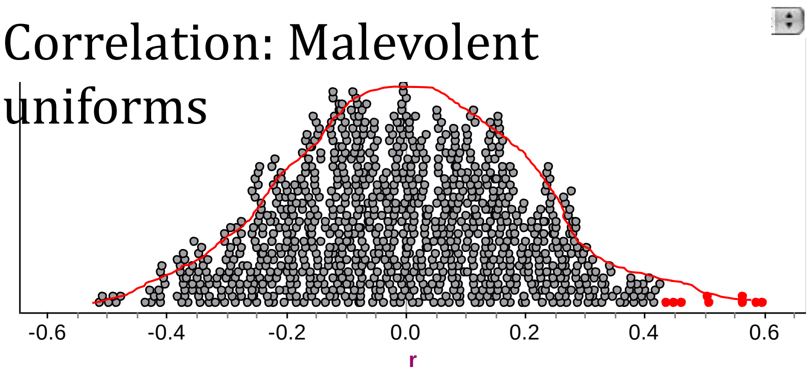

Mean :Body Temperatures

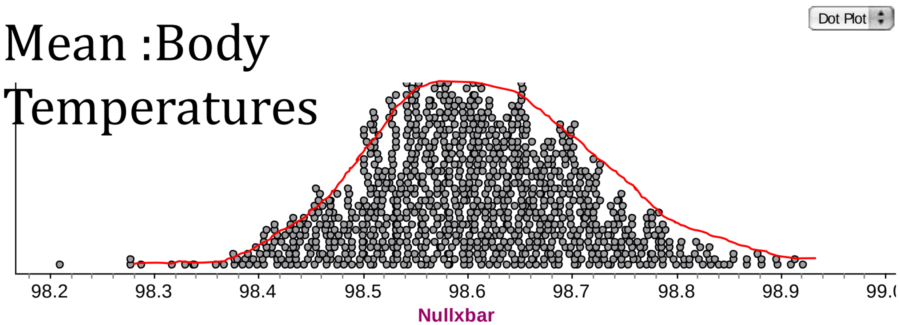

Diff means: Finger taps

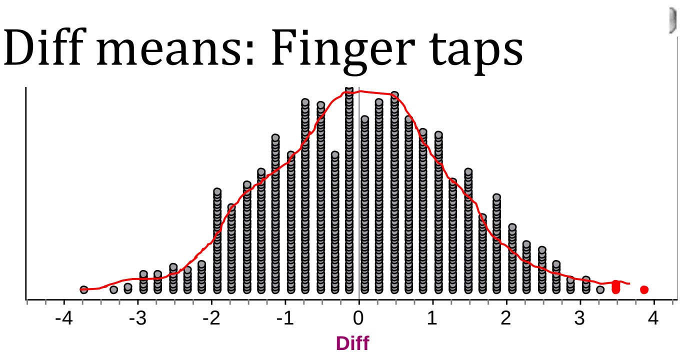

Proportion : Owners/dogs

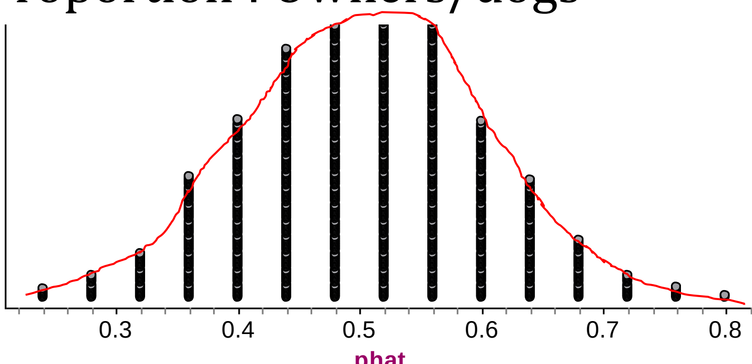

Mean : Atlanta commutes

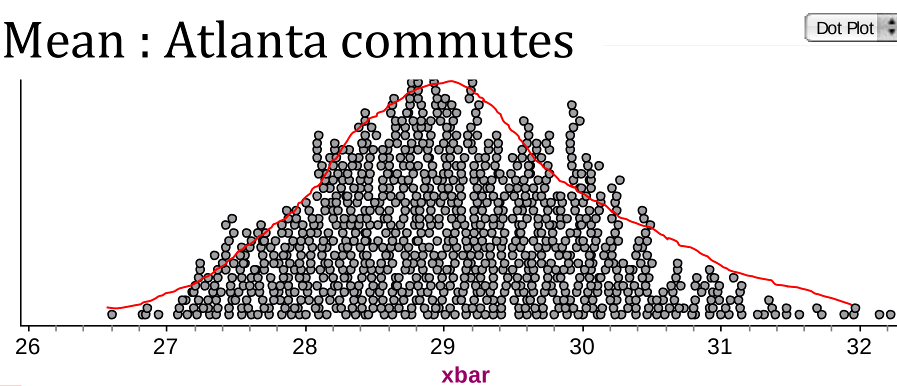

Central Limit Theorem

• The central limit theorem holds for ANY original distribution, although "sufficiently large sample size" varies

• The more skewed the original distribution is, the larger n has to be for the CLT to work

• For quantitative variables that are not very skewed, $n \geq 30$ is usually sufficient

• For categorical variables, counts of at least 10 within each category is usually sufficient

Hearing Loss

- In a random sample of 1771 Americans aged 12 to 19, 19.5% had some hearing loss (this is a dramatic increase from a decade ago!)

- What proportion of Americans aged 12 to 19 have some hearing loss? Give a 95% CI.

Rabin, R. "Childhood: Hearing Loss Grows Among Teenagers," www.nytimes.com, 8/23/10.

Hearing Loss

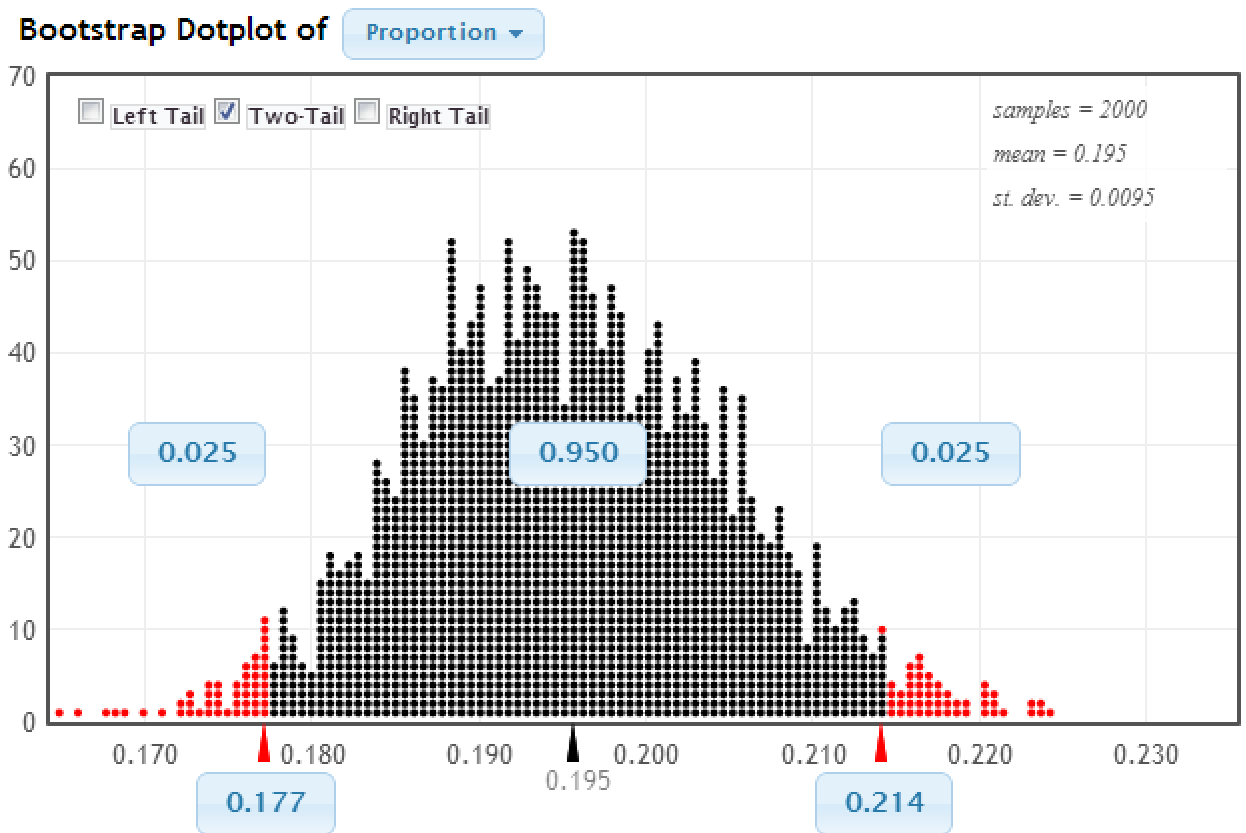

Original Sample

Count	Sample Size	Proportion
345	1771	0.195

Bootstrap Sample

Count	Sample Size	Proportion
307	1771	0.173

$$(0.177, 0.214)$$

Hearing Loss

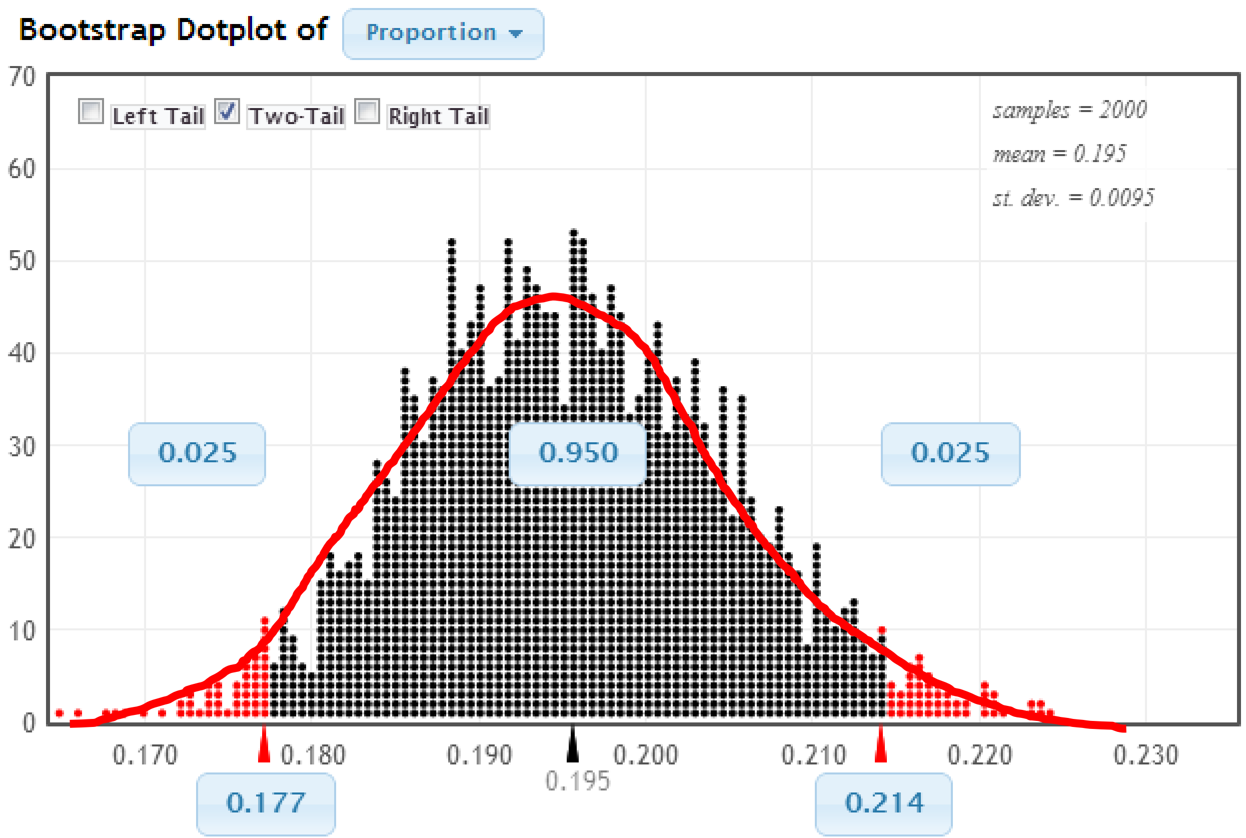

Original Sample

Count	Sample Size	Proportion
345	1771	0.195

Bootstrap Sample

Count	Sample Size	Proportion
307	1771	0.173

Bootstrap Distributions

If a bootstrap distribution is approximately normally distributed, we can write it as

a) N(parameter, sd)
b) N(statistic, sd)
c) N(parameter, se)
d) N(statistic, se)

sd = standard deviation of variable
se = standard error = standard deviation of statistic

Confidence Intervals

If the bootstrap distribution is normal:

To find a P% confidence interval , we just need to find the middle P% of the distribution

$$N(\text{statistic}, \text{SE})$$

Hearing Loss

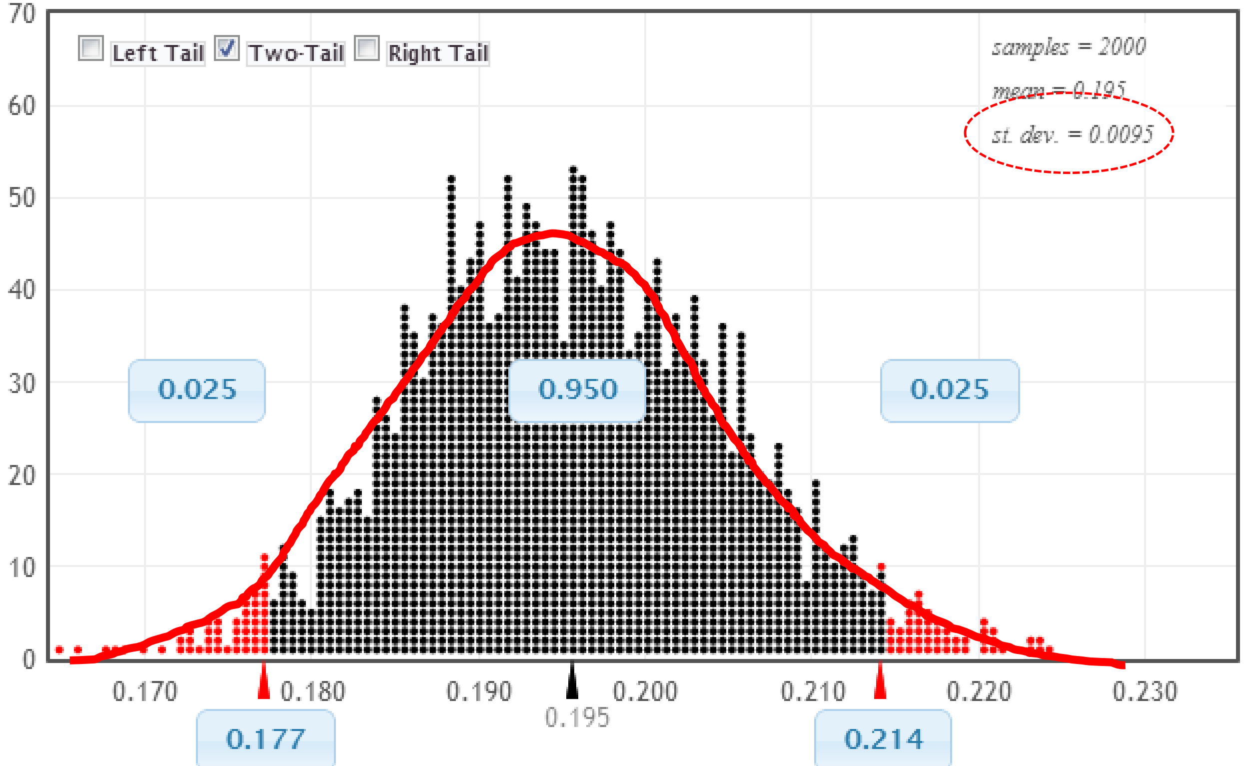

Original Sample

Count	Sample Size	Proportion
345	1771	0.195

Bootstrap Sample

Count	Sample Size	Proportion
307	1771	0.173

N(0.195, 0.0095)

Hearing Loss

www.lock5stat.com/statkey

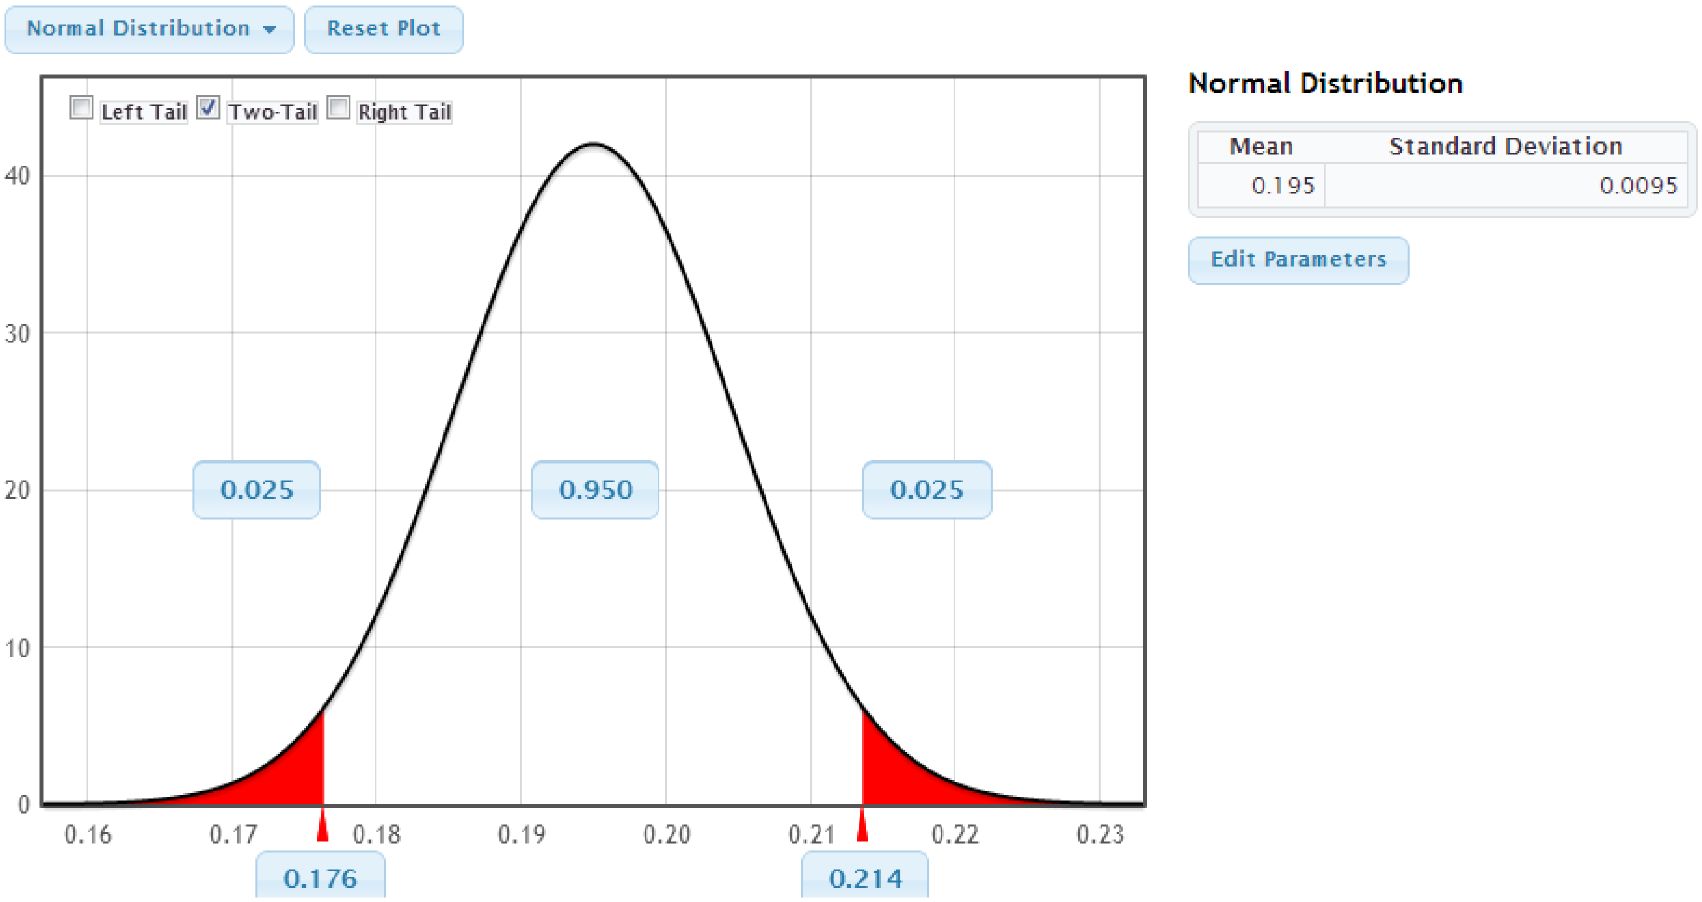

(0.176, 0.214)

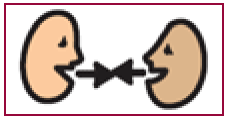

Confidence Intervals

For normal bootstrap distributions, the formula $\textit{statistic} \pm z^* \cdot SE$ also gives a 95% confidence interval.

How would you use the N(0,1) normal distribution to find the appropriate multiplier for other levels of confidence?

Confidence Interval using N(0,1)

If a statistic is normally distributed, we find a confidence interval for the parameter using

$$statistic \pm z^* \cdot SE$$

where the area between $-z^*$ and $+z^*$ in the standard normal distribution is the desired level of confidence.

P% Confidence Interval

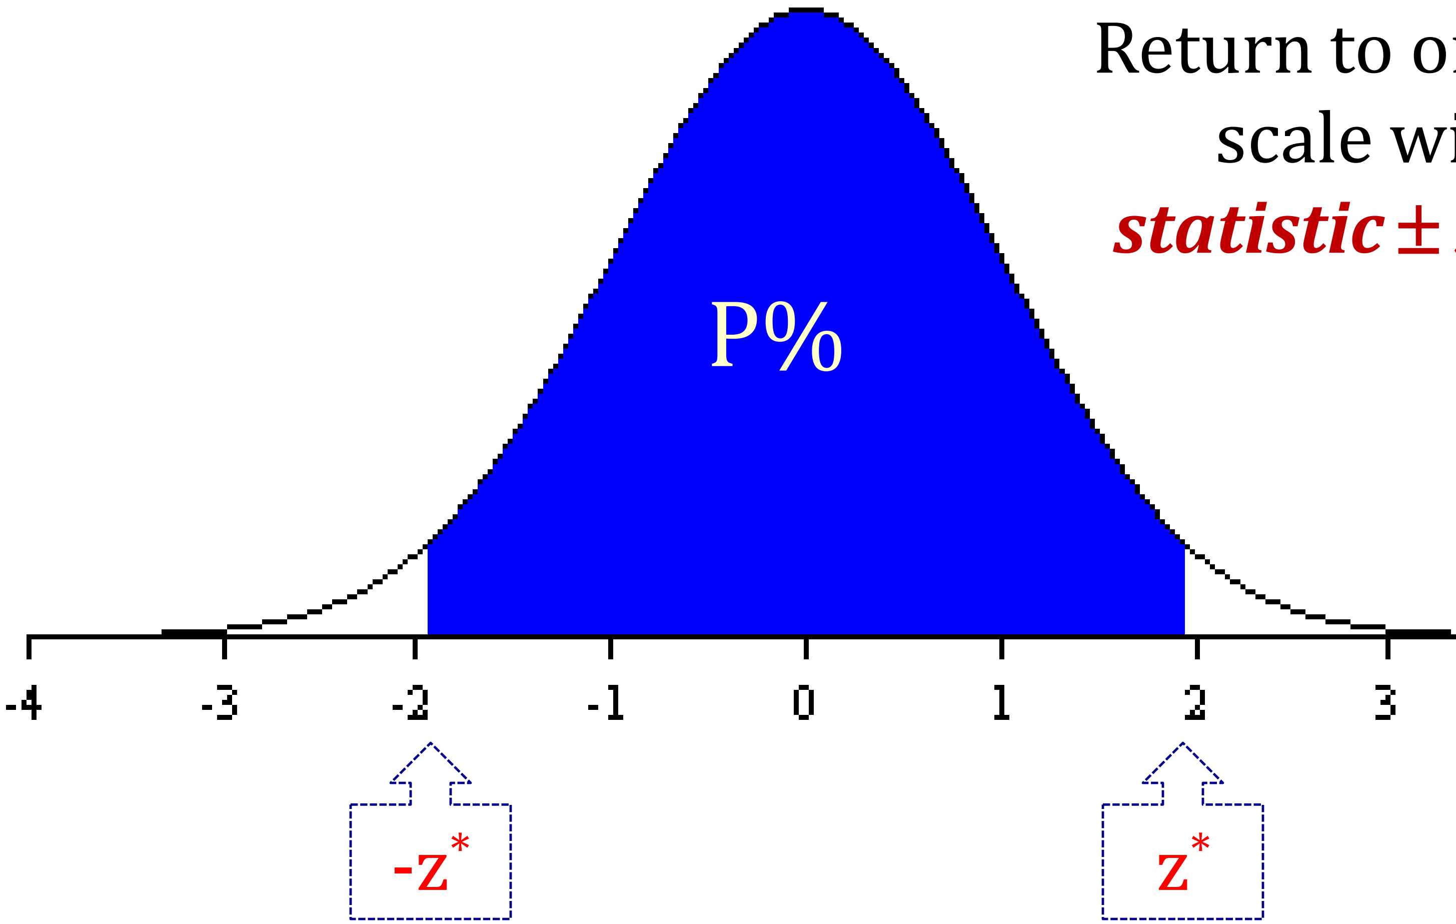

Return to original scale with

statistic* $\pm$ *z** $\cdot$ *SE

Confidence Intervals

Find z^* for a 99% confidence interval.

www.lock5stat.com/statkey

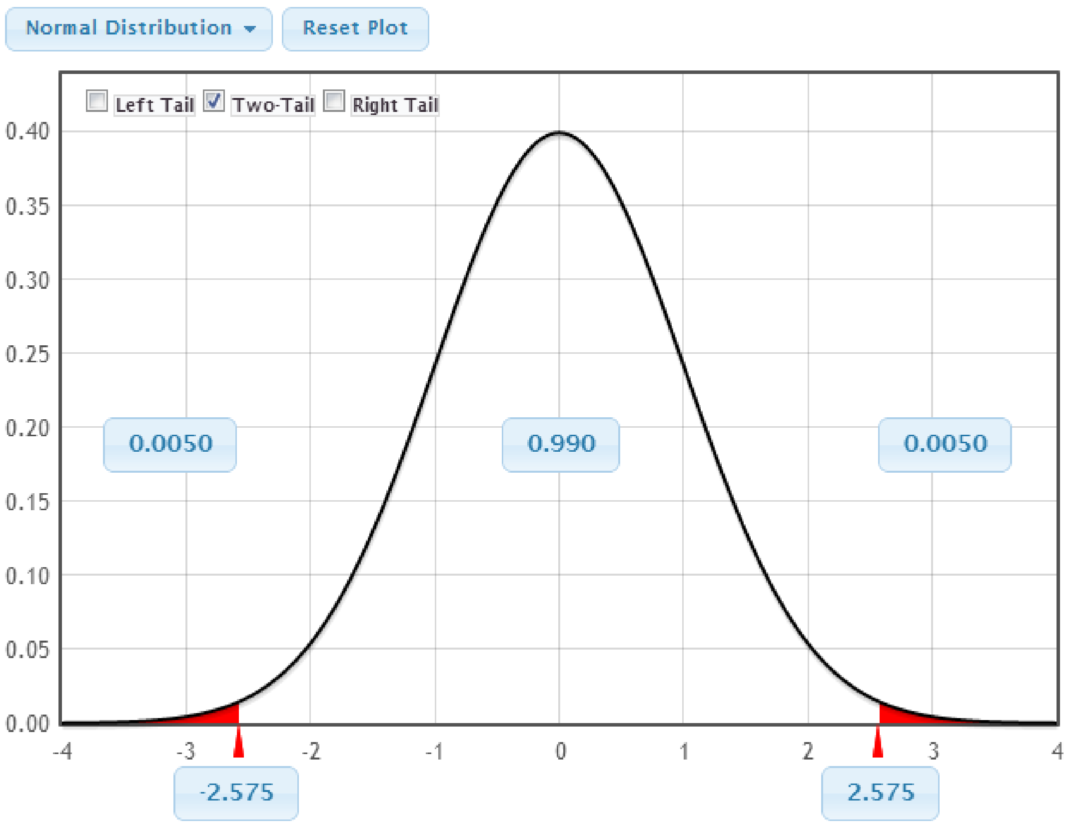

$z^* = 2.575$

Hearing Loss

- Find a 99% confidence interval for the proportion of Americans aged 12-19 with some hearing loss.

$$statistic \pm z^* \cdot SE$$

$$0.195 \pm 2.575 \cdot 0.0095$$

$$(0.171, 0.219)$$

Other Levels of Confidence

www.lock5stat.com/statkey

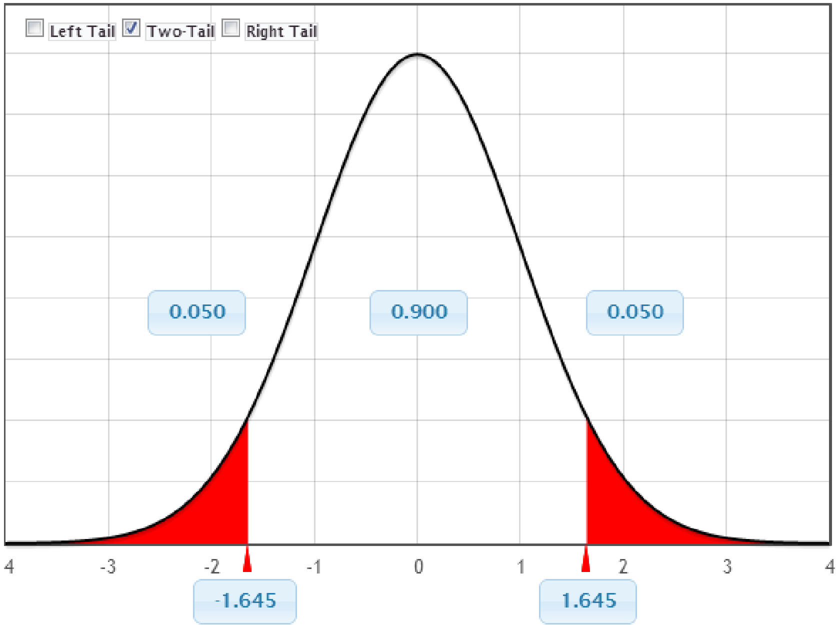

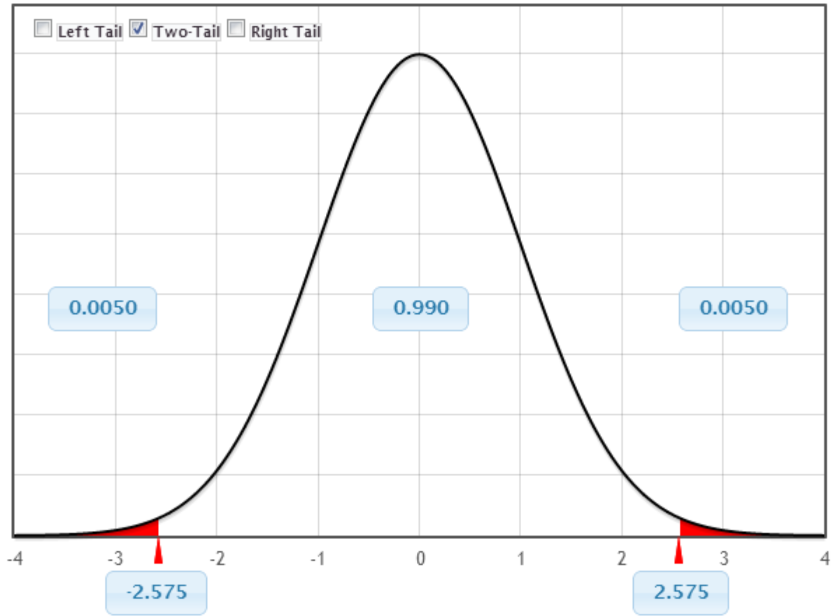

90% Confidence

$z^* = 1.645$

99% Confidence

$z^* = 2.576$

Technically, for 95% confidence, z* = 1.96, but 2 is much easier to remember, and close enough

News Sources

- "A new national survey shows that the majority (64%) of American adults use at least three different types of media every week to get news and information about their local community"
- The standard error for this statistic is 1%
- Find a 90% CI for the true proportion.

$$statistic \pm z^* \cdot SE$$
$$0.64 \pm 1.645 \cdot 0.01$$
$$(0.624, 0.656)$$

Source: http://pewresearch.org/databank/dailynumber/?NumberID=1331

First Born Children

- Are first born children actually smarter?

- Explanatory variable: first born or not
- Response variable: combined SAT score

- Based on a sample of college students, we find $\bar{x}_{first\ born} - \bar{x}_{not\ first\ born} = 30.26$

- From a randomization distribution, we find $SE = 37$

First Born Children

$$\bar{x}_{first\ born} - \bar{x}_{not\ first\ born} = 30.26$$

$$SE = 37$$

What normal distribution should we use to find the p-value?

a) N(30.26, 37)
b) N(37, 30.26)
c) N(0, 37)
d) N(0, 30.26)

Because this is a hypothesis test, we want to see what would happen if the null were true, so the distribution should be centered around the null. The variability is equal to the standard error.

p-values

If the randomization distribution is normal:

To calculate a p-value, we just need to find the area in the appropriate tail(s) beyond the observed statistic of the distribution

$$N(\text{null value, SE})$$

Hypothesis Testing

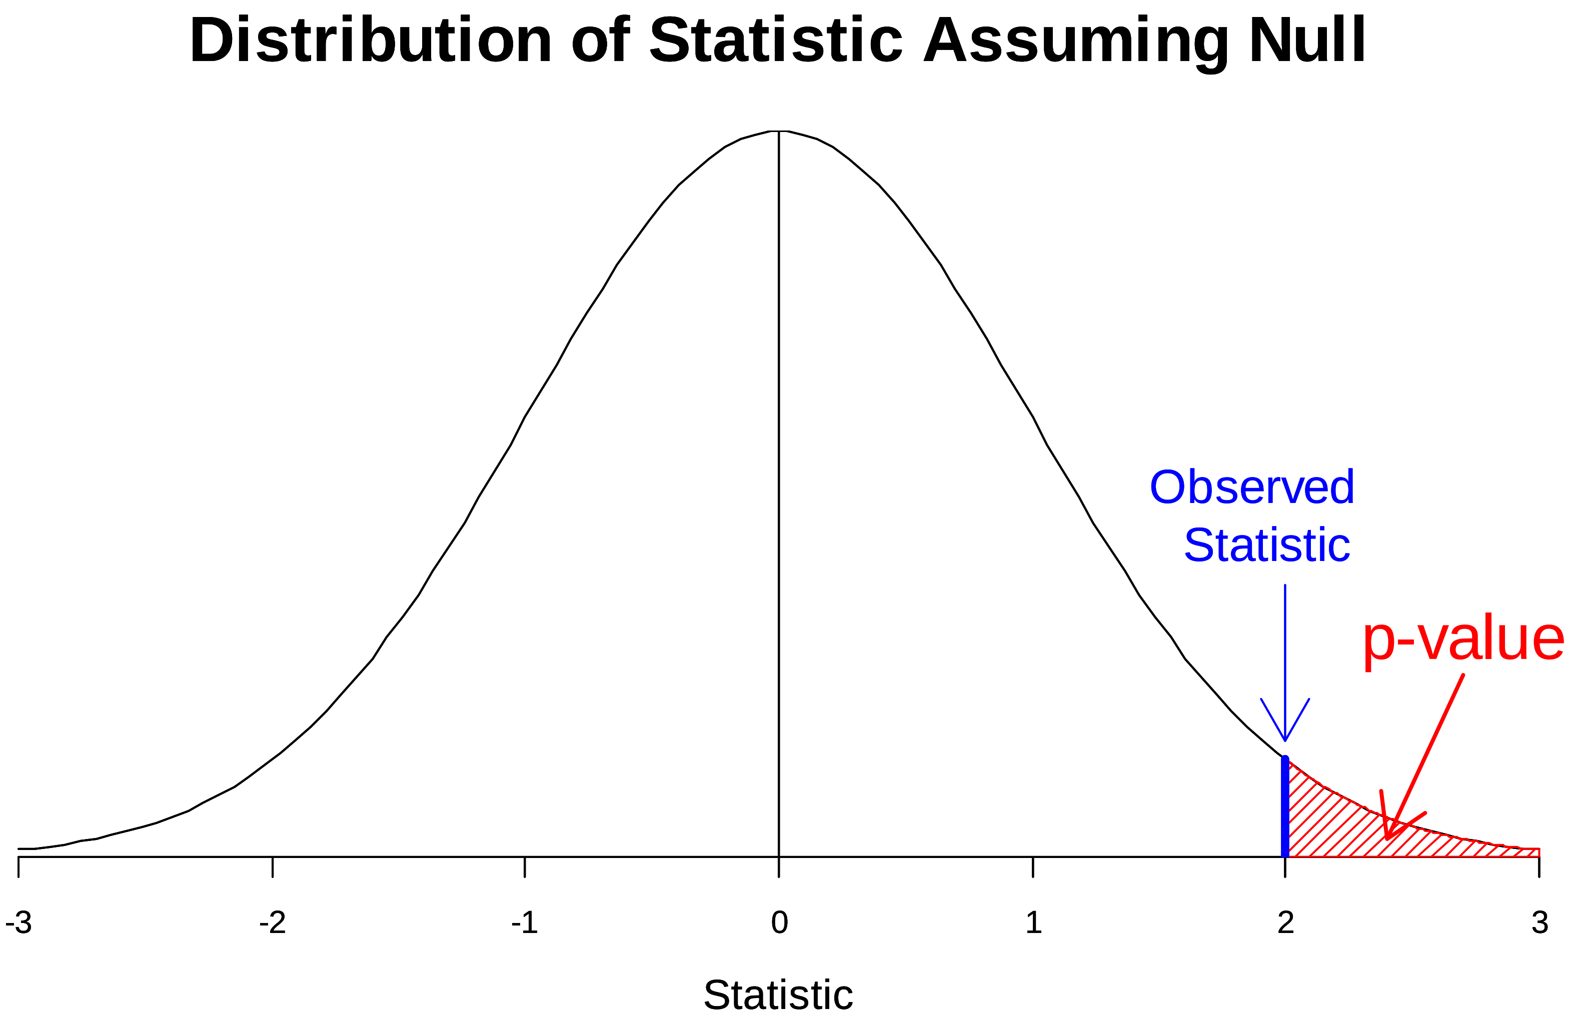

First Born Children

N(0, 37)

www.lock5stat.com/statkey

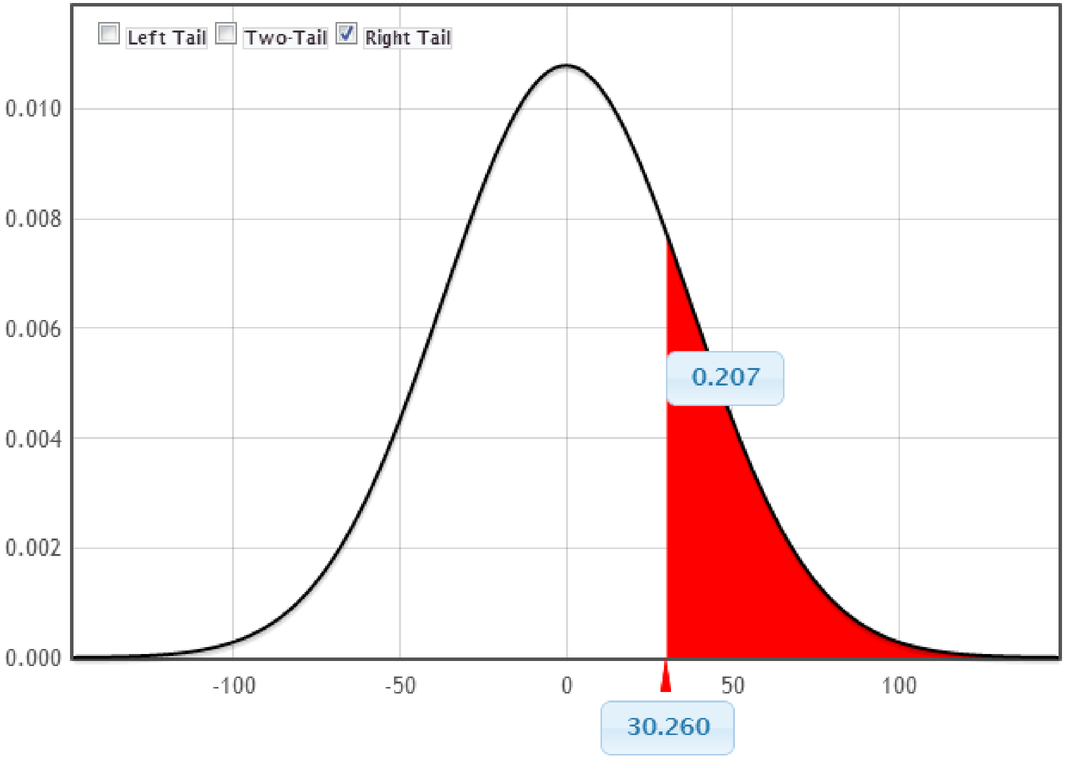

p-value = 0.207

Standardized Test Statistic

The standardized ***test statistic*** is the number of standard errors a statistic is from the null:

$$z = \frac{sample\ statistic - null\ parameter}{SE}$$

- Calculating the number of standard errors a statistic is from the null value allows us to assess extremity on a common scale

p-value using N(0,1)

If a statistic is normally distributed under H_0, the ***p-value*** is the probability a standard normal is beyond

$$z = \frac{sample\ statistic\ - null\ parameter}{SE}$$

First Born Children

$$\bar{x}_{first\ born} - \bar{x}_{not\ first\ born} = 30.26, SE = 37$$

1) Find the standardized test statistic

$$z = \frac{statistic - null}{SE} = \frac{30.26 - 0}{37} = 0.818$$

2) Compute the p-value

p-value $= 0.207$

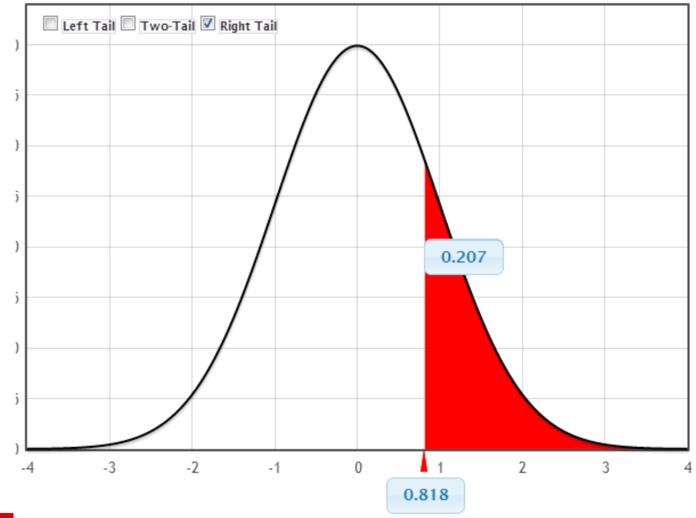

z-statistic

If $z = -3$, using $\alpha = 0.05$ we would

(a) Reject the null
(b) Not reject the null
(c) Impossible to tell
(d) I have no idea

About 95% of z-statistics are within -2 and +2, so anything beyond those values will be in the most extreme 5%, or equivalently will give a p-value less than 0.05.

Summary: Confidence Intervals

From N(0,1)

$$statistic \pm z^* \cdot SE$$

From original data

From bootstrap distribution

Summary: p-values

From original data

From H_0

$$\frac{\boldsymbol{statistic} - \boldsymbol{null}}{\boldsymbol{SE}}$$

From randomization distribution

Compare to N(0,1) for p-value

Standard Error

- Wouldn't it be nice if we could compute the standard error *without* doing thousands of simulations?

- We can!!!

- Or rather, we'll be able to next class!