Lecture 9

Introduction to Inference

Recall:**Statistical inference** is making a decision or a conclusion based on the data.

Example: Researchers want to know if a new drug is more effective than a placebo. Twenty patients receive the new drug, and 20 receive a placebo.

| Twelve (60%) of those taking the drug show improvement |

| only eight (40%) of the placebo patients |

VERSUS

We would conclude that the new drug is better. However, probability calculations yell us that a difference this large or larger between the results in the two groups would occur about one time in five simply because of chance variation. Since this probability is not very small, it is better to conclude that the observed difference is due to chance rather than a real difference between two treatments.

Example: Suppose you want to estimate the mean SAT score for the more than 420,000 high school seniors in California. You know better than to trust data from the students who choose to take the SAT. Only 45% of California students take the SAT. These self-selected students are planning to attend college and are not representative of all California seniors. At considerable effort and expense, you give the test to an SRS of 500 California high school seniors.

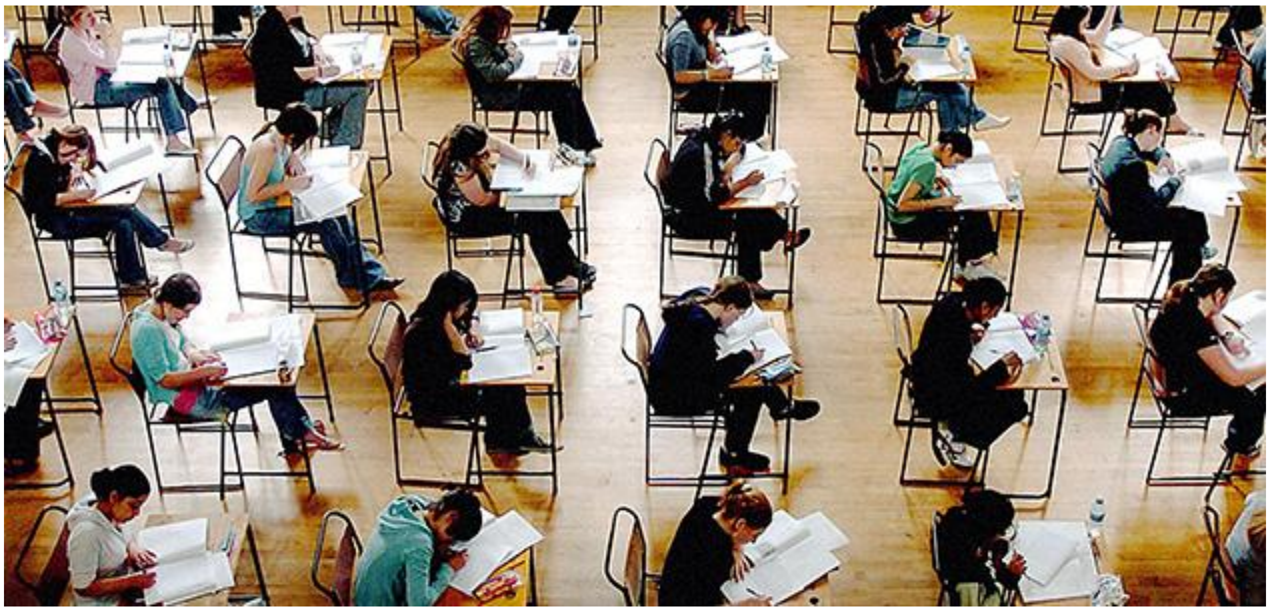

The mean score for your sample is $\bar{x} = 461$. What can you say about the mean score μ in the population of all 420,000 seniors?

Statistical Confidence

Suppose the entire population of SAT scores had mean μ and standard deviation $\sigma = 100$.

$$\bar{X} \underset{CLT}{\sim} N\left(\mu, \frac{\sigma}{\sqrt{n}}\right) = N\left(\mu, \frac{100}{\sqrt{500}}\right)$$

$$= N(\mu, 4.5)$$

By 68-95-99.7% Rule,

$$P(|\bar{X} - \mu| \leq 2\sigma_{\bar{X}}) = 0.95$$

$$P(|\bar{X} - \mu| \leq 9) = 0.95$$

$$|\bar{X} - \mu| \leq 9$$

$$-9 \leq \bar{X} - \mu \leq 9$$

$$P(-9 \le \bar{x} - \mu \le 9) = 0.95$$

$$P(\bar{x} - 9 \le \mu \le \bar{x} + 9) = 0.95$$

$$\bar{x} = 461$$

$$P(461 - 9 \le \mu \le 461 + 9) = 0.95$$

$$P(452 \le \mu \le 470) = 0.95$$

$$(452, 470)$$

We are 95% confident that true parameter μ is in this interval, which is equivalent to saying that we arrived at these numbers by a method that gives correct results 95% of the time.

Confidence Intervals (CI)

Figure below illustrates the behaviour of 95% CIs in repeated sampling. The center of each interval is at $\bar{x}$ and therefore varies from sample to sample. The sampling distribution of $\bar{x}$ appears on the top of the figure to show the long-term pattern of this variation.

95% CI

$\bar{x} \pm 2\sigma_{\bar{x}}$

estimate

$\pm$

margin of error

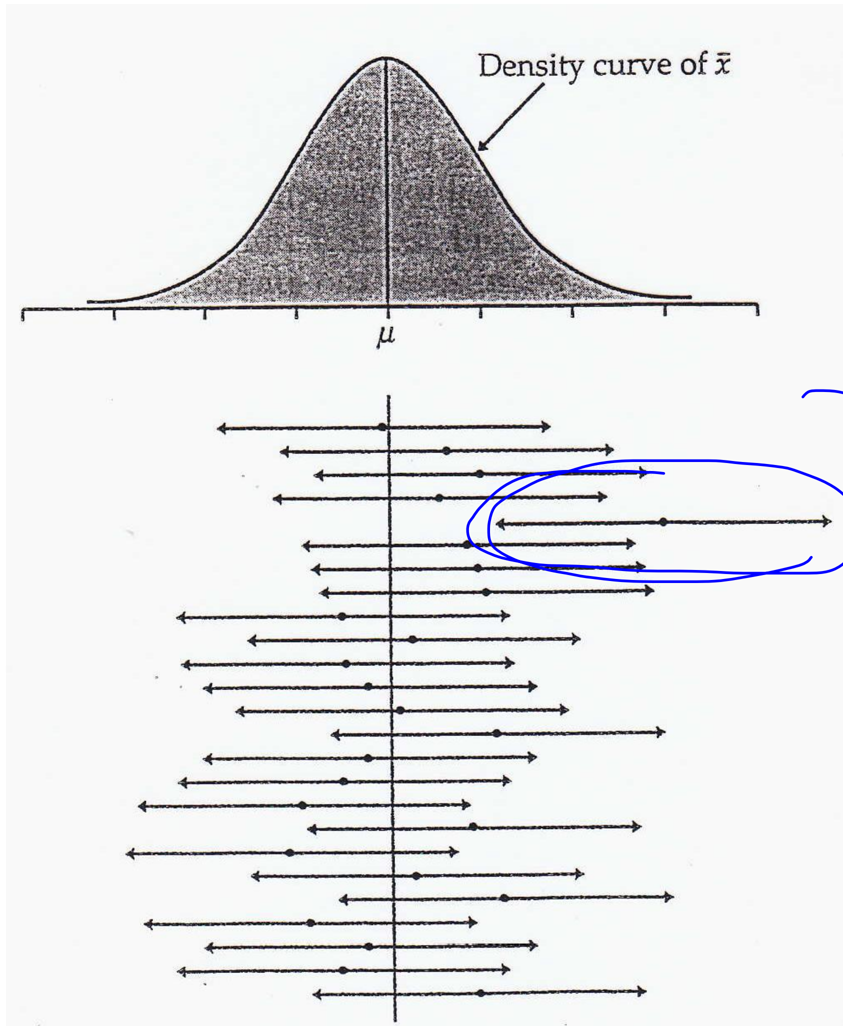

$\frac{24}{25} = 0.96$

25 CIs

CI for a Population Mean

Let's denote a confidence level by C, e.g. 95% confidence means $C = 0.95$.

Let $X_1, \ldots, X_n$ be an SRS from a population with mean μ and standard deviation σ.

Denote $1 - C = \alpha$.

$$\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i \sim N\left(\mu, \frac{\sigma}{\sqrt{n}}\right)$$

$$\frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \sim N(0, 1)$$

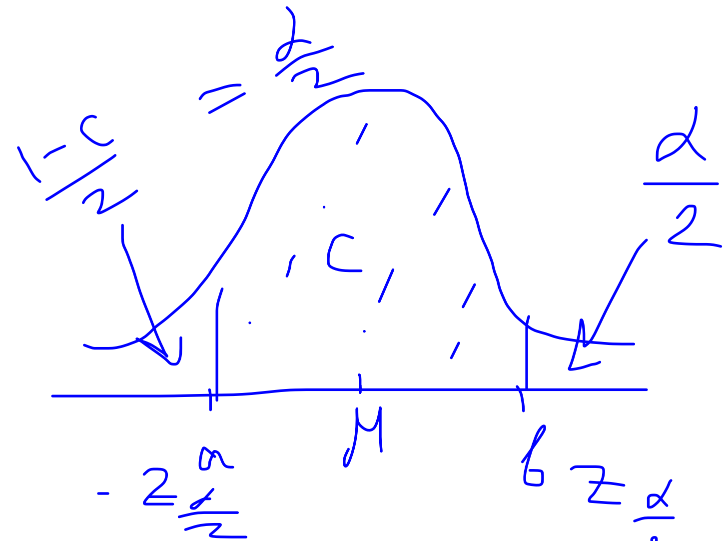

$$P(a \le \bar{X} \le b) = C$$

z-score for $b = \frac{b - \mu}{\sigma/\sqrt{n}} = z_{\frac{\alpha}{2}}$

z-score for $a = \frac{a - \mu}{\sigma/\sqrt{n}} = -z_{\frac{\alpha}{2}}$

$$C = P\left(-z_{\frac{\alpha}{2}} \le \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \le z_{\frac{\alpha}{2}}\right) = P\left(-z_{\frac{\alpha}{2}} \cdot \frac{\sigma}{\sqrt{n}} \le \bar{X} - \mu \le z_{\frac{\alpha}{2}} \cdot \frac{\sigma}{\sqrt{n}}\right)$$

$$C = P\left(\bar{X} - z_{\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}} \le \mu \le \bar{X} + z_{\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}}\right)$$

Definition: Choose an SRS of size n from a population having unknown mean μ and known standard deviation σ. The level C or $100(1-\alpha)\%$ **confidence interval** for μ is

$$\boxed{\bar{x} \pm z_{\alpha/2} \cdot \frac{\sigma}{\sqrt{n}}}$$

The quantity $ME = z_{\alpha/2} \cdot \frac{\sigma}{\sqrt{n}}$ is called the **margin of error,** and $z_{\alpha/2}$ is the value on the standard Normal curve with area $1-\alpha$ between the critical points $-z_{\alpha/2}$ and $z_{\alpha/2}$.

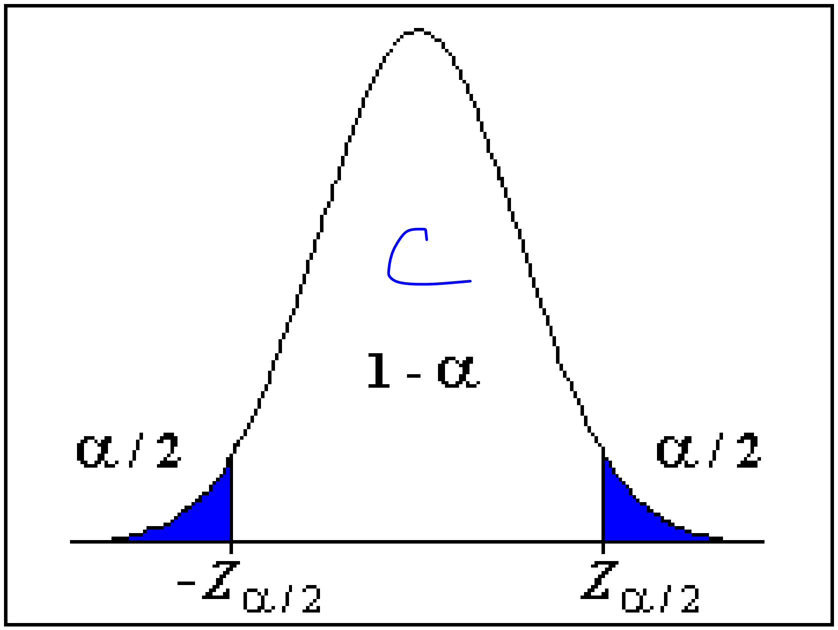

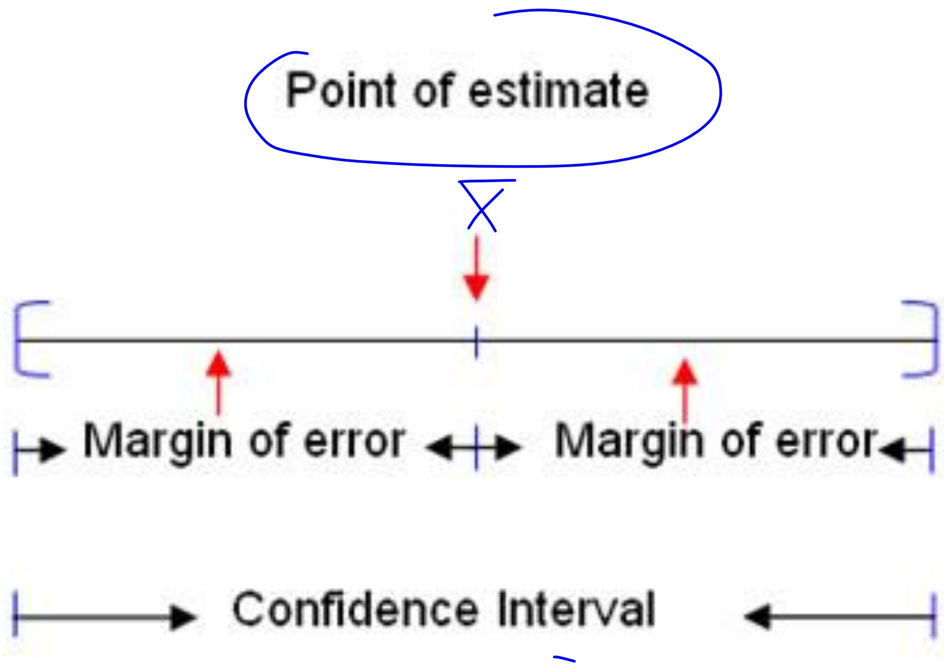

This interval is exact when the population distribution is Normal and is approximately correct when n is large in other cases.

Example: The National Student Loan Survey collected data about the amount of money that borrowers owe.

n

The survey selected a random sample of 1280 borrowers who began repayment of their loans between four to six months prior to the study.

The mean debt for the selected borrowers was $\bar{x} = \$18{,}900$. Assume $\sigma = \$49{,}000$.

Find a 95% CI for the mean debt for all borrowers.

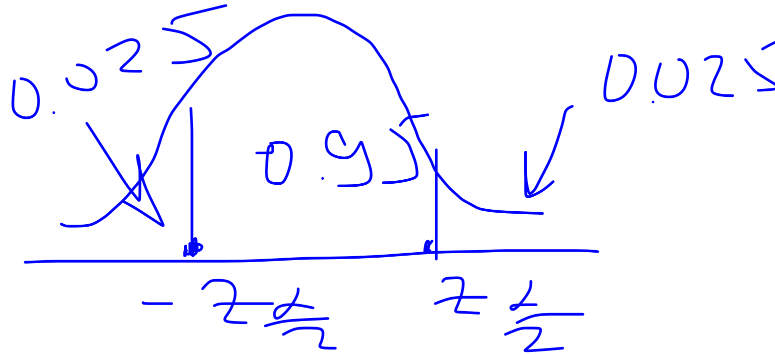

$$\bar{x} \pm z_{\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}}$$

$\bar{x} = 18{,}900$, $n = 1280$, $\sigma = 49{,}000$

$\alpha = 1 - C = 1 - 0.95 = 0.05$, $\frac{\alpha}{2} = 0.025$

$z_{\frac{\alpha}{2}} = 1.96$

95% CI: $18{,}900 \pm 1.96 \cdot \frac{49{,}000}{\sqrt{1280}}$

$= 18{,}900 \pm \underbrace{2700}_{} \rightarrow ME$

$= (16{,}200, 21{,}600)$

<u>Conclusion</u>: We are 95% confident that the mean debt of all borrowers is between \$16,200 and \$21,600.

How does a sample size affect the CI? 95% CI

Let n = 320.

$$ME = z_{\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}} = 1.96 \cdot \frac{49000}{\sqrt{320}}$$

$$= 5400$$

$$\bar{x} \pm ME = 18,900 \pm 5400$$

$$= (13,500, 24,300)$$

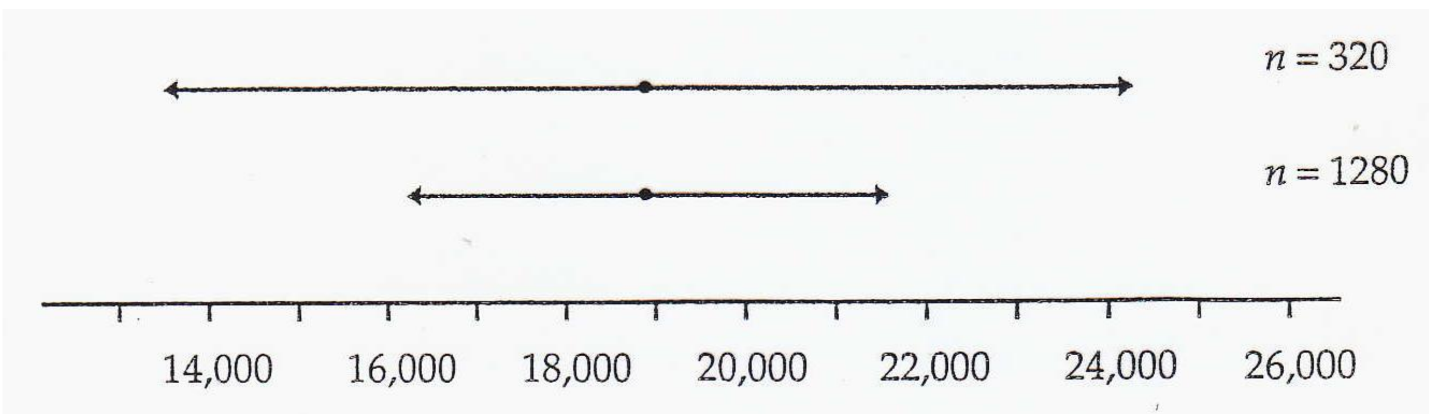

How does a confidence level affect the CI?

$C = 0.99$

99% CI , $z_{\frac{\alpha}{2}} = 2.575$

$$ME = z_{\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}} = 2.575 \cdot \frac{49000}{\sqrt{1280}} = 3,500$$

$$\bar{x} \pm ME = (15,400, \ 22,400)$$

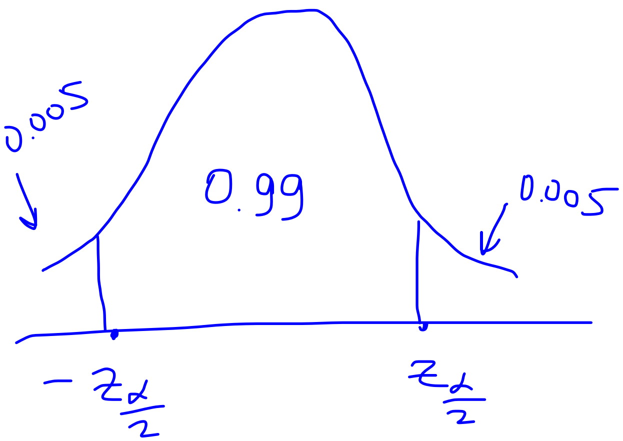

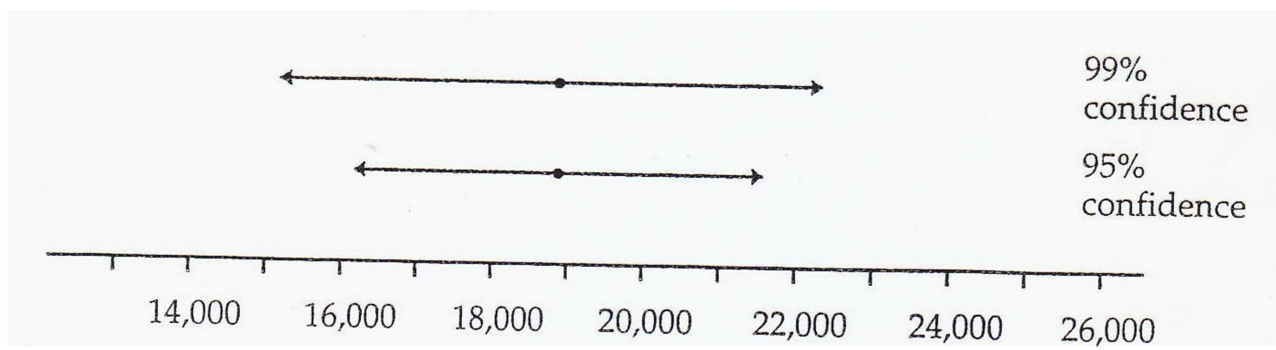

To reduce ME:

- Increase n
- Reduce σ
- Choose smaller C

$$ME = z_{\frac{\alpha}{2}} \frac{\sigma}{\sqrt{n}}$$

What if we don't know σ?

***t* distribution**

Let $X_1, X_2, \dots, X_n$ be an SRS from $N(\mu, \sigma)$, μ, σ are both unknown.

Then $\bar{X} \sim N(\mu, \frac{\sigma}{\sqrt{n}})$.

We use s (sample standard deviation) to estimate σ.

Definition: When the standard deviation of a statistic is estimated from the data, the result is called the **standard error** of the statistic. The standard error of the sample mean is

$$\boxed{SE_{\bar{x}} = \frac{s}{\sqrt{n}}}$$

$$Z = \frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \sim N(0,1)$$

$$T = \frac{\bar{X} - \mu}{s/\sqrt{n}} \sim \underline{t_{n-1}}$$

$$s = \frac{1}{n-1} \sum (x_i - \bar{x})^2$$

Definition: Suppose that an SRS of size n is drawn from $N(\mu, \sigma)$ population. Then the **one-sample *t* statistic**

$$t = \frac{\bar{x} - \mu}{s/\sqrt{n}}$$

has the ***t* distribution (Student's *t*-distribution)** with n - 1 **degrees of freedom**.

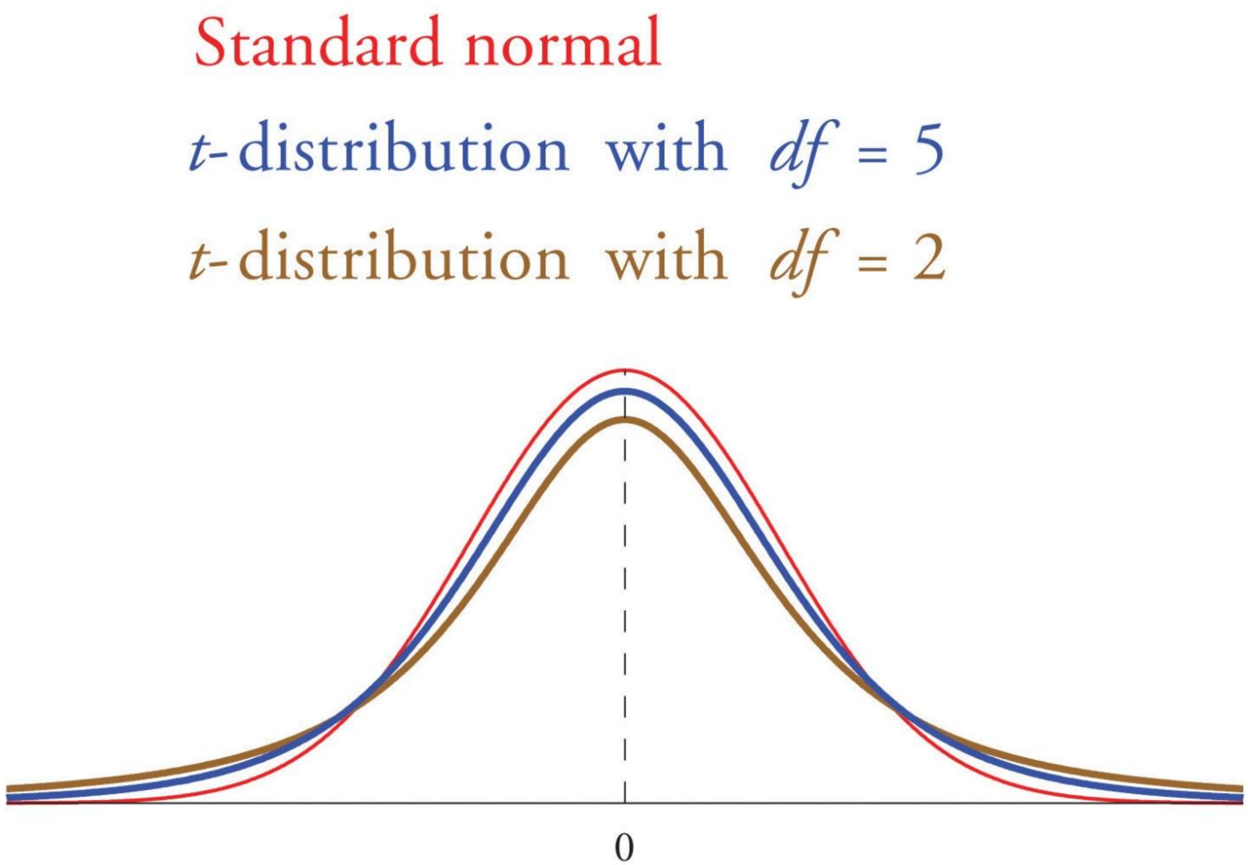

One-Sample *t* CI: Suppose that an SRS of size n is drawn from a population having unknown mean μ. A 100(1-α)% confidence interval for μ is

$$\boxed{\bar{x} \pm t_{\alpha/2} \frac{s}{\sqrt{n}}}$$

Where $t_{\alpha/2}$ is the value for the t_{n-1} density curve with area $1-\alpha$ between $-t_{\alpha/2}$ and $t_{\alpha/2}$.

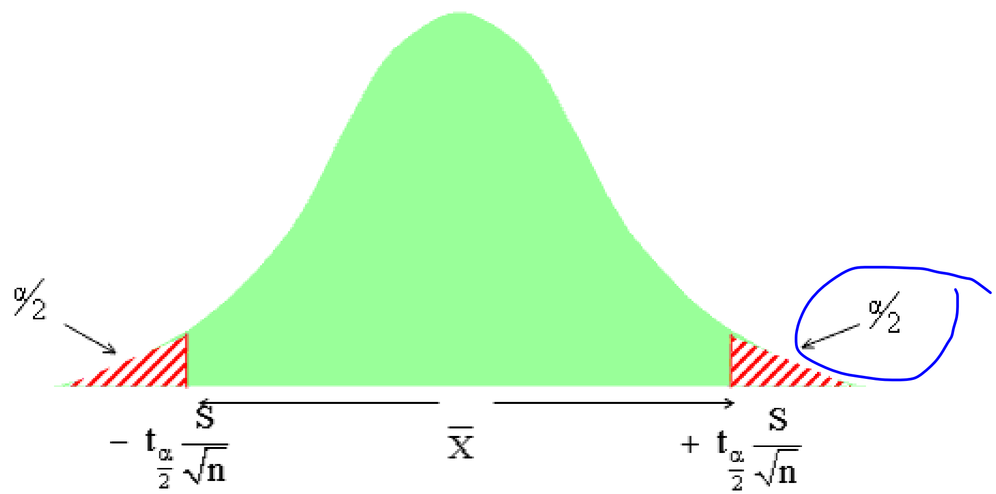

The quantity

$$t_{\alpha/2} \frac{s}{\sqrt{n}}$$

is the **margin of error**. This interval is exact when the population distribution is Normal and is approximately correct for large n in other cases.

Example: Founded in 1998, Telephia provides a wide variety of information on cellular phone use.

TELEPHIA®

In 2006, Telephia reported that, on average, United Kingdom (U.K.) subscribers with third-generation technology (3G) phones spent an average of 8.3 hours per month listening to full-track music on their cell phones.

Suppose we want to determine a 95% CI for the U.S. average and draw the following random sample of size 8 from the U.S. population of 3G subscribers:

5 6 0 4 11 9 2 3

$n = 8$

The sample mean is $\boxed{\bar{x} = 5}$ and the standard deviation $s = 3.63$ with degrees of freedom $n - 1 = 7$.

5 6 0 4 11 9 2 3

$\bar{x} = 5, \quad s = 3.63, \quad n - 1 = 7$

95% CI : $\bar{x} \pm t_{\frac{\alpha}{2}} \frac{s}{\sqrt{n}}$

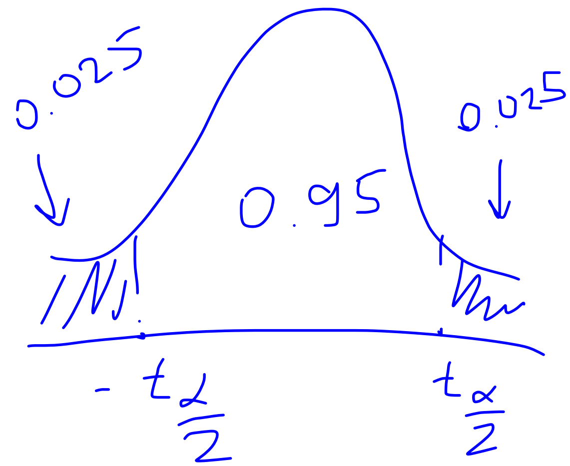

$t_{\frac{\alpha}{2}} = 2.365$

95% CI $= 5 \pm 2.365 \cdot \frac{3.63}{\sqrt{8}}$

$= (2, 8)$

Conclusion: We are 95% confident that U.S. population spends on average between 2 and 8 hours listening music on cell phones.

Inference for a Single Proportion

Choose an SRS of size n from a large population with unknown proportion p of successes.

The **sample proportion** is

$$\hat{p} = \frac{X}{n}$$

where X is the number of successes

$X \sim Bin(n, p)$

$\hat{p} \sim$ appr. $N\left(p, \sqrt{\frac{p(1-p)}{n}}\right)$

Use $\hat{p}$ to estimate p

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

A $100(1-\alpha)\%$ CI for p is

$$\hat{p} \pm z_{\alpha/2}\sqrt{\hat{p}(1-\hat{p})/n}$$

where $z_{\alpha/2}$ is the value for the standard Normal density curve with area $(1-\alpha)$ between $-z_{\alpha/2}$ and $z_{\alpha/2}$.

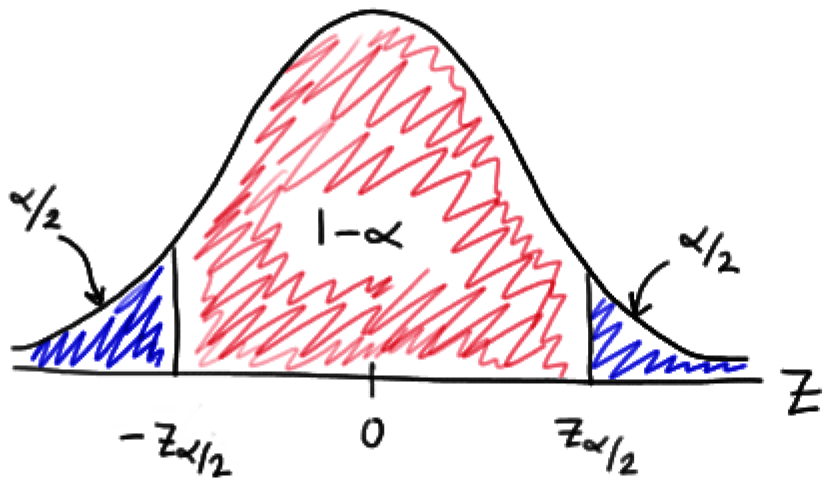

The value $SE_{\hat{p}} = \sqrt{\hat{p}(1-\hat{p})/n}$ is called the **standard error** of $\hat{p}$ and $ME = z_{\alpha/2}\sqrt{\hat{p}(1-\hat{p})/n}$ is the **margin of error** for the given confidence level.

<u>Example</u>:

Alcohol abuse has been described by college presidents as the number one problem on campus, and it is a major cause of death in young adults. How common is it?

A survey of 13,819 students in four-year colleges collected information on drinking behavior and alcohol-related problems.

The researchers defined "binge drinking" as having five or more drinks in a row for men and four or more drinks in a row for women. "Frequent binge drinking" was defined as binge drinking three or more times in the past two weeks. According to this definition, 3140 students were classified as frequent binge drinkers.

LIFESTYLE

"Alcohol may be man's worst enemy, but the Bible says love your enemy."

-Frank Sinatra

WEDNESDAY
OCTOBER 1, 2008

Lifestyle Editor: Abigail Naumann
Assistant Editor: Victoria Vlisides

PAGE 8

BINGE DRINKING

ALCOHOL RELATED INCIDENTS ARE ON THE RISE.
REPORTED PROTECTIVE CUSTODY INCIDENTS FOR SEPTEMBER, HAVE ALREADY MATCHED THE ENTIRETY OF LAST YEAR'S INCIDENTS .

By Victoria Vlisides

Jon Good photo/Goodj22@uww.edu

Travis Borchert, Alex Baxter, Bryan Lehman, Ryan Schlipp, Blaine Patino, Adam Pohlkamp, and Ben Wilson celebrate, what is for many students, an old past-time, Thirsty Thursdays.

The college drinking culture is undeniable. Dates like St. Patrick's Day, Homecoming and just about any 21st birthday exemplify this culture because all three usually contain mass amounts of alcohol consumption, also known as binge drinking.

This semester, protective custody incidents have already reached as many as the last year's total, with four occurrences from Sept. 5 to Sept. 26, making it imperative students remain aware of their drinking habits and when they've gone from healthy to out-of-control.

"We take them into protective custody because they are too intoxicated to protect themselves," Operations Program Associate Deanne Thomsen said.

Sophomore Jonta Brown, thought binge drinking was the norm for college life in Wisconsin.

"Everybody knows the UW-systems are well-known for drinking," Brown said.

According to the Web site www.intheknowzone.com, 60 percent of college women who have acquired sexually transmitted diseases, including AIDS, were under the influence of alcohol at the time they had intercourse. Over 30,000 students each year need emergency health care for alcohol overdose.

College students spend $5.5 billion on alcohol per year. This is more than is spent on books, soda, coffee, juice and milk combined, according to the Web site.

University Police Chief Matt Kiederlen said the drinking habits in Wisconsin are much more drastic than where he previously was an officer in Illinois.

"There's a social ideal around alcohol that makes it a focus of an event," Kiederlen said.

Kiederlen and senior Megan Repp both thought of binge drinking as drinking in access for the purpose to get drunk.

"When people drink for a loss of control, that's where heavy regulation comes into play," Kiederlen said.

Repp, on a typical night of drinking, has anywhere from "three to I can't remember drinks."

Sophomore Brittany Hereford thought a student would have about three to four drinks on an average Thursday night.

Binge drinking is formally defined as, "a pattern of drinking that brings a person's blood alcohol concentration to 0.08 grams percent or above," according to the National Institute of Alcohol Abuse and Alcoholism Web site, www.cdc.gov. "This typically happens when men consume 5 or more drinks, and when women consume 4 or more drinks, in about 2 hours."

Many students have recognized alcohol as an issue in their lives.

SEE DRINKING PAGE 9

"I'LL KILL SIX TO EIGHT BEERS. I DON'T CONSIDER IT BINGE DRINKING BECAUSE I DO IT IN A SOCIAL SETTING, AND I DON'T CHUG."
- Adam Cumbee

"I STARTED DRINKING AT 4 P.M. LAST SATURDAY WHILE PLAYING 'LANDMINES.'"
-Marcus Shellberg

"ONE NIGHT I GOT REALLY HAMMERED. WE SPENT ABOUT $400 AT THE BAR, AND WE WOKE UP WITH SOME RANDOM GUY IN OUR APARTMENT."
-Drew Wolfgram

$x = 3140, \ n = 13819$

$$\hat{p} = \frac{3140}{13819} = 0.227$$

95% CI : $z_{\frac{\alpha}{2}} = 1.96$

$$\frac{\alpha}{2} = \frac{0.05}{2} = 0.025$$

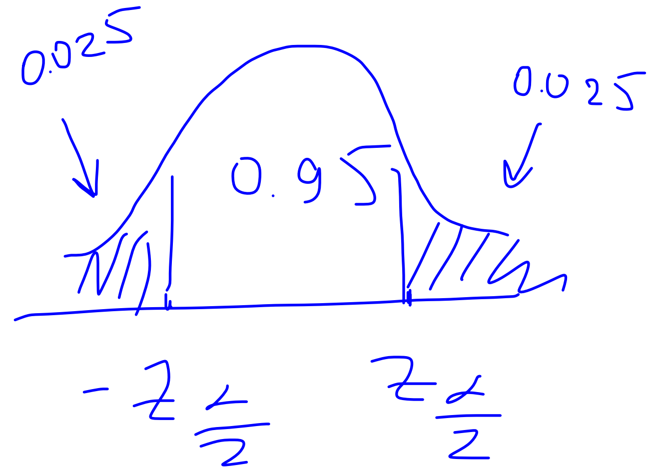

$$\boxed{\hat{p} \pm z_{\frac{\alpha}{2}} \cdot \sqrt{\frac{\hat{p}(1-\hat{p})}{n}}}$$

$$= 0.227 \pm 1.96 \cdot \sqrt{\frac{0.227(1-0.227)}{13,819}}$$

$$= (0.220, 0.234)$$

<u>Conclusion</u>: We estimate with 95% confidence that between 22% and 23.4% of college students are binge drinkers.

Hypothesis Testing

A hypothesis is a conjecture about the distribution of some randomvariables. For example, a claim about the value of a parameter of the statistical model.

There are two types of hypotheses:

- The null hypothesis, H_0, is the current belief.
- The alternative hypothesis, H_a, is your belief; it is what you want to show.

Examples: Each of the following situations requires a significance test about a population mean. State the appropriate null hypothesis H_0 and alternative hypothesis H_a in each case.

(a) The mean area of the several thousand apartments in a new development is advertised to be 1250 square feet. A tenant group thinks that the apartments are smaller than advertised. They hire an engineer to measure a sample of apartments to test their suspicion.

$H_0: \mu = 1250$

$H_a: \mu < 1250$

→ one-sided alternative

(b) Larry's car consume on average 32 miles per gallon on the highway. He now switches to a new motor oil that is advertised as increasing gas mileage. After driving 3000 highway miles with the new oil, he wants to determine if his gas mileage actually has increased.

$H_0 : \mu = 32$

$H_a : \mu > 32$

↳ one-sided alternative

(c) The diameter of a spindle in a small motor is supposed to be 5 millimeters. If the spindle is either too small or too large, the motor will not perform properly. The manufacturer measures the diameter in a sample of motors to determine whether the mean diameter has moved away from the target.

$H_0 : \mu = 5$

$H_a : \mu \neq 5$

↳ two-sided alternative

Guidelines for Hypothesis testing

Hypothesis testing is a proof by contradiction. The testing process has four steps:

Step 1: Assume H_0 is true.

Step 2: Use statistical theory to make a statistic (function of thedata) that includes H_0. This statistic is called the **test statistic**.

Step 3: Find the probability that the test statistic would take a value as extreme or more extreme than that actually observed. Think of this as: probability of getting our sample assuming $H_0$is true. This is what we'll call a **P-value**.

Step 4: If the probability we calculated in step 3 is high it meansthat the sample is likely under H_0 and so we have no evidenceagainst H_0. If the probability is low, there are two possibilities:

- we observed a very unusual event, or
- our assumption is wrong

$\hookrightarrow$ reject H_0

Test Statistic

• The test is based on a statistic that estimates the parameter that appears in the hypotheses. Usually this is the same estimate we would use in a confidence interval for the parameter. When H_0 is true, we expect the estimate to take a value near the parameter value specified in H_0.

• Values of the estimate far from the parameter value specified by H_0 give evidence against H_0. The alternative hypothesis determines which directions count against H_0.

• A **test statistic** measures compatibility between the null hypothesis and the data.

• To assess how far the estimate is from the parameter, standardize the estimate. In many common situations the test statistics has the form

$$\boxed{\frac{estimate - hypothesized\ value}{standard\ deviation\ of\ the\ estimate}}$$

Example: An air freight company wishes to test whether or not the mean weight of parcels shipped on a particular root exceeds 10 pounds. A random sample of 49 shipping orders was examined and found to have average weight of 11 pounds. Assume that the standard deviation of the weights is 2.8 pounds.

$n = 49$, $\bar{X} = 11$, $\sigma = 2.8$

$$\sigma_{\bar{X}} \underset{CLT}{=} \frac{\sigma}{\sqrt{n}} = \frac{2.8}{\sqrt{49}}$$

$H_0: \mu = 10$

$H_a: \mu > 10$

Test statistic: $\dfrac{\bar{X} - 10}{2.8/\sqrt{49}} \overset{H_0}{\sim} N(0,1)$

$$\frac{11 - 10}{2.8/\sqrt{49}} = 2.5$$

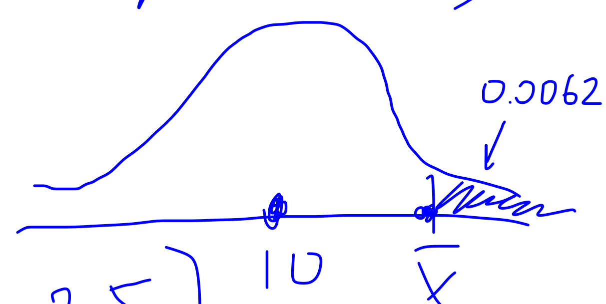

$$P(Z \geq 2.5) = 1 - P(Z \leq 2.5)$$
$$= 1 - 0.9938 = 0.0062 \Rightarrow \text{reject } H_0$$

small

P-values

Definition: The probability, assuming H_0 is true, that the test statistic would take a value as extreme or more extreme than that actually observed is called the **P-value** of the test. The smaller the P-value, the stronger the evidence against H_0 provided by the data.

Guideline for how small is "small":

- P-value > 0.1 provides no evidence against H_0.
- 0.05 < P-value < 0.1 provides weak evidence against H_0.
- 0.01 < P-value < 0.05 provides moderated evidence against H_0.
- P-value < 0.01 provides strong evidence against H_0.

We can compare the P-value we calculate with a fixed value that we regard as decisive. The decisive value of P is called the **significance level** (denoted byα). Most common values for α are 0.1, 0.05, 0.01.

If the P-value is as small or smaller than α, we say that the data are **statistically significant at level $\boldsymbol{\alpha}$**.

For example, P-value = 0.03 is significant at the level $\alpha = 0.05$, but not significant at the level $\alpha = 0.01$.

Example: 85% of the general public is right-handed. A survey of 300 chief executive officers of large corporations found that 95% were right-handed. Is this difference in percentages statistically significant? Use $\alpha = 0.01$. Find the P-value for the test.

$H_0: p = 0.85$ $\qquad \hat{p} = 0.95 \qquad n = 300$

$H_a: p > 0.85$ $\qquad \sigma_{\hat{p}} = \sqrt{\dfrac{p(1-p)}{n}}$

p = proportion of right-handed executives

$$z = \frac{\hat{p} - 0.85}{\sqrt{\frac{0.85(1-0.85)}{300}}} = \frac{0.95 - 0.85}{\sqrt{\frac{0.85(1-0.85)}{300}}} = 4.43$$

If H_0 is true, $z \sim N(0,1)$

$$\text{P-value} = P(z \geq 4.43) = 0.0001 < \alpha = 0.01$$

$\Rightarrow$ reject H_0

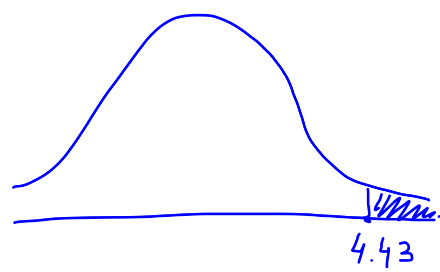

Tests for a Population Mean (σ is unknown)

<u>One-Sample t Test:</u> Suppose that an SRS of size n is drawn from a population having unknown mean μ. To test the hypothesis $H_0: \mu = \mu_0$ based on an SRS of size n, compute the one-sample t statistic

$$t = \boxed{\frac{\bar{x} - \mu_0}{s/\sqrt{n}}}$$

In terms of a random variable T having t_{n-1} distribution, the P-value for a test of H_0 against

$H_a: \mu > \mu_0$ is $P(T \geq t)$ 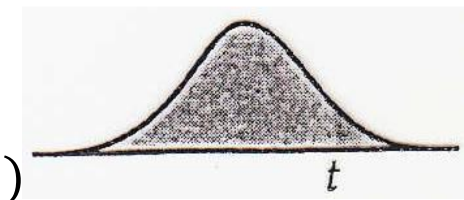

$H_a: \mu < \mu_0$ is $P(T \leq t)$ 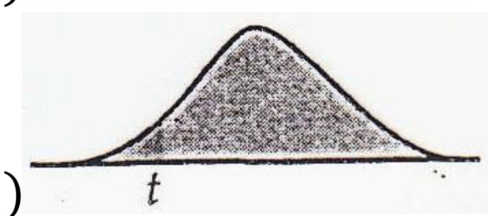

$H_a: \mu \neq \mu_0$ is $2P(T \geq |t|)$ 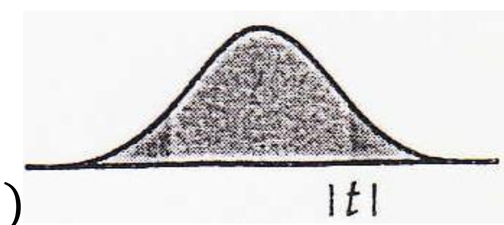

These P-values are exact if the population distribution is Normal and are approximately correct for large n in other cases.

Example: Suppose that, for the U.S. data in example before we want to test whether the U.S. average is different from the reported U.K. average.

$$H_0: \mu = 8.3 = \mu_0$$
$$H_a: \mu \neq 8.3$$

$$\boxed{n = 8, \qquad \bar{x} = 5, \qquad s = 3.63}$$

$$t = \frac{\bar{X} - \mu_0}{s/\sqrt{n}} = \frac{5 - 8.3}{3.63/\sqrt{8}} = -2.57$$

$$\text{P-value} = 2 \cdot P(T \geq |-2.57|)$$
$$= 2 \underbrace{P(T \geq 2.57)}$$

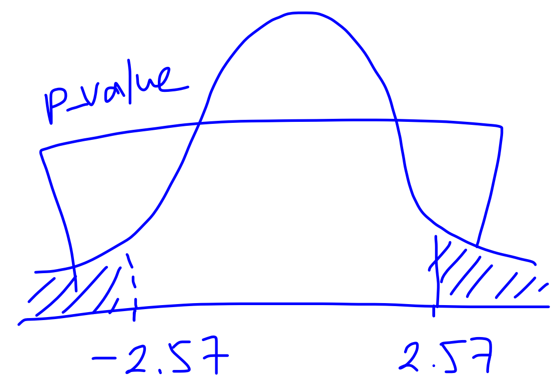

p	0.02	0.01
t_7	2.517	2.998

$$2 \cdot 0.01 \leq \text{P-value} \leq 2 \cdot 0.02$$

$$0.02 \leq \text{P-value} \leq 0.04 < 0.05$$

$\boxed{\text{Software P-value} = 0.037}$

$\Rightarrow$ reject H_0 at $\alpha = 0.05$

What if we want to test whether the U.S. average is smaller than the UK average?

$$H_0: \mu = 8.3$$
$$H_a: \mu < 8.3$$

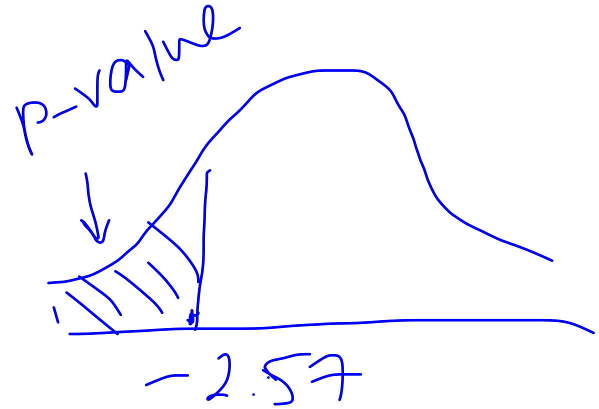

$t = -2.57$

$\text{p-value} = P(T \leq -2.57)$

$.01 \leq \text{p-value} \leq 0.02 < 0.05$

Software P-value = 0.0185

$\Rightarrow$ reject H_0

At $\alpha = 0.05$ level, we conclude that the U.S. average is smaller than the U.K. average.

Robustness of the *t* procedure

The results of one-sample t procedures are exactly correct only when the population is Normal. In practice, the usefulness of the t procedures depends on how strongly they are affected by non-Normality.

Definition: A statistical inference procedure is called **robust** if the required probability calculations are insensitive to violations of the assumptions made.

Here are practical guidelines for inference on a single mean:

- Sample size is less than 15: Use t procedures if data are close to Normal. If data are clearly non-Normal or if outliers are present, do not use t.
- Sample size at least 15: The t procedure can be used except in the presence of outliers or strong skewness.
- Large samples: The t procedures can be used even for clearly skewed distributions when the sample is large, roughly $n \geq 40$.

Testing Hypotheses on a Proportion

Draw an SRS of size n from a large population with unknown proportion p of successes. To test the hypothesis $H_0: p = p_0$, compute the z statistic

$$z = \frac{\hat{p} - p_0}{\sqrt{\frac{p_0(1-p_0)}{n}}}$$

In terms of a standard Normal random variable Z, the approximate P-value for a test of H_0 against

$H_a: p > p_0$ is $P(Z \geq z)$

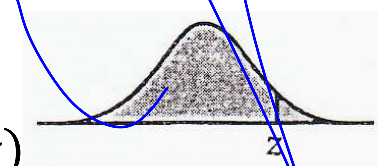

$H_a: p < p_0$ is $P(Z \leq z)$

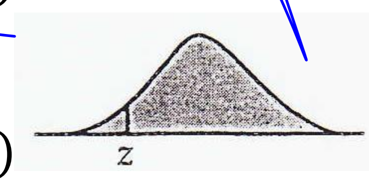

$H_a: p \neq p_0$ is $2P(Z \geq |z|)$

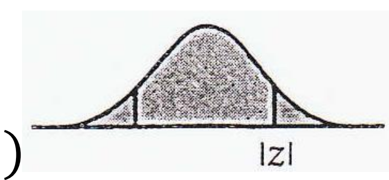

(We assume that $np_0 \geq 10, n(1-p_0) \geq 10$)

Example: According to the National Institute for Occupational Safety and Health, job stress poses a major threat to the health of workers.

A national survey of restaurant employees found that 75% said that work stress had a negative impact on their personal lives.

A sample of 100 employees of a restaurant chain finds that 68 answer "Yes" when asked, "Does work stress have a negative impact on your personal life?"

Is this good reason to think that the proportion of all employees of this chain who would say "Yes", differs from the national proportion p_0=0.75?

SKIP

<u>Conclusion</u>: The restaurant data are compatible with the survey results.