

Stat 13, Intro. to Statistical Methods for the Life and Health Sciences.

0. Submit hw4 to statgrader.
1. Common problems with regression.
 - a. Inferring causation.
 - b. Extrapolation.
 - c. Curvature.
2. Testing significance of correlation or slope.
3. ANOVA and F-test, ch9.
4. Review and examples.

The final is Fri in class and will be on ch 1-7, 10, and at most 1 question on ch9. Bring a PENCIL and CALCULATOR and any books or notes you want. No computers. <http://www.stat.ucla.edu/~frederic/13/sum22> .


As mentioned last 2 classes, if you cannot take it because of Covid or other health reasons, then email me to arrange a time to take an oral, in-person, 10-15 minute final exam in my office.

Alternatively, you get an incomplete in the course and take the Fall stat 13 final, or I may also be able to arrange a makeup exam date in Sept. But don't come if you have covid or just recently had covid.

1. Common problems with regression.

- a. Correlation is not causation.

ESPECIALLY WITH OBSERVATIONAL DATA!



Common problems with regression.



Common problems with regression.

Holmes and Willett (2004) reviewed all prospective studies on fat consumption and breast cancer with at least 200 cases of breast cancer. "Not one study reported a significant positive association with total fat intake.... Overall, no association was observed between intake of total, saturated, monounsaturated, or polyunsaturated fat and risk for breast cancer."

They also state "The dietary fat hypothesis is largely based on the observation that national per capita fat consumption is highly correlated with breast cancer mortality rates. However, per capita fat consumption is highly correlated with economic development. Also, low parity and late age at first birth, greater body fat, and lower levels of physical activity are more prevalent in Western countries, and would be expected to confound the association with dietary fat."

Common problems with regression.

- b. Extrapolation.



If the birthrate remains at **1.19** children per woman, South Korea could face natural extinction by **2750**.

Source:
<http://blogs.wsj.com/korearealtime/2014/08/26/south-korea-birthrate-hits-lowest-on-record/>

BROOKINGS

Common problems with regression.

- b. Extrapolation.
- Often researchers extrapolate from high doses to low.

D.M. Odom et al.




Figure 4. Relationship between diclofenac daily dose and the estimated risk ratio of a cardiovascular event. Numbers correspond to the observations in [Table III](#).


Common problems with regression.

- b. Extrapolation.

The relationship can be nonlinear though.

Researchers also often extrapolate from animals to humans.

Zaichkina et al. (2004) on hamsters



Common problems with regression.

- c. Curvature.

The best fitting line might fit poorly. Port et al. (2005).










FIGURE 4. Adjusted 2-year rates of death from all causes for men (upper panel) and women (lower panel) separately, by glucose level, predicted by three models, Framingham Heart Study, 1948–1978. Linear model (dashed curve); optimal spline models (solid curve). The horizontal dashed

Common problems with regression.

- c. Curvature.


The best fitting line might fit poorly. Wong et al. (2011).



Common problems with regression.

- d. Statistical significance.

Could the observed correlation just be due to chance alone?



5. Inference for the Regression Slope: Theory-Based Approach

Section 10.5

Do students who spend more time
in non-academic activities tend to
have lower GPAs?

Example 10.4

Do students who spend more time in non-academic activities tend to have lower GPAs?


- The subjects were 34 undergraduate students from the University of Minnesota.
- They were asked questions about how much time they spent in activities like work, watching TV, exercising, non-academic computer use, etc. as well as what their current GPA was.
- We are going to test to see if there is a significant **negative** association between the number of hours per week spent on nonacademic activities and GPA.

Hypotheses

- Null Hypothesis: There is no association between the number of hours students spend on nonacademic activities and student GPA in the population.
- Alternative Hypothesis: There is a negative association between the number of hours students spend on nonacademic activities and student GPA in the population.

Descriptive Statistics

- $\widehat{\text{GPA}} = 3.60 - 0.0059(\text{nonacademic hours})$.
- Is the slope significantly different from 0?



Shuffle to Develop Null Distribution

- We are going to shuffle just as we did with correlation to develop a null distribution.
- The only difference is that we will be calculating the slope each time and using that as our statistic.
- **a test of association based on slope is equivalent to a test of association based on a correlation coefficient.**

Beta vs Rho


- Testing the slope of the regression line is equivalent to testing the correlation (same p-value, but obviously different confidence intervals since the statistics are different)
- Hence these hypotheses are equivalent.
 - $H_0: \beta = 0$ $H_a: \beta < 0$ (Slope)
 - $H_0: \rho = 0$ $H_a: \rho < 0$ (Correlation)
- Sample slope (b) Population (β : beta)
- Sample correlation (r) Population (ρ : rho)
- When we do the theory based test, we will be using the t -statistic which can be calculated from either the slope or correlation.

Introduction

- Our null distributions are again bell-shaped and centered at 0 (for either correlation or slope as our statistic).


Example 10.2: Exercise and mood intensity

● Correlation ○ Slope ○ *t*-statistic



Example 10.4: GPA and nonacademic hours

○ Correlation ● Slope ○ *t*-statistic



The book on p549 finds a p value of 3.3% by simulation.

Validity Conditions


- Under the usual conditions: relationship is linear, observations are iid, both variables are normally distributed, and data are homoscedastic, theory-based inference for correlation or slope of the regression line uses the t -distribution.
- We could use simulations or the theory-based methods for the slope of the regression line.
- We would get exactly the same p-value if we used correlation as our statistic.

Predicting Heart Rate from Body Temperature

Example 10.5A

Heart Rate and Body Temp

- Earlier we looked at the relationship between heart rate and body temperature with 130 healthy adults
- Predicted Heart Rate = $-166.3 + 2.44(\text{Temp})$
- $r = 0.257$




Heart Rate and Body Temp

- We tested to see if we had convincing evidence that there is a positive association between heart rate and body temperature in the population using a simulation-based approach. (We will make it 2-sided this time.)
- **Null Hypothesis:** There is no association between heart rate and body temperature in the population. $\beta = 0$
- **Alternative Hypothesis:** There is an association between heart rate and body temperature in the population. $\beta \neq 0$

Heart Rate and Body Temp

We get a very small p-value (0.0036). Anything as extreme as our observed slope of 2.44 happening by chance is very rare.





Heart Rate and Body Temp

- We can also approximate a 95% confidence interval
observed statistic $\pm$ multiplier $\times$ SE
 $2.44 \pm 1.96 \times 0.842 = 0.790$ to 4.09
- When both variables are normally distributed (scatterplot is elliptical), use the t-multiplier instead of 1.96, but when n is large it makes very little difference.
- This means we are 95% confident that, in the population of healthy adults, each 1° increase in body temp is associated with an increase in heart rate of between 0.790 to 4.09 beats per minute.

Heart Rate and Body Temp

- The theory-based approach should work well since the distribution of the slopes has a nice bell shape
- Also check the scatterplot



Heart Rate and Body Temp

- We will use the t-statistic to get our theory-based p-value.
- We will find a theory-based confidence interval for the slope.
- On p554, the book notes the formula $t = \frac{r}{\sqrt{\frac{1-r^2}{n-2}}}$.
- Here the t statistic is 2.97.
- The p-value is 0.36%. So the correlation is statistically significantly greater than zero.

Smoking and Drinking

Example 10.5B

Validity Conditions

Remember our validity conditions for theory-based inference for slope of the regression equation.



1. The scatterplot should follow a linear trend.
2. The observations should be iid.
3. For the t-test, both variables should be normal.

In particular, there should be approximately the same number of points above and below the regression line (symmetry).

4. The variability of vertical slices of the points should be similar. This is called homoskedasticity.

Validity Conditions

- Let's look at some scatterplots that do not meet the requirements.




Smoking and Drinking

The relationship between number of drinks and cigarettes per week for a random sample of students at Hope College.

The dot at (0,0)
represents 524
students


Are the conditions met?
Hard to say. The book
says no.



Smoking and Drinking

- When the conditions are not met, applying simulation-based inference is preferable to theory-based t-tests and CIs.

○ Correlation ○ Slope ⊙ *t*-statistic



Validity Conditions

- What do you do when validity conditions aren't met for theory-based inference?
 - Use the simulated-based approach.
- Another strategy is to “transform” the data on a different scale so conditions are met.
 - The logarithmic scale is common.
- One can also fit a different curve, not necessarily a line.

6. ANOVA and F-test.

Section 9.2

ANOVA

- ANOVA stands for ANalysis Of VAriance.
- Useful when comparing more than 2 means.
- If I have 2 means to compare, I just look at their difference to measure how far apart they are.
- Suppose I wanted to compare three means. I have the mean for group A, the mean for group B, and the mean for group C.

F test statistic

- The analysis of variance F test statistic is:

$$F = \frac{\text{variability between groups}}{\text{variability within groups}}$$

- This is similar to the t-statistic when we were comparing just two means. $t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{\frac{s_1^2}{n_1} + \frac{s_2^2}{n_2}}}$

Recalling Ambiguous Prose

Example 9.2

Comprehension Example

(**Don't** follow along in your book or look ahead on the PowerPoint until after I read you the passage.)

- Students were read an ambiguous prose passage under one of the following conditions:
 - Students were given a picture that could help them interpret the passage **before** they heard it.
 - Students were given the picture **after** they heard the passage.
 - Students were **not** shown any picture before or after hearing the passage.
- They were then asked to evaluate their comprehension of the passage on a 1 to 7 scale.

Comprehension Example

- This experiment is a partial replication done at Hope College of a study done by Bransford and Johnson (1972).
- Students were randomly assigned to one of the 3 groups.
- Listen to the passage and see if it makes sense. Would a picture help?

If the balloons popped, the sound wouldn't be able to carry since everything would be too far away from the correct floor. A closed window would also prevent the sound from carrying, since most buildings tend to be well insulated. Since the whole operation depends on a steady flow of electricity, a break in the middle of the wire would also cause problems. Of course, the fellow could shout, but the human voice is not loud enough to carry that far. An additional problem is that a string could break on the instrument. Then there could be no accompaniment to the message. It is clear that the best situation would involve less distance. Then there would be fewer potential problems. With face to face contact, the least number of things could go wrong.



Hypotheses

- **Null:** In the population there is no association between whether or when a picture was shown and comprehension of the passage
- **Alternative:** In the population there is an association between whether and when a picture was shown and comprehension of the passage

Hypotheses


- **Null:** All three of the long term mean comprehension scores are the same.

$$\mu_{\text{no picture}} = \mu_{\text{picture before}} = \mu_{\text{picture after}}$$

- **Alternative:** At least one of the mean comprehension scores is different.

Recall Score

- Students rated their comprehension, and the researchers also had the students recall as many ideas from the passage as they could. They were then graded on what they could recall and the results are shown.



Summary Statistics:

	n	Mean	SD
None	19	6.63	2.01
Before	19	8.26	1.82
After	19	5.37	1.46
Pooled	57	6.75	1.78

Observed MAD = 1.930

Validity Conditions

- Just as with the simulation-based method, we are assuming we have independent groups.
- Two extra conditions must be met to use traditional ANOVA:
 - Normality: If sample sizes are small within each group, data shouldn't be very skewed. If it is, use simulation approach. (Sample sizes of at least 30 is a good guideline.)
 - Equal variation: Largest standard deviation should be no more than twice the value of the smallest.

ANOVA Output

- This is the kind of output you would see in most statistics packages when doing ANOVA.
- The variability between the groups is measured by the mean square treatment (40.02).
- The variability within the groups is measured by the mean square error (3.16).
- The F statistic is $40.02/3.16 = 12.67$.

Source	df	SS	MS	F	p-value
Treatment	2	80.04	40.02	12.67	0.0000
Error	54	170.53	3.16		
Total	56	250.56			

Conclusion

- Since we have a small p-value we have strong evidence against the null and can conclude at least one of the long-run mean recall scores is different.

Review list.

1. Meaning of SD.
2. Parameters and statistics.
3. Z statistic for proportions.
4. Simulation and meaning of pvalues.
5. SE for proportions.
6. What influences pvalues.
7. CLT and validity conditions for tests.
8. 1-sided and 2-sided tests.
9. Reject the null vs. accept the alternative.
10. Sampling and bias.
11. Significance level.
12. Type I, type II errors, and power.
13. CIs for a proportion.
14. CIs for a mean.
15. Margin of error.
16. Practical significance. (causation, extrapolation, curvature, heteroskedasticity).
17. Confounding.
18. Observational studies and experiments.
19. Random sampling and random assignment.
20. Two proportion CIs and testing.
21. IQR and 5 number summaries.
22. CIs for 2 means and testing.
23. Paired data.
24. Placebo effect, adherer bias, and nonresponse bias.
25. Prediction and causation.
26. Multiple testing and publication bias
27. Regression.
28. Correlation.
29. Calculate & interpret a & b.
30. Goodness of fit for regression.
31. Common regression problems
32. ANOVA and F-test.

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

1. What does the correlation of 0.82 imply?
 - a. 82% of the variation in weight is explained by height.
 - b. The typical variation in people's heights is 82% as large as the typical variation in their weights.
 - c. There is strong association between height and weight in this sample.
 - d. For every inch of increase in one's height, we would predict a 0.82 lb. increase in weight.
 - e. If a person weighs 100 pounds, then we typically would expect the person to be about 82 inches tall.

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

1. What does the correlation of 0.82 imply?

a. 82% of the variation in weight is explained by height.

b. The typical variation in people's heights is 82% as large as the typical variation in their weights.

c. There is strong association between height and weight in this sample.

d. For every inch of increase in one's height, we would predict a 0.82 lb. increase in weight.

e. If a person weighs 100 pounds, then we typically would expect the person to be about 82 inches tall.

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

2. What is the estimated slope, in lbs/inch, of the regression line for predicting weight from height?

a. 6.56. b. 7.12. c. 8.04. d. 9.92. e. 10.2. f. 11.4.

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

2. What is the estimated slope, in lbs/inch, of the regression line for predicting weight from height?

a. 6.56. b. 7.12. c. 8.04. d. 9.92. e. 10.2. f. 11.4.

$$r s_y / s_x = .82 \times 40 / 5 = 6.56.$$

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

3. How much would a prediction using this regression line typically be off by?

- a. 12.7 lbs. b. 13.5 lbs. c. 14.4lbs. d. 20.2 lbs. e. 22.9 lbs.

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

3. How much would a prediction using this regression line typically be off by?

- a. 12.7 lbs. b. 13.5 lbs. c. 14.4lbs. d. 20.2 lbs. e. **22.9 lbs.**

$$\sqrt{(1-r^2)} s_y = \sqrt{(1-.82^2)} \times 40 = 22.9.$$

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

4. If we were to randomly take one adult from this sample, how much would his/her height typically differ from 65 by?

a. 0.05 in. b. 0.1 in. c. 0.5 in. d. 1.0 in. e. 2.5 in. f. 5.0 in.

example problems.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

4. If we were to randomly take one adult from this sample, how much would his/her height typically differ from 65 by?

a. 0.05 in. b. 0.1 in. c. 0.5 in. d. 1.0 in. e. 2.5 in. **f. 5.0 in.**

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the median height is 64.5 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

5. Why shouldn't one trust this regression line to predict the weight of someone who is 25 inches tall?
- a. The sample size is insufficiently large.
 - b. The sample SD of weight is too small.
 - c. The value of 25 inches is too far outside the range of most observations.
 - d. The correlation of the ANOVA is a t-test confidence interval with statistical significance.
 - e. The data come from an observational study, so there may be confounding factors.
 - f. The height values are heavily right skewed, so the prediction errors are large.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the median height is 64.5 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs.

5. Why shouldn't one trust this regression line to predict the weight of someone who is 25 inches tall?
- a. The sample size is insufficiently large.
 - b. The sample SD of weight is too small.
 - c. The value of 25 inches is too far outside the range of most observations.**
 - d. The correlation of the ANOVA is a t-test confidence interval with statistical significance.
 - e. The data come from an observational study, so there may be confounding factors.
 - f. The height values are heavily right skewed, so the prediction errors are large.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the median height is 64.5 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs. The estimated slope in predicting weight from height is 6.56 lbs/in.

6. How should one interpret the estimated slope of 6.56?

- a. Each extra inch you grow causes you to increase your weight by 6.56 lbs on average.
- b. Each extra lb. you weigh causes you to grow 6.56 inches.
- c. The amount of weight Americans average is 6.56 standard errors above the mean.
- d. The Z-score corresponding to the correlation between height and weight is 6.56.
- e. For each extra inch taller you are, your predicted weight increases by 6.56 lbs.
- f. The proportion of variance in weight explained by the regression equation is 6.56%.

Suppose that among a sample of 100 adults in a given town, the correlation between height (inches) and weight (lbs.) is 0.82, and the mean height is 65 inches, the median height is 64.5 inches, the sd of height is 5 inches, the mean weight is 160 lbs., and the sd of weight is 40 lbs. The estimated slope in predicting weight from height is 6.56 lbs/in.

6. How should one interpret the estimated slope of 6.56?

- a. Each extra inch you grow causes you to increase your weight by 6.56 lbs on average.
- b. Each extra lb. you weigh causes you to grow 6.56 inches.
- c. The amount of weight Americans average is 6.56 standard errors above the mean.
- d. The Z-score corresponding to the correlation between height and weight is 6.56.
- e. For each extra inch taller you are, your predicted weight increases by 6.56 lbs.**
- f. The proportion of variance in weight explained by the regression equation is 6.56%.

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

7. What is the pooled sample percentage with acne, in both groups combined?
- a. 16.1%. b. 17.2%. c. 18.4%. d. 19.7%. e. 21.2%. f. 23.0%.

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

7. What is the pooled sample percentage with acne, in both groups combined?
a. 16.1%. b. 17.2%. c. 18.4%. **d. 19.7%.** e. 21.2%. f. 23.0%.

$$118/600 = 19.7\%.$$

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

8. Using this pooled sample percentage, under the null hypothesis that the two groups have the same acne rate, what is the standard error for the difference between the two percentages?

- a. 1.42%. b. 1.88%. c. 2.02%. d. 2.99%. e. 3.08%. f. 3.44%.

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

8. Using this pooled sample percentage, under the null hypothesis that the two groups have the same acne rate, what is the standard error for the difference between the two percentages?

a. 1.42%. b. 1.88%. c. 2.02%. d. 2.99%. e. 3.08%. **f. 3.44%.**

$$\sqrt{(.197 * (1-.197)/400 + .197*(1-.197)/200)} = .0344.$$

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

9. What is the Z statistic for the difference between the two group percentages?

- a. 1.02. b. 1.23. c. 1.55. d. 1.88. e. 2.03. f. 2.43.

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

9. What is the Z statistic for the difference between the two group percentages?

- a. 1.02. b. 1.23. c. 1.55. d. 1.88. **e. 2.03.** f. 2.43.

$$(88/400 - 30/200) \div 0.0344 = 2.03.$$

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

10. Using the unpooled standard error for the difference between the two percentages, find a 95% confidence interval for the percentage with acne among those age 21-30 minus the percentage of acne among those age 31-40.

a. 7% +/- 5.02%. b. 7% +/- 5.54%. c. 7% +/- 5.92%. d. 7% +/- 6.03%. e. 7% +/- 6.41%.

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

10. Using the unpooled standard error for the difference between the two percentages, find a 95% confidence interval for the percentage with acne among those age 21-30 minus the percentage of acne among those age 31-40.

a. 7% +/- 5.02%. b. 7% +/- 5.54%. c. 7% +/- 5.92%. d. 7% +/- 6.03%. **e. 7% +/- 6.41%.**

$$\begin{aligned} & 88/400 - 30/200 \pm 1.96 * \sqrt{(88/400 * (1-88/400) / 400 + 30/200 * (1-30/200) / 200)} \\ & = 7\% \pm 6.41\%. \end{aligned}$$

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

11. What can we conclude from the 95% CI in the problem above?

- a. There is no statistically significant difference in the percentage with acne among those age 21-30 and those age 31-40.
- b. There is statistically significant correlation between the sample size and the effect size for the confounding factors of a randomized controlled experiment.
- c. Of those with acne, there is no statistically significant difference between the percentage who are age 21-30 and the percentage who are 31-40.
- d. A statistically significantly higher percentage of people age 21-30 have acne than those age 31-40.
- e. The sample sizes are too small for a conclusion to be valid here.

Suppose a researcher studying acne, or oily skin, takes a simple random sample of 400 Americans age 21-30 and calls them group A, and a simple random sample of 200 Americans age 31-40 and calls them group B. In group A, 88 people have acne, and in group B, 30 people have acne.

11. What can we conclude from the 95% CI in the problem above?

- a. There is no statistically significant difference in the percentage with acne among those age 21-30 and those age 31-40.
- b. There is statistically significant correlation between the sample size and the effect size for the confounding factors of a randomized controlled experiment.
- c. Of those with acne, there is no statistically significant difference between the percentage who are age 21-30 and the percentage who are 31-40.
- d. A statistically significantly higher percentage of people age 21-30 have acne than those age 31-40.**
- e. The sample sizes are too small for a conclusion to be valid here.

For the following two problems, suppose a researcher studies a treatment for a certain condition. The researcher divides 230 subjects randomly into three groups. Group A receives a placebo. Group B receives the treatment at a low dose, and group C receives the treatment at a high dose. The table below shows the output from an ANOVA F-test on the mean levels of the condition in the 3 groups.

Analysis of Variance Table

Response: outcome

	Df	Sum Sq	MeanSq	F value	Pr(>F)
group	2	0.467	0.233	0.221	0.803
Residuals	227	28.5	1.06		

12. In the table, what number is a measure of the variability between groups?

a. 2. b. 0.233 c. 0.221 d. 0.803 e. 28.5 f. 1.06.

For the following two problems, suppose a researcher studies a treatment for a certain condition. The researcher divides 230 subjects randomly into three groups. Group A receives a placebo. Group B receives the treatment at a low dose, and group C receives the treatment at a high dose. The table below shows the output from an ANOVA F-test on the mean levels of the condition in the 3 groups.

Analysis of Variance Table

Response: outcome

	Df	Sum Sq	MeanSq	F value	Pr(>F)
group	2	0.467	0.233	0.221	0.803
Residuals	227	28.5	1.06		

12. In the table, what number is a measure of the variability between groups?

a. 2. **b. 0.233** c. 0.221 d. 0.803 e. 28.5 f. 1.06.

For the following two problems, suppose a researcher studies a treatment for a certain condition. The researcher divides 230 subjects randomly into three groups. Group A receives a placebo. Group B receives the treatment at a low dose, and group C receives the treatment at a high dose. The table below shows the output from an ANOVA F-test on the mean levels of the condition in the 3 groups.

Analysis of Variance Table

Response: outcome

	Df	Sum Sq	MeanSq	F value	Pr(>F)
group	2	0.467	0.233	0.221	0.803
Residuals	227	28.5	1.06		

13. What can we conclude from the result of the F test?

- a. The treatment has a statistically significant effect, though we cannot be sure which dose is attributable to this significant effect or in which direction the effect goes.
- b. The treatment seems to have significantly greater effect at large doses than at small doses.
- c. The treatment does not seem to have a statistically significant effect.
- d. We fail to reject the null hypothesis that the treatment has a significant effect on the condition.
- e. The correlation between dose and outcome appears to be statistically significant.

For the following two problems, suppose a researcher studies a treatment for a certain condition. The researcher divides 230 subjects randomly into three groups. Group A receives a placebo. Group B receives the treatment at a low dose, and group C receives the treatment at a high dose. The table below shows the output from an ANOVA F-test on the mean levels of the condition in the 3 groups.

Analysis of Variance Table

Response: outcome

	Df	Sum Sq	MeanSq	F value	Pr(>F)
group	2	0.467	0.233	0.221	0.803
Residuals	227	28.5	1.06		

13. What can we conclude from the result of the F test?

- a. The treatment has a statistically significant effect, though we cannot be sure which dose is attributable to this significant effect or in which direction the effect goes.
- b. The treatment seems to have significantly greater effect at large doses than at small doses.
- c. The treatment does not seem to have a statistically significant effect.**
- d. We fail to reject the null hypothesis that the treatment has a significant effect on the condition.
- e. The correlation between dose and outcome appears to be statistically significant.