AP Statistics 1-Day Workshop

David Ferris, College Board Consultant
dave_ferris@nobl.k12.in.us www.noblestatman.com

Workshop Warm-ups:

1. Take Socrative survey at socrative.com or upload and launch the Socrative (Student) app. Click on Student login…Room Name is **NOBLESTATMAN**

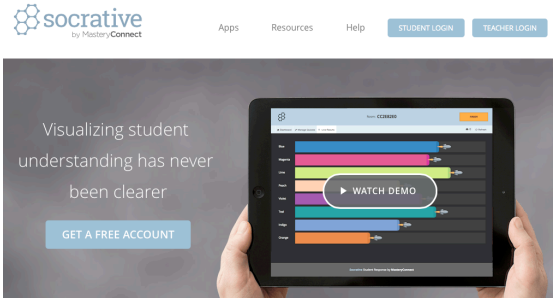

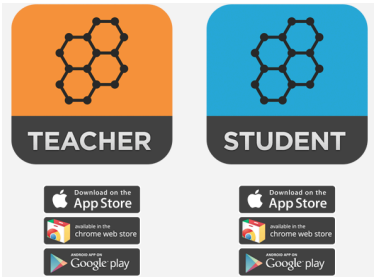

2. Check out resources at **thisisstatistics.org**

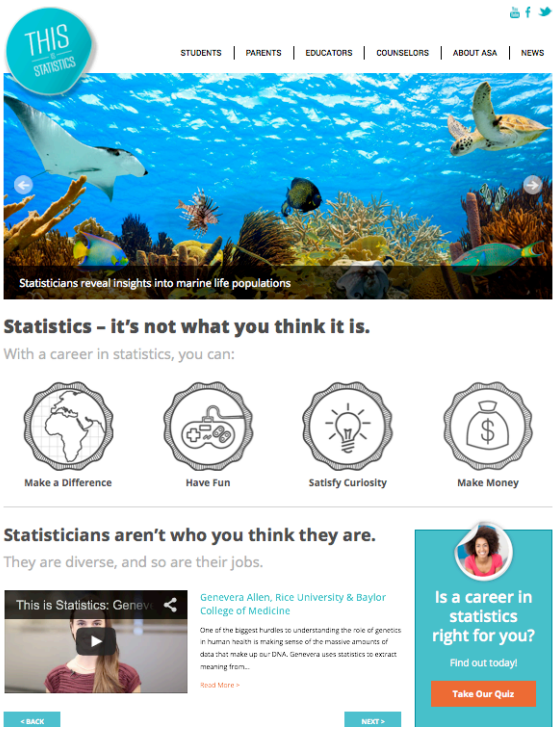

3. The Downtown Marriott in Kansas City employs an elevator assignment system to increase efficiency. There are 5 elevators, so the chances of randomly picking the elevator you will be assigned is 20%. A random sample of 20 elevator rides was taken, and the rider correctly predicted only 2 using a random number generator. Is this evidence that the software is not randomly assigning elevators?

4. Verify the probability claims in this article.

POINT AFTER

America's Wait Problem

→ BY MICHAEL ROSENBERG

Tim Hudson stood in the Giants' clubhouse Sunday, talking about pitching in his first World Series in a 16-year major league career, and he knew he was lucky. Fans in Section 334, in the upper deck in leftfield of AT&T Park, had traveled all the way from Kansas City to San Francisco just to chant "Let's Go, Royals!" After a 28-year World Series drought, they knew they were lucky too.

But did anybody realize *how* lucky?

Championship droughts are not what they used to be. They are longer, drier and more painful. Every major sports league has expanded to at least 30 teams, but they all still only give out one championship trophy per season. In many cities sports fans will be on hold. Enjoy the music while your party is delayed.

We made such a big deal of Kansas City making the playoffs for the first time since 1985 because it's been a really long time, and for most of that stretch the Royals were terrible. There were moments when they did not seem to understand things like the infield fly rule or why home plate was where it was. A 29-year gap between playoff appearances may be unusual, but in an era of 30-team leagues, 29 years between championships is nothing.

There are 30 major league teams (and 30 in the NHL and NBA too; the NFL has 32). Let's assume leagues resist the urge to further expand. If every league distributes championships equally, the average wait for a title will be 30 years. But that's just the average wait. Factor in variables such as competence and some teams will be cooling their, say, Jets longer than others.

So how long would you have to wait for your parade? Despite my advanced math skills, I found somebody else to answer the question: Jeffrey Rosenthal, a professor of statistics at the University of Toronto and the author of the bestselling book *Struck by Lightning: The Curious World of Probabilities.*

Rosenthal completed some statistical calculations under the assumption that all 30 teams had an equal chance of winning. The result: a team would have a 25.8% chance of going without a title for 40 years, with an 18.4% chance that the drought would extend to 50 seasons. The probability of not seeing your team win another World Series before 2100 would be 5.4%.

Sound bad? Reality is worse. That model assumes all teams have the same shot. But if all things were equal, cucumbers would taste like chocolate cake. We don't live in that world. We live in a world where the Spurs somehow land both David Robinson and Tim Duncan, brilliant coach Bill Belichick finds brilliant quarterback Tom Brady, LeBron James teams up with his buddies in Miami and the Yankees charge fans $100 for Officially Licensed Used Tobacco, then spend that money on payroll. Some teams will be consistently better than others, for reasons that are mental, financial, geographical or coincidental. That sound you hear is Cubs fans crying.

> We live in an era of **instant everything,** but sports fans have to wait longer than ever for their share of happiness.

In Rosenthal's model, the most likely number of champions in an 18-year span is 14. In the 18 baseball seasons before this one, 10 franchises won championships. The chance of that happening, if all teams were equal, is 1.0%.

We live in an era of instant everything, but many sports fans have to wait longer than ever for their share of happiness. Millions of us can expect to spend at least 40 years in the desert, hoping there is good Wi-Fi there.

The Royals have a young team, with more talent on the way, but that only guarantees pretty preseason predictions. And Giants fans should enjoy this era because they are unlikely to see anything like it again. The Giants went 56 years without a championship, then won in 2010 and '12 before making the World Series again this season. The 56-year wait seemed excruciating. In the future it will seem normal.

Still, you should keep hope alive, and while you're at it, keep yourself alive. Americans are now expected to live an average of 78.8 years. I'm not saying there is a good chance you will die before your favorite team wins a championship. But I strongly advise you to eat your vegetables. □

Which sports drought will end next? Join the discussion on Twitter by using **#SIPointAfter** and following **@Rosenberg_Mike**

CARLOS M. SAAVEDRA FOR SPORTS ILLUSTRATED

Textbooks:

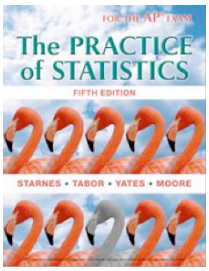

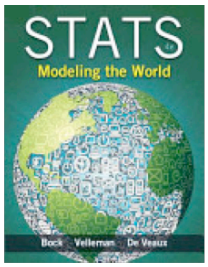

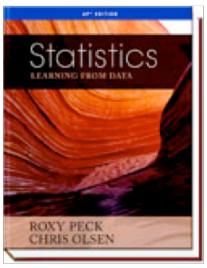

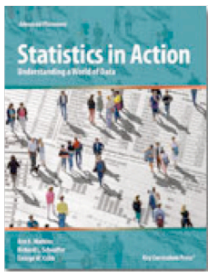

Other books:

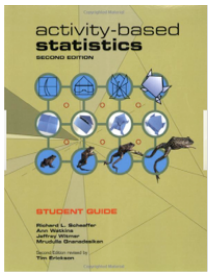

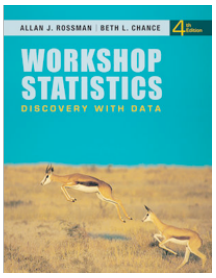

Web Site List:

Resources

noblestatman.com ("ehandouts" and other resources)

AP Central (bookmark course **home page** and **exam page**)

apstatsmonkey.com (clearinghouse for many useful resources)

amstat.org (Census At School, STEW lesson plans)

www.causeweb.org/sbi (Simulation Based Inference discussions/blog. This is a "trending" topic among high school and college statistics teachers.)

openintro.org/stat (free pdfs of statistics textbooks)

fathom.concord.org (Fathom is the best software for ***teaching*** statistics)

Applets, demos, games, simulations

Rossman Chance applets (many good simulation applets)

onlinestatbook.com/stat_sim/sampling_dist/index.html (sampling distribution)

StatCrunch (teacher account is free, student can subscribe for minimal fee)

Against All Odds statistics videos (can stream for free--learner.org)

socrative.com (formative assessment tool)

getkahoot.com (create engaging online review)

quizizz.com (create engaging online review)

StatKey (simulation website app)

Classifying Statistics Problems (ltcconline; practice choosing correct procedure)

fivethirtyeight.com (great current articles and graphs with a statistical slant)

tuvalabs.com (online tool for analyzing distributions and scatterplots)

gapminder.org (amazing online analysis tool of United Nations data)

Collecting Student Data:

socrative.com

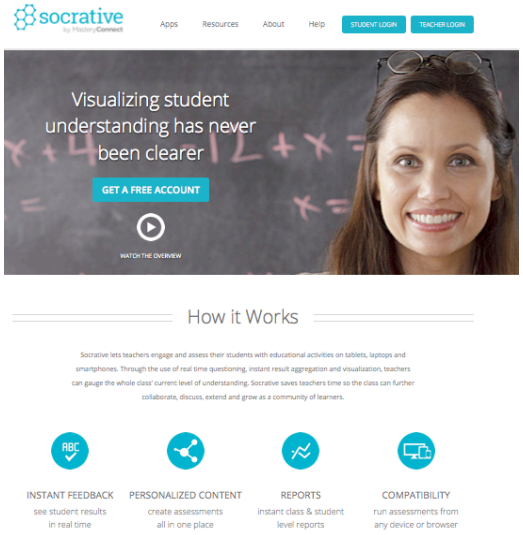

statcrunch.com

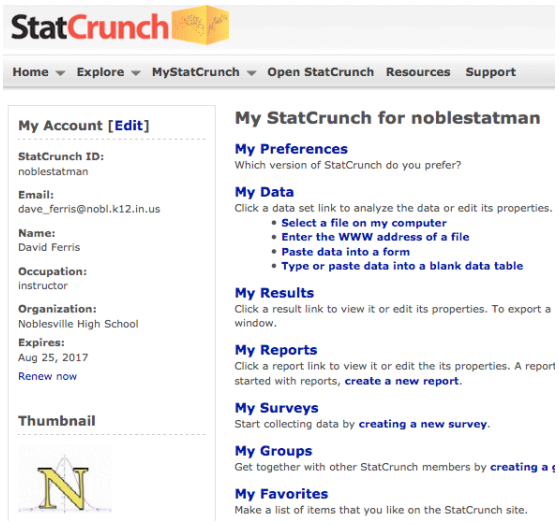

Google Forms

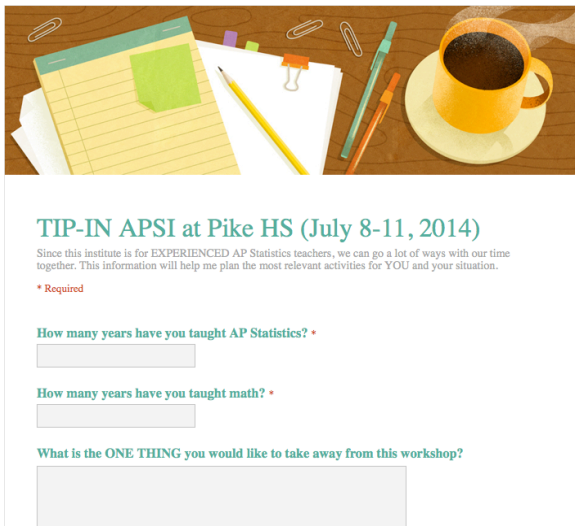

nearpod.com (for iPads, tablets, laptops)

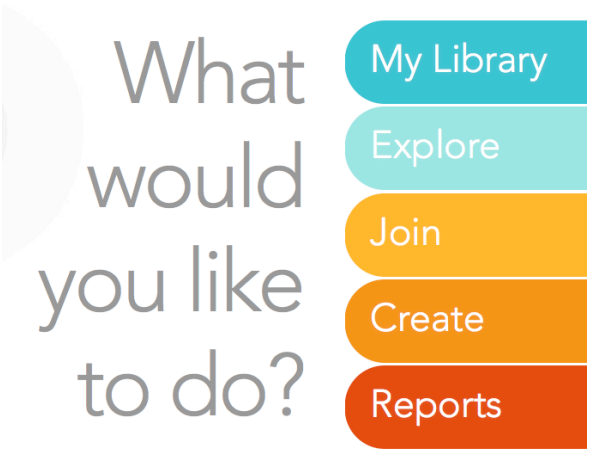

Online Data Analysis:

tuvalabs.com

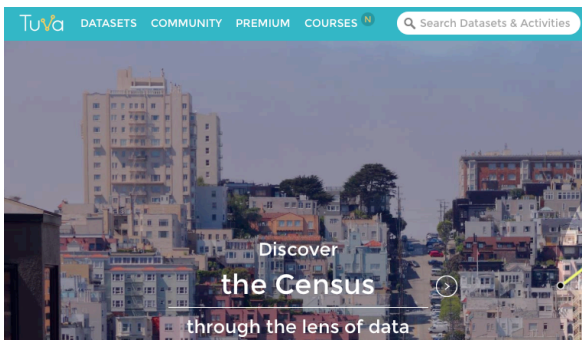

lock5stat.com/statkey/

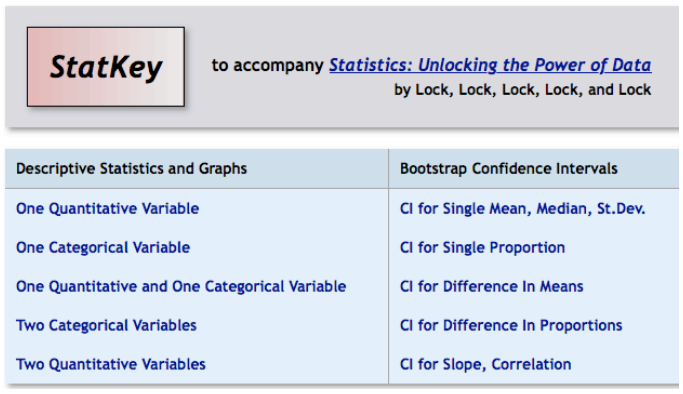

(Use Safari) Class ID: _____ ***Password:*** _____

www.amstat.org/censusatschool

Questions to be completed before taking the online survey:

4. How tall are you without your shoes on? Answer to the nearest centimeter. _____

5. What is the length of your right foot (without your shoe on)? Answer to the nearest centimeter.

_____

6. What is your arm span? (Open arms wide and measure distance across your back from tip of right hand middle finger to tip of left hand middle finger.) Answer to the nearest centimeter.

_____

9. How long does it usually take you to travel to school? Answer to the nearest minute. _____

14. What is the length of your left foot (without your shoe on)? Answer to the nearest centimeter.

_____

16. What is the length of your index finger (finger next to your thumb) on your left hand? Answer to the nearest centimeter.

_____

17. What is the length of your ring finger? (located between your middle finger and little finger) on your left hand? Answer to the nearest millimeter (there are 10 millimeters in one centimeter).

_____

26. How many hours of sleep do you usually get when you have school the next day? _____

27. How many hours of sleep do you usually get when you don't have school the next day? _____

Trick or Treat!

Fun size M&M's and Skittles NAME_____

1. Question: How many MM's are there in a Fun Size bag? Guess:_____

   Graph the class guesses on the back side of this sheet.

2. Number of M&M's in your bag: _____

3. Make a dot plot of the class data on the back side of this sheet.

4. What was the population of interest?

5. What was the sample?

6. What was the observational unit (subject)?

7. So how many M&M's are in a fun size bag?

8. Do you think the typical number of Skittles in a Fun Size bag is more than the typical number of M&M's in a Fun Size bag? Explain.

9. Number of Skittles in your bag: _____

10. Make a dot plot of the class data on the back side of this sheet.

11. So how many Skittles are in a fun size bag?

Dot plot of class guesses:

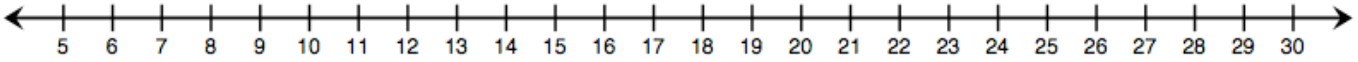

Dot plot of class bags of **M&M's**:

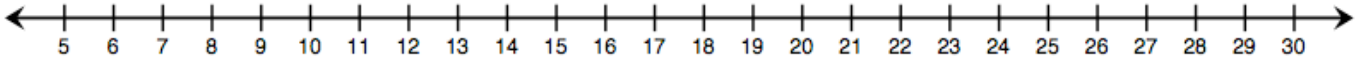

Dot plot of class bags of **Skittles**:

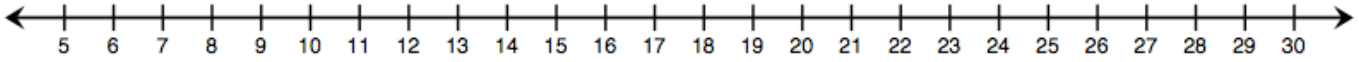

12. Compare and contrast the three distributions in context. Be sure to talk about the shapes, centers, spreads and outliers.

2016 AP Statistics Exam

1. Robin works as a server in a small restaurant, where she can earn a tip (extra money) from each customer she serves. The histogram below shows the distribution of her 60 tip amounts for one day of work.

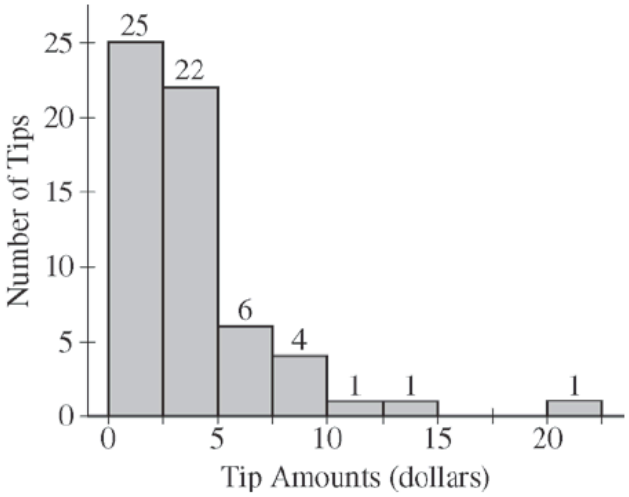

(a) Write a few sentences to describe the distribution of tip amounts for the day shown.

(b) One of the tip amounts was $8. If the $8 tip had been $18, what effect would the increase have had on the following statistics? Justify your answers.

The mean:

The median:

Student Performance Q&A:
**(Chief Reader's Report)
2016 AP® Statistics Free-Response Questions

The following comments on the 2016 free-response questions for AP® Statistics were written by the Chief Reader, Jessica Utts of the University of California, Irvine. They give an overview of each free- response question and of how students performed on the question, including typical student errors. General comments regarding the skills and content that students frequently have the most problems with are included. Some suggestions for improving student performance in these areas are also provided. Teachers are encouraged to attend a College Board workshop to learn strategies for improving student performance in specific areas.

Question #1:

Based on your experience of student responses at the AP® Reading, what message would you like to send to teachers that might help them to improve the performance of their students on the exam?

Teachers should give more exposure to ***analyzing distributions*** of quantitative data based on graphical displays in which some or most of the data values are unknown.

Teachers should provide more opportunities to discuss distributions ***in context*** without simply listing features, and stress that context is about the variable(s) of interest and not the units of measurement for those variable(s).

Teachers should give students ample practice ***writing justifications*** and receiving ***critical feedback*** about them.

Teachers should be sure students can distinguish among the effects of the following changes to the data:

- ***Changing one value*** in a data set.
- ***Adding or removing a value*** from a data set.
- ***Transforming*** all the values in a data set via addition, subtraction, multiplication, division.

Teachers should build in ***spiral review*** throughout the course, revisiting topics and emphasizing connections, comparisons and contrasts. Some students seem to have forgotten the basics from the beginning of the course by the time they reach the end.

Entire document can be found at AP Central:
https://secure-media.collegeboard.org/digitalServices/pdf/ap/ap16_statistics_student_performance_qa.pdf

FRAPPY!

NAME_____

Comparing Distributions: 2015 #1

1. Two large corporations, A and B, hire many new college graduates as accountants at entry-level positions. In 2009 the starting salary for an entry-level accountant position was $36,000 a year at both corporations. At each corporation, data were collected from 30 employees who were hired in 2009 as entry-level accountants and were still employed at the corporation five years later. The yearly salaries of the 60 employees in 2014 are summarized in the boxplots below.

(a) Write a few sentences comparing the distributions of the yearly salaries at the two corporations.

(b) Suppose both corporations offered you a job for $36,000 a year as an entry-level accountant.

(i) Based on the boxplots, give one reason why you might choose to accept the job at corporation A.

(ii) Based on the boxplots, give one reason why you might choose to accept the job at corporation B.

FRAPPY!

NAME_____

Student Solutions: 2015 #1

Student Sample C:

Two large corporations, A and B, hire many new college graduates as accountants at entry-level positions. In 2009 the starting salary for an entry-level accountant position was $36,000 a year at both corporations. At each corporation, data were collected from 30 employees who were hired in 2009 as entry-level accountants and were still employed at the corporation five years later. The yearly salaries of the 60 employees in 2014 are summarized in the boxplots below.

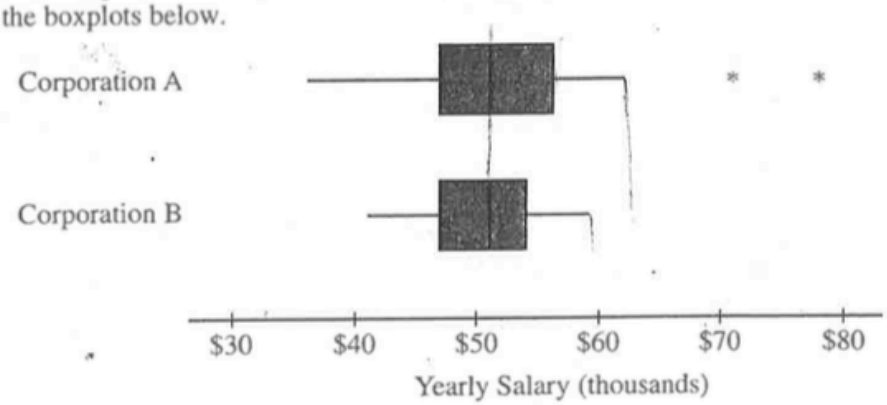

(a) Write a few sentences comparing the distributions of the yearly salaries at the two corporations.

- The shape of both boxplots appear to be roughly symmetric
- Corporation A appears to have 2 outliers, while corporation B appears to have no outliers
- While corporation A and corporation B appear to have the same median salary at about $50,000, the interquartile range appears to be larger for corporation a than corporation B.

(b) Suppose both corporations offered you a job for $36,000 a year as an entry-level accountant.

(i) Based on the boxplots, give one reason why you might choose to accept the job at corporation A.

Based on the boxplots, I might choose to accept a job at corporation A because it has a greater maximum salary than corporation B, making it a higher chance of me earning more money in the future.

(ii) Based on the boxplots, give one reason why you might choose to accept the job at corporation B.

I would choose to accept a job at corporation B because the minimum salary is greater than the minimum salary than corporation A.

Student Sample F:

1. Two large corporations, A and B, hire many new college graduates as accountants at entry-level positions. In 2009 the starting salary for an entry-level accountant position was $36,000 a year at both corporations. At each corporation, data were collected from 30 employees who were hired in 2009 as entry-level accountants and were still employed at the corporation five years later. The yearly salaries of the 60 employees in 2014 are summarized in the boxplots below.

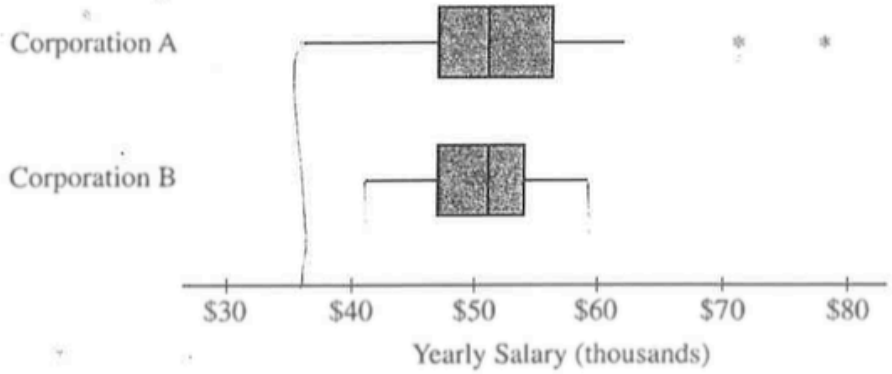

(a) Write a few sentences comparing the distributions of the yearly salaries at the two corporations.

Corp A

Center at about $50

Range (spread) of about $42 ($78 - $36)

The distribution is skewed right

Corp B

- Center at about $51
- Range (spread) of about $17 ($59-$42)
- The distribution is roughly symmetrical

(b) Suppose both corporations offered you a job for $36,000 a year as an entry-level accountant.

(i) Based on the boxplots, give one reason why you might choose to accept the job at corporation A.

Corp A would be a good choice because there is a possibility of earning over $60,000 per year in 5 years time. That is impossible as an employee in Corp B

(ii) Based on the boxplots, give one reason why you might choose to accept the job at corporation B.

Corp B would be a good choice because you are garunteed to make over $40,000 per year after 5 years. As an employee at corp A it is possible to make less than that after 5 years.

Scoring Guidelines:

Part (a):

The median salary is approximately the same for both corporations. The range and interquartile range of the salaries are greater for Corporation A than for Corporation B. The two highest salaries at Corporation A are outliers while Corporation B has no outliers.

Part (b):

(i) Five years after starting, at least 3 out of 30 (10%) of the salaries at Corporation A are greater than the maximum salary at Corporation B. If I accept the offer from Corporation A, I might be able to make a higher salary at Corporation A than at Corporation B.

(ii) Five years after starting, the minimum salary at Corporation B is greater than at Corporation A. In fact, at Corporation A it looks like some people are still making the starting salary of $36,000 and never received a raise in the five years since they were hired. So if I work at Corporation A, I might never receive a raise in salary.

Part (a) is scored as follows:

Essentially correct (E) if the response includes the following four components:

1. A correct comparison of center.
2. A correct comparison of spread.
3. A discussion of the outliers for Corporation A.
4. The response is in context.

Partially correct (P) if the response includes only three of the four components.

Incorrect (I) if the response includes at most two of the four components.

Note: Any mention of shape should be ignored because complete shape information cannot be determined from a boxplot.

Part (b) is scored as follows:

Essentially correct (E) if the response includes the following four components:

1. In part (b-i) a relevant statistical measure is identified (or described) or a relevant statistical comparison is provided that supports the choice of Corporation A.
2. In part (b-i) an explanation is provided for why the measure or comparison is relevant.
3. In part (b-ii) a relevant statistical measure is identified (or described) or a relevant statistical comparison is provided that supports the choice of Corporation B.
4. In part (b-ii) an explanation is provided for why the measure or comparison is relevant.

Partially correct (P) if the response includes only two or three of the four components.

Incorrect (I) if the response includes none or one of the four components.

Note: If a response does not provide a statistical measure or comparison in part (b-i) or (b-ii), the second and fourth components can still be satisfied if an acceptable explanation is provided that would follow from a relevant statistical measure or comparison. For example, if the response in part (b-i) only states "At Corporation A, I have the potential to earn a higher salary," the second component is satisfied.

AP Statistics Exam General Information:

- The AP Statistics Exam is Thursday, May 11, 2017 from 12-4 PM
- There are _____ multiple choice questions on the AP Statistics Exam.
- There are _____ free response questions on the AP Statistics Exam.
- Each part weighs 50% (scaled to _____ points each)
- Students have _____ minutes for each section of the exam.
    - #_____ of the Free Response part is called the Investigative Task.
    - Students should reserve about _____ minutes to complete this problem.
    - It is worth _____ of the free response section, or _____ of the entire test!
    - It typically covers several _____ _____.
    - It typically introduces something _____.
    - Students should have a strategy:
        1. S_____/s_____ the entire test, and r_____ problems.
        2. Find the two _____ and do those first (≈20 minutes)
        3. **Then do #6 (up to 25 minutes)**
        4. Then spend the rest of your time (≈45 minutes) on the last three.

The investigative Task is designed to _____ you. Try to _____ read, think and give your best answers.

_____ credit earned on the Investigative Task will help your overall score!

Mean scores on recent Investigative Tasks (2011-2016): 1.31, 1.50, 2.14, 1.29, 1.08, 1.61

Age Guessing Activity

1. Guess the ages of the following people, and put your guess in this column:

Name:	Actual Age:	Your Guess:
Barack Obama	_____	_____
Harrison Ford	_____	_____
Bill Gates	_____	_____
Johnny Depp	_____	_____
Bono	_____	_____
Alex Trebek	_____	_____
Oprah Winfrey	_____	_____
Miley Cyrus	_____	_____
Peyton Manning	_____	_____
J. K. Rowling	_____	_____
Mick Jagger	_____	_____
Mark Zuckerberg	_____	_____

2. Put the actual ages of each person in the first column.

3. Type both lists into your calculator. (Actual ages in L1, your guesses in L2.)

4. Make a scatterplot for these two lists. (x-axis is L1, y-axis is L2)

5. Calculate r: _____

6. Describe below what you discovered about your age guesses:

Commonly Asked Regression Questions

(as seen on previous AP Statistics exams)

1. Describe the association in context.

2. Is a linear model appropriate to describe this relationship? Explain.

3. Write the equation for the linear model on this data.

4. Explain the meaning of the slope in this linear model

5. Explain the meaning of the y-intercept in this linear model

6. Find the value and explain the meaning of the correlation coefficient.

7. Find the value of and interpret r-squared

8. a. Using the linear model, predict _____ when _____ is _____.

   b. Is the **residual for this data point** positive or negative? Is the model over- or underestimating? Explain.

9. Comment on any outliers present. Fully describe their effect on analyses, if any.

10. Interpret regression and model information from a computer printout.

2016 AP Statistics Exam

6. A newspaper in Germany reported that the more semesters needed to complete an academic program at the university, the greater the starting salary in the first year of a job. The report was based on a study that used a random sample of 24 people who had recently completed an academic program. Information was collected on the number of semesters each person in the sample needed to complete the program and the starting salary, in thousands of euros, for the first year of a job. The data are shown in the scatterplot below.

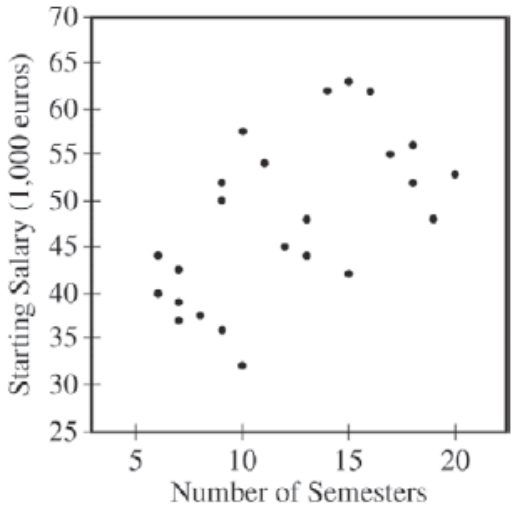

(a) Does the scatterplot support the newspaper report about number of semesters and starting salary? Justify your answer.

The table below shows computer output from a linear regression analysis on the data.

Predictor	Coef	SE Coef	T	P
Constant	34.018	4.455	7.64	0.000
Semesters	1.1594	0.3482	3.33	0.003
S = 7.37702	R-Sq = 33.5%		R-Sq(adj) = 30.5%	

(b) Identify the slope of the least-squares regression line, and interpret the slope in context.

An independent researcher received the data from the newspaper and conducted a new analysis by separating the data into three groups based on the major of each person. A revised scatterplot identifying the major of each person is shown below.

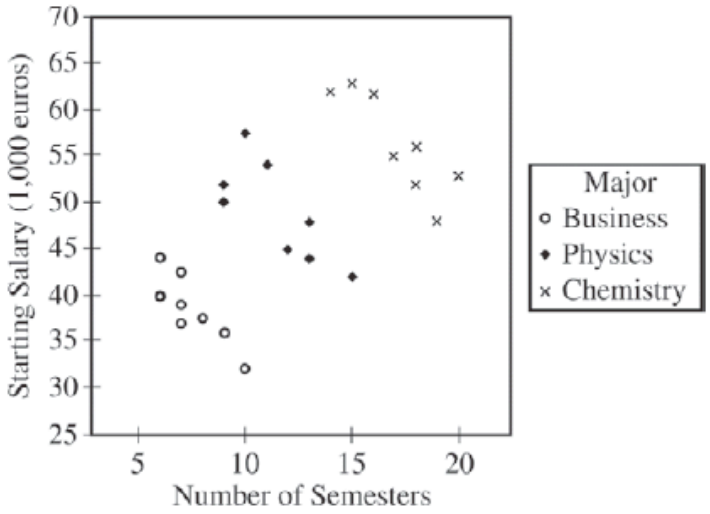

(c) Based on the people in the sample, describe the association between starting salary and number of semesters for the business majors.

(d) Based on the people in the sample, compare the median starting salaries for the three majors.

(e) Based on the analysis conducted by the independent researcher, how could the newspaper report be modified to give a better description of the relationship between the number of semesters and the starting salary for the people in the sample?

2015 #5: (condensed here…a really good problem for later in year, perhaps when reviewing for exam)

5. A student measured the heights and the arm spans, rounded to the nearest inch, of each person in a random sample of 12 seniors at a high school. A scatterplot of arm span versus height for the 12 seniors is shown.

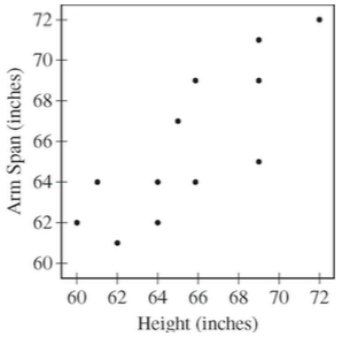

(a) Based on the scatterplot, describe the relationship between arm span and height for the sample of 12 seniors.

Let x represent height, in inches, and let y represent arm span, in inches. Two scatterplots of the same data are shown below. Graph 1 shows the data with the least squares regression line $\hat{y} = 11.74 + 0.8247x$, and graph 2 shows the data with the line $y = x$.

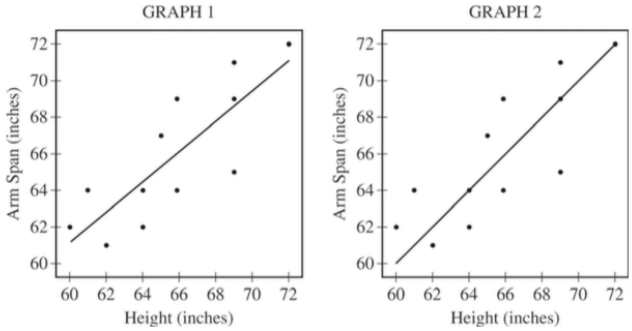

(b) The criteria described in the table below can be used to classify people into one of three body shape categories: square, tall rectangle, or short rectangle.

Square	Tall Rectangle	Short Rectangle
Arm span is equal to height.	Arm span is less than height.	Arm span is greater than height.

(i) For which graph, 1 or 2, is the line helpful in classifying a student's body shape as square, tall rectangle, or short rectangle? Explain.

(ii) Complete the table of classifications for the 12 seniors.

Classification	Square	Tall Rectangle	Short Rectangle
Frequency			

(c) Using the best model for prediction, calculate the predicted arm span for a senior with height 61 inches.

Sampling Activity

An Exercise in Sampling: Rolling Down the River

Name: _____

A farmer has just cleared a new field for corn. It is a unique plot of land in that a river runs along one side. The corn looks good in some areas of the field but not others. The farmer is not sure that harvesting the field is worth the expense. He has decided to harvest 10 plots and use this information to estimate the total yield. Based on this estimate, he will decide whether to harvest the remaining plots.

A. Method Number One: Convenience Sample

The farmer began by choosing 10 plots that would be easy to harvest. They are marked on the grid below:

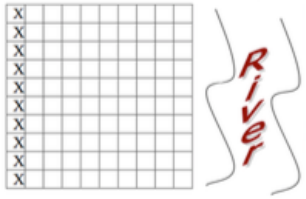

Since then, the farmer has had second thoughts about this selection and has decided to come to you (knowing that you are an AP statistics student, somewhat knowledgeable, but far cheaper than a professional statistician) to determine the approximate yield of the field.

You will still be allowed to pick 10 plots to harvest early. Your job is to determine which of the following methods is the best one to use – and to decide if this is an improvement over the farmer's original plan.

B. Method Number Two: Simple Random Sample

Use your calculator or a random number table to choose 10 plots to harvest. Mark them on the grid below, and describe your method of selection.

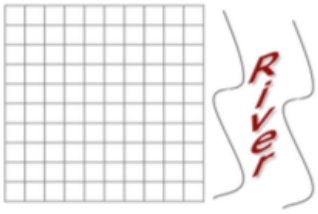

C. Method Number Three: Stratified Sample

Consider the field as grouped in vertical columns (called strata). Using your calculator or a random number table, randomly choose one plot from each vertical column and mark these plots on the grid.

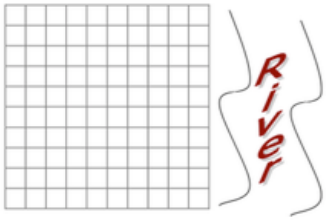

D. Method Number Four: Stratified Sample

Consider the field as grouped in horizontal rows (also called strata). Using your calculator or a random number table, randomly choose one plot from each horizontal row and mark these plots on the grid.

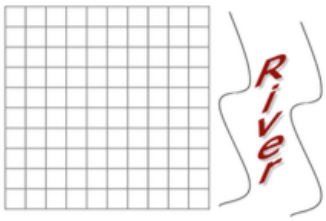

OK, the crop is ready. Below is a grid with the yield per plot. Estimate the average yield per plot based on each of the four sampling techniques.

6	17	20	38	47	55	69	76	82	97
7	14	23	34	43	56	63	75	81	92
2	14	28	30	50	50	62	80	85	96
9	15	27	34	43	51	65	72	88	91
4	15	28	32	44	50	64	76	82	97
5	16	27	31	48	59	69	72	86	99
5	18	28	34	50	60	62	75	90	90
8	15	20	38	40	54	62	77	88	93
7	17	29	39	44	53	61	77	80	90
7	19	22	33	49	53	67	76	86	97

Sampling Method	Mean yield per plot	Estimate of total yield
Convenience Sample (farmer's)		
Simple Random Sample		
Vertical Strata		
Horizontal Strata		

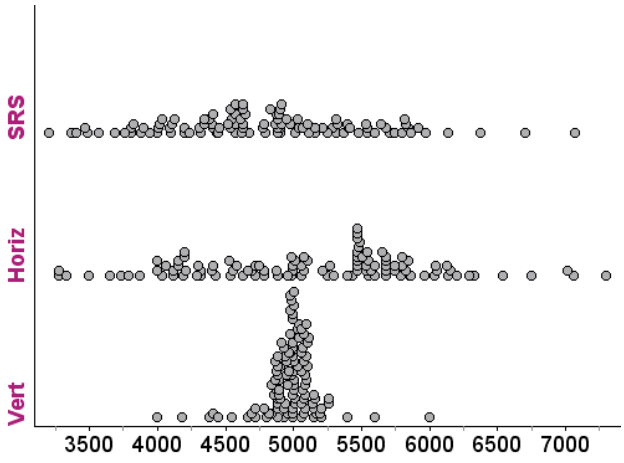

2013 #2

2. An administrator at a large university wants to conduct a survey to estimate the proportion of students who are satisfied with the appearance of the university buildings and grounds. The administrator is considering three methods of obtaining a sample of 500 students from the 70,000 students at the university.

(a) Because of financial constraints, the first method the administrator is considering consists of taking a convenience sample to keep the expenses low. A very large number of students will attend the first football game of the season, and the first 500 students who enter the football stadium could be used as a sample. Why might such a sampling method be biased in producing an estimate of the proportion of students who are satisfied with the appearance of the buildings and grounds?

(b) Because of the large number of students at the university, the second method the administrator is considering consists of using a computer with a random number generator to select a simple random sample of 500 students from a list of 70,000 student names. Describe how to implement such a method.

(c) Because stratification can often provide a more precise estimate than a simple random sample, the third method the administrator is considering consists of selecting a stratified random sample of 500 students. The university has two campuses with male and female students at each campus. Under what circumstance(s) would stratification by campus provide a more precise estimate of the proportion of students who are satisfied with the appearance of the university buildings and grounds than stratification by gender?

Student Sample:

(a) Because of financial constraints, the first method the administrator is considering consists of taking a convenience sample to keep the expenses low. A very large number of students will attend the first football game of the season, and the first 500 students who enter the football stadium could be used as a sample. Why might such a sampling method be biased in producing an estimate of the proportion of students who are satisfied with the appearance of the buildings and grounds?

This could be biased because the students who are attending the football game are therefore → obviously very proud of their school and support their team and school. These students most likely will not have a problem with their school's appearance because they're so proud of the school itself and what it stands for.

(b) Because of the large number of students at the university, the second method the administrator is considering consists of using a computer with a random number generator to select a simple random sample of 500 students from a list of 70,000 student names. Describe how to implement such a method.

First each student should be assigned a a number that falls between 00000 and 69,999. Now all 70,000 student have a number assigned to them. Then use a comuter to generate a random digits table and pick out the first 500 numbers that fall between 00000 and 69,999. When you have your 500 numbers, correspond the correct number to the correct student and you have a sample of 500 random students.

(c) Because stratification can often provide a more precise estimate than a simple random sample, the third method the administrator is considering consists of selecting a stratified random sample of 500 students. The university has two campuses with male and female students at each campus. Under what circumstance(s) would stratification by campus provide a more precise estimate of the proportion of students who are satisfied with the appearance of the university buildings and grounds than stratification by gender?

Stratification by campus would give a more precise estimate than gender because if one campus is nicer than another, that proportion will obviously be more satisfied with the appearance of the university, compared to the less appealing campus.

a) P. No mention of why/how survey is biased.
b) P. No mention of "without replacement"
c) P. No comparison to gender.
PPP =2 (Developing)

2016 AP Statistics Exam

3. Alzheimer's disease results in a loss of cognitive ability beyond what is expected with typical aging. A local newspaper published an article with the following headline.

> Study Finds Strong Association Between Smoking and Alzheimer's

The article reported that a study tracked the medical histories of 21,123 men and women for 23 years. The article stated that, for those who smoked at least two packs of cigarettes a day, the risk of developing Alzheimer's disease was 2.57 times the risk for those who did not smoke.

(a) Identify the explanatory and response variables in the study.

Explanatory variable:

Response variable:

(b) Is the study described in the article an observational study or an experiment? Explain.

(b) Exercise status (regular weekly exercise versus no regular weekly exercise) was mentioned in the article as a possible confounding variable. Explain how exercise status could be a confounding variable in the study.

3. Alzheimer's disease results in a loss of cognitive ability beyond what is expected with typical aging. A local newspaper published an article with the following headline.

> Study Finds Strong Association Between Smoking and Alzheimer's

The article reported that a study tracked the medical histories of 21,123 men and women for 23 years. The article stated that, for those who smoked at least two packs of cigarettes a day, the risk of developing Alzheimer's disease was 2.57 times the risk for those who did not smoke.

(a) Identify the explanatory and response variables in the study.

Explanatory variable:
# of packs of cigarettes smoked per day

Response variable:
the risk of developing alzheimers

(b) Is the study described in the article an observational study or an experiment? Explain.

This study is an observational study because treatments were not given to the men and women in the study. Instead the effects of smoking on the risk of developing Alzheimer's disease were observed.

(c) Exercise status (regular weekly exercise versus no regular weekly exercise) was mentioned in the article as a possible confounding variable. Explain how exercise status could be a confounding variable in the study.

Usually exercise leads to better overall health. Those who exercise on a regular basis may be more healthy than those who do not. This could possibly effect a person's risk of developing Alzheimer's disease.

a) P-"risk" is not a variable
b) P-no context @ levels of smoking
c) P-no connecting exercise to Alzheimer's

Paper Helicopter Experiment

Name:_____

1. Make a paper helicopter according to the directions and the photo below. You will need scissors, two staples, and a template from which to cut out the helicopter.

2. Using three different heights (low, medium and high—you choose the distances) and two different helicopters (short rotor and long rotor), launch helicopters and measure how far they land from a target point on the floor. Conduct four launches at each position for each type of helicopter. This will give you 24 total launches. Collect the data in the table below, and answer the questions on the next page.

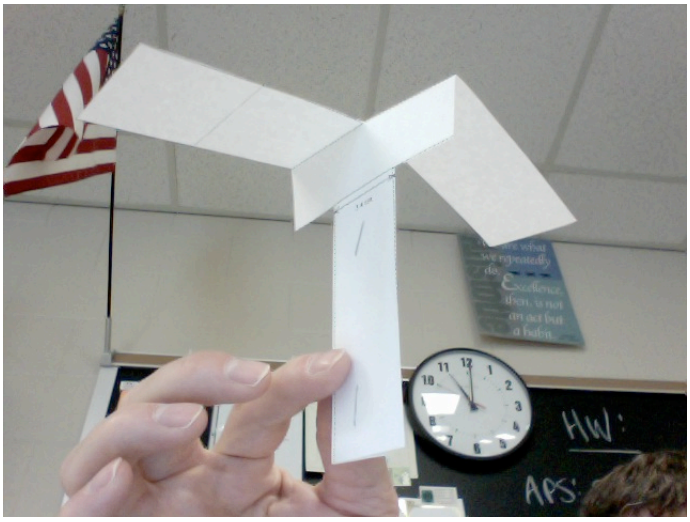

(Long rotor helicopter pictured above.)

Type	Height	Distance from target				Mean	SD
Long rotor	(Low)						
Short rotor	(Low)						
Long rotor	(Medium)						
Short rotor	(Medium)						
Long rotor	(High)						
Short rotor	(High)						

1. What sources of variation are in this experiment?

2. How could we have accounted for these sources of variation?

3. What are the factors in the experiment?

4. What are the levels of each factor?

5. How many total treatments did you have? _____

6. Name at least three possible sources of bias in this experiment (be sure to remember the definition of bias).

F.R.A.P.P.Y!

(See apstatsmonkey.com for more FRAPPY's)

As dogs age, diminished joint and hip health may lead to joint pain and thus reduce a dog's activity level. Such a reduction in activity can lead to other health concerns such as weight gain and lethargy due to lack of exercise. A study is to be conducted to see which of two dietary supplements, glucosamine or chondroitin, is more effective in promoting joint and hip health and reducing the onset of canine osteoarthritis. Researchers will randomly select a total of 300 dogs from ten different large veterinary practices around the country. All of the dogs are more than 6 years old, and their owners have given consent to participate in the study. Changes in joint and hip health will be evaluated after 6 months of treatment.

(a) What would be an advantage to adding a control group in the design of this study?

E P I

(b) Assuming a control group is added to the other two groups in the study, explain how you would assign the 300 dogs to these three groups for a completely randomized design.

E P I

(c) Rather than using a completely randomized design, one group of researchers proposes blocking on clinics, and another group of researchers proposes blocking on breed of dog. How would you decide which <u>one</u> of these two variables to use as a blocking variable?

E P I

TOTAL: _____ / 4

Student Responses for 2007 #2: Dogs' Hip Health

(a) What would be an advantage to adding a control group in the design of this study?

An advantage to adding a control group to this design would be that it gives the experiment something to compare its results to, to see how much of a difference the treatments make.

(b) Assuming a control group is added to the other two groups in the study, explain how you would assign the 300 dogs to these three groups for a completely randomized design.

For every dog that is chosen roll a die. IF the die is a 1 or 2 give the dog the glucosamine. IT the die is a 3 or 4 give the dog the chondroitin. IF the die is a 5 or 6 put the dog in the control group. This will completely randomize the design.

(c) Rather than using a completely randomized design, one group of researchers proposes blocking on clinics, and another group of researchers proposes blocking on breed of dog. How would you decide which one of these two variables to use as a blocking variable?

I would decide to use the blocking on breed of dog. The clinic the dog is in should not affect the medicine the dog is given. However, different breeds of dogs might respond to the medicines differently. Therefore, the blocking on breed of dog should be used.

(a) What would be an advantage to adding a control group in the design of this study?

The advantage to adding a control group in the design of this study would be to have something to compare the results to. This helps to reduce the effects of confounding variables. For example the weather which can affect joint pain.

(b) Assuming a control group is added to the other two groups in the study, explain how you would assign the 300 dogs to these three groups for a completely randomized design.

To obtain a compleatly randomized design I would number each dog 1 to 300 and then using a random number generator I would select 100 numbers ignoring repeats the 100 dogs corresponding to those 100 numbers will be placed in the first treatment group and will recieve glucosamine. I will repeat this process selecting 100 new numbers, these 100 dogs will be placed in the second treatment group and will recieve chondroitin and the remaining 100 dogs will be the control group and will recieve a placeb.

(c) Rather than using a completely randomized design, one group of researchers proposes blocking on clinics, and another group of researchers proposes blocking on breed of dog. How would you decide which one of these two variables to use as a blocking variable?

Whichever variable has more variation should be used as a block. I think breed of dog will cause more variation in the experiment because different kinds of dogs can respond differently to the treatment, but which clinic the dogs came from probably will have less effect on the experiment.

(a) What would be an advantage to adding a control group in the design of this study?

It would be an advantage to add a control group to this study because then after 6 months, you have a group to compare with the treated dogs in the study, to see if the treatments really had an impact in promoting joint and hip health and reducing the onset of canine osteoarthritis.

(b) Assuming a control group is added to the other two groups in the study, explain how you would assign the 300 dogs to these three groups for a completely randomized design.

For a completely randomized design, I would assign each of the 300 dogs a number, 1-300, and then put all the numbers into a hat. Then draw 100 numbers out of the hat and assign them to group 1, the control group. Then pull out 100 more numbers and assign them to group 2, the glucosamine treatment. Then with the 100 left over dogs assign them to group 3, for the chondroitin treatment. That way, you will have three groups for a completely randomized design.

(c) Rather than using a completely randomized design, one group of researchers proposes blocking on clinics, and another group of researchers proposes blocking on breed of dog. How would you decide which one of these two variables to use as a blocking variable?

Rather than using a completely randomized design, I would encorporate blocking on the specific breed of dog, because the different treatments could possibly have a different effect on the different types of dogs, therefore I would use the blocking on breed of dog so it would eliminate any variables that could change the actual data. I would rather block on breed of dog than clinic, because breeds of dogs seem to be much more different than different clinics, therefore I would block on breeds of dogs.

Probability:

1. BIG vs. small Game:
    1. Two players, “BIG” and “small,” play a game with a single die.
    2. It does not matter who rolls the single die each time (players could alternate).
    3. If a 5 or 6 is rolled on the die, “BIG” receives that number of points. If 1, 2, 3, or 4 is rolled, “small” gets that number of points.
    4. The first player to 20 points wins the game. Play 10 games and record the results.
    5. Is this game fair? If not, who has the advantage? Explain.

2. Casino Match War Activity:

    Each student has a deck. Shuffle each deck. Students deal--simultaneously--one card at a time onto two piles. What is the probability that there will be an EXACT (suit and value) match by the time they reach the 52nd pair of dealt cards?

3. Sicherman dice:

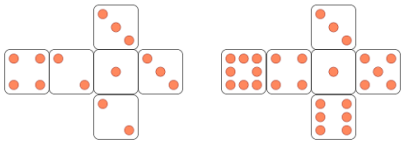

a) Would you rather roll Sicherman dice or standard dice in Monopoly? Explain.

b) What is the probability distribution model for Sicherman dice?

4. Strange Dice Problem: A game is played with 2 strange dice.
    - The six faces of Die A show a 1 and five 3’s.
    - Die B has four 2’s and two 6’s.

    a. Create a probability model for the sum when you roll both dice.

    b. Find the mean and standard deviation.

5. Roll two standard dice. Let X = the roll on which you get “doubles.”
    1. What is E(X)? (“On which roll will you typically get the first doubles?”)

    2. On which roll is it most likely to get doubles? The 3rd? The 6th?

    3. WITP that X = 3? (You get your first

    4. WITPO getting doubles BY the 3rd roll?

    5. P(X > 3 rolls)

    6. Simulate this situation.

AP Statistics: Random Variables

Developing the Rules for +/– Name_____

Let X be a random variable with equally likely outcomes {2, 4, 6, 8}

Let Y be a random variable with equally likely outcomes {1, 3, 5}

E(X) = _____ E(Y) = _____

SD(X) = _____ SD(Y) = _____

VAR(X)= _____ VAR(Y)= _____

Now find the set of outcomes for X + Y: {_____}

Store these values in a list and calculate the mean, SD and variance.

E(X + Y) = _____ SD(X + Y) = _____ VAR(X + Y) = _____

Find the set of outcomes for X – Y: {_____}

Store these values in a list and calculate the mean, SD and variance.

E(X – Y) = _____ SD(X – Y) = _____ VAR(X – Y) = _____

You should be able to verify that for the sum or difference of ***independent*** random variables,

_____ _____.

Dice Problems

Random Variables NAME_____

1. You have two non-standard fair dice, one a 10-sided die numbered with 00, 10, 20…90 and a dodecahedron die numbered 1-12. Let *T* represent the random variable "outcome of a roll on the ten-sided die," and *D* represent the random variable "outcome of a roll on the dodecahedron die."

Calculate the following:

E(D) = _____ SD(D) = _____ VAR(D) = _____

E(T) = _____ SD(T) = _____ VAR(T) = _____

Imagine a game where you roll both dice and calculate the sum. Compute the following:

E(T + D) = _____ SD(T + D) = _____ VAR(T + D) = _____

Now imagine a game where you roll both dice and calculate the difference in the rolls (T – D).

E(T – D) = _____ SD(T – D) = _____ VAR(T – D) = _____

2. Yahtzee is played by rolling five standard dice. What is the mean, standard deviation and variance of the sum of five standard (fair) dice?

Mean = _____

SD = _____

VAR = _____

3. Farkle is played with six standard die. Find the mean, standard deviation and variance of the sum of six die.

Mean = _____

SD = _____

VAR = _____

The 12 Days of Christmas

On the first day of Christmas, my true love gave to me: A Partridge in a Pear Tree.

If the probability of getting a partridge is 0.58 and the probability of getting a pear tree is 0.76, and these are independent events, find the probability of getting a partridge and a pear tree.

On the second day of Christmas, my true love gave to me: Two Turtle Doves.

If the probability of a female turtle dove is 0.53, find the probability of at least one female turtle dove in the pair.

On the third day of Christmas, my true love gave to me: Three French Hens.

If the probability of a hen truly having French citizenship is 0.81, find the probability of exactly two French hens out of the three.

On the fourth day of Christmas, my true love gave to me: Four Calling Birds.

If the probability of a bird actually calling is 0.63, find the probability of finding the first calling bird on the third attempt.

On the fifth day of Christmas, my true love gave to me: Five Golden Rings.

If the probability of getting a real golden ring is 0.72, find the probability of getting three or fewer golden rings in the five.

On the sixth day of Christmas, my true love gave to me: Six Geese A-laying.

If the probability of getting an authentic laying goose is 0.83, find the probability of getting a laying goose on or before the fourth trial.

On the seventh day of Christmas, my true love gave to me: Seven Swans A-swimming.

If the probability that a swan drowns is 0.23, find the probability of exactly 4 out of the 7 swans drown.

On the eighth day of Christmas, my true love gave to me: Eight Maids A-milking.

If the probability of getting a sour maid A-milking is 0.38, find the expected number of sour maids A-milking.

On the ninth day of Christmas, my true love gave to me: Nine Ladies Dancing.

If the probability of a dancing lady accepting an invitation to dance is 0.18, what is the expected number of ladies you would have to ask, before one accepts.

On the tenth day of Christmas, my true love gave to me: Ten Lords A-leaping.

If the probability of a lame leaping lord is 0.24, find the probability of getting your first lame leaping lord after the sixth attempt.

On the eleventh day of Christmas, my true love gave to me: Eleven Pipers Piping.

If the probability of frozen pipes is 0.63, find the probability of 8 or more frozen pipes.

On the twelfth day of Christmas, my true love gave to me: Twelve Drummers Drumming.

If the probability of a dribbling drummer is 0.48, find the standard deviation of the dribbling drummers drumming.

From Jim Luhring, Cherry Creek HS

2016 AP Statistics Exam

4. A company manufactures model rockets that require igniters to launch. Once an igniter is used to launch a rocket, the igniter cannot be reused. Sometimes an igniter fails to operate correctly, and the rocket does not launch. The company estimates that the overall failure rate, defined as the percent of all igniters that fail to operate correctly, is 15 percent.

   A company engineer develops a new igniter, called the super igniter, with the intent of lowering the failure rate. To test the performance of the super igniters, the engineer uses the following process.

   Step 1: One super igniter is selected at random and used in a rocket.
   Step 2: If the rocket launches, another super igniter is selected at random and used in a rocket.

   Step 2 is repeated until the process stops. The process stops when a super igniter fails to operate correctly or 32 super igniters have successfully launched rockets, whichever comes first. Assume that super igniter failures are independent.

   (a) If the failure rate of the super igniters is 15 percent, what is the probability that the first 30 super igniters selected using the testing process successfully launch rockets?

   (b) Given that the first 30 super igniters successfully launch rockets, what is the probability that the first failure occurs on the thirty-first or the thirty-second super igniter tested if the failure rate of the super igniters is 15 percent?

   (c) Given that the first 30 super igniters successfully launch rockets, is it reasonable to believe that the failure rate of the super igniters is less than 15 percent? Explain.

Avg. score: 0.99
a) 0.0076
b) 0.2775
c) Yes. If the rate is 0.15, there is a very small (0.0076) chance of 30 non-failures in a row.

(2002 #3) There are four runners on the New High School team. The team is planning to participate in a race in which each runner runs a mile. The team time is the sum of the individual times for the four runners. Assume that the individual times for the four runners are all independent of each other. The individual times, in minutes, of the runners in similar races are approximately normally distributed with the following means and standard deviations.

	Mean	Standard Deviation
Runner 1	4.9	0.15
Runner 2	4.7	0.16
Runner 3	4.5	0.14
Runner 4	4.8	0.15

a) Runner 3 thinks he can run a mile in less than 4.2 minutes in the next race. Is that likely to happen? Explain.

b) The distribution of possible team times is approximately normal. Find the mean and standard deviation of this distribution.

c) Suppose the team's best time to date is 18.4 minutes. What is the probability that the team will beat its own best time in the next race?

Inference: The Pièce de Résistance!

the most important or remarkable feature…literally "piece which has staying power"

1. The foundation of inference: s_____ d_____

2. German Tanks Problem:

During WWII, Allied spies were asked to estimate the numbers of tanks the Germans had of various types. At about the same time, the Allies were able to capture a number of German tanks, and it was discovered that part numbers on the tanks had coded information that almost certainly indicated serial numbers from the same factories. The part numbers were decoded, and British mathematicians were given the serial numbers and asked to estimate the number of tanks. The mathematicians came up with estimates quite a bit lower than those given by the spies. Long after the war, it was discovered that the spies had been deceived by the Germans repainting their tanks to increase their apparent numbers. The mathematicians were much closer to getting the number of tanks right.[1]

a. Take a random sample of 7 tank numbers from a population of an unknown number of tank numbers. Assume the tank numbers are consecutive integers. (You can also assume the lowest number is "1.")

b. In your group, create three reasonable statistics from these seven numbers that would give an estimate of the largest number in the population (think about what you have learned about distributions). Discuss and decide which of the three is your "best" estimator of the maximum tank number.

c. Calculate your group's estimates of the maximum tank number using this method.

d. Now let's think: assume one group's estimate was only 5 off from the correct answer. Explain why this does not necessarily mean that they have the best method.

e. What would happen if we took another sample of seven tank numbers and your group estimated the maximum tank number?

f. How could we decide which group's method was the best in the long run? What would be the features of such a method?

g. Let's use Fathom to simulate some of these statistics from your groups.

h. Examine the handout of several statistics and their sampling distributions.

For more on the German Tanks Problem, see AP Central AP Statistics Course web site, under Special Focus, Sampling Distributions: http://apcentral.collegeboard.com/apc/public/repository/Statistics_Sampling_Distributions_SF.pdf (pages 19 and 48)

[1] This problem was first introduced to the world in 1947, shortly after many documents concerning WWII became declassified. The original article was *An Empirical Approach to Economic Intelligence in World War II* by Richard Ruggles and Henry Brodie, published in *the Journal of the American Statistical Association,* Vol. 42, No. 237. (Mar., 1947), pp. 72–91. Much has been published about it since then, and information can readily be found on the Web by searching for "German Tank Problem."

Reese's Pieces Activity: Sampling Distribution of $\hat{p}$

Inspired by Activity 13-1 in *Workshop Statistics: Discovery with Data*
Allan J. Rossman and Beth L. Chance

1. Have you noticed that orange is a prominent color with the Reese's brand of candy? What percent of Reese's Pieces so you suppose are orange? Go ahead, guess: _____

2. I will now take a random sample of 25 Reese's Pieces.

   My sample proportion is _____. We have a special symbol for this: $\hat{p}$ ("p-hat")

   So now we know, right? Well, maybe. How likely is it that MY ONE SAMPLE is a PERFECT representation of the population of ALL Reese's Pieces?

3. We should also examine my method. If a sample is to represent the population, it must be _____. And one of the best ways to guarantee that is with a sample that is _____. How did I do?

4. Another condition that must be true (remember binomial/Bernoulli conditions?):

   I_____ How did I do on this one?

5. One more question: what if my sample contained only three pieces? There's no way I could have detected an accurate estimate of the percent orange. So my sample must also be _____ enough. Was it?

6. OK. I think we're ready to believe that I did in fact have a good representative sample that also fits a binomial/Bernoulli model. But even if my sample is a good one, what do you think will happen if I take another sample of 25 and calculate the percent orange?

7. And better yet, what would happen if we took LOTS of samples of size 25 and calculated the percent orange?

8. So let's find out how LOTS of sample proportions behave. Take as many samples of 25 as you can, and calculate the percent orange.

Sample #1: _____% Sample #2: _____% Sample #3: _____%

Make a dot plot of the ***class*** $\hat{p}$'s below. Use $\hat{p}$'s instead of dots.

REMEMBER: This is a distribution of st_____, NOT d_____.

Therefore, it is a very different distribution called a s_____ distribution.

Describe our class *sampling distribution* below: (remember shape, center, spread?)

Notice that you now have some idea of the *sampling variability:* values of statistics ($\hat{p}$'s) from many different samples will vary somewhat from sample to sample.

According to mathematicians who studied this a bit more deeply, the mean of a sampling distribution of sample proportions (i.e. the distribution of $\hat{p}$'s) is always _____ (where _____ is the TRUE population proportion.)

The standard deviation of this sampling distribution of $\hat{p}$'s is _____.

2015 #2: To increase business, the owner of a restaurant is running a promotion in which a customer's bill can be randomly selected to receive a discount. When a customer's bill is printed, a program in the cash register randomly determines whether the customer will receive a discount on the bill. The program was written to generate a discount with a probability of 0.2, that is, giving 20 percent of the bills a discount in the long run. However, the owner is concerned that the program has a mistake that results in the program not generating the intended long-run proportion of 0.2.

The owner selected a random sample of bills and found that only 15 percent of them received discounts. A confidence interval for p, the proportion of bills that will receive a discount in the long run, is 0.15 ± 0.06. All conditions for inference were met.

(a) Consider the confidence interval 0.15 ± 0.06.

(i) Does the confidence interval provide convincing statistical evidence that the program is not working as intended? Justify your answer.

(ii) Does the confidence interval provide convincing statistical evidence that the program generates the discount with a probability of 0.2 ? Justify your answer.

A second random sample of bills was taken that was four times the size of the original sample. In the second sample 15 percent of the bills received the discount.

(b) Determine the value of the margin of error based on the second sample of bills that would be used to compute an interval for p with the same confidence level as that of the original interval.

(c) Based on the margin of error in part (b) that was obtained from the second sample, what do you conclude about whether the program is working as intended? Justify your answer.

2016 AP Statistics Exam:

5. A polling agency showed the following two statements to a random sample of 1,048 adults in the United States.

> Environment statement: Protection of the environment should be given priority over economic growth.
>
> Economy statement: Economic growth should be given priority over protection of the environment.

The order in which the statements were shown was randomly selected for each person in the sample. After reading the statements, each person was asked to choose the statement that was most consistent with his or her opinion. The results are shown in the table.

	Environment Statement	Economy Statement	No Preference
Percent of sample	58%	37%	5%

(a) Assume the conditions for inference have been met. Construct and interpret a 95 percent confidence interval for the proportion of all adults in the United States who would have chosen the economy statement.

(b) One of the conditions for inference that was met is that the number who chose the economy statement and the number who did not choose the economy statement are both greater than 10. Explain why it is necessary to satisfy that condition.

(c) A suggestion was made to use a two-sample z-interval for a difference between proportions to investigate whether the difference in proportions between adults in the United States who would have chosen the environment statement and adults in the United States who would have chosen the economy statement is statistically significant. Is the two-sample z-interval for a difference between proportions an appropriate procedure to investigate the difference? Justify your answer.

2016 AP Statistics Exam:

2. Product advertisers studied the effects of television ads on children's choices for two new snacks. The advertisers used two 30-second television ads in an experiment. One ad was for a new sugary snack called Choco-Zuties, and the other ad was for a new healthy snack called Apple-Zuties.

   For the experiment, 75 children were randomly assigned to one of three groups, A, B, or C. Each child individually watched a 30-minute television program that was interrupted for 5 minutes of advertising. The advertising was the same for each group with the following exceptions.

   - The advertising for group A included the Choco-Zuties ad but not the Apple-Zuties ad.
   - The advertising for group B included the Apple-Zuties ad but not the Choco-Zuties ad.
   - The advertising for group C included neither the Choco-Zuties ad nor the Apple-Zuties ad.

   After the program, the children were offered a choice between the two snacks. The table below summarizes their choices.

Group	Type of Ad	Number Who Chose Choco-Zuties	Number Who Chose Apple-Zuties
A	Choco-Zuties only	21	4
B	Apple-Zuties only	13	12
C	Neither	22	3

   (a) Do the data provide convincing statistical evidence that there is an association between type of ad and children's choice of snack among all children similar to those who participated in the experiment?

   (b) Write a few sentences describing the effect of each ad on children's choice of snack.

Floyd Bullard's Beads Activity:

1. H_0: The proportion of green beads = 0.50
2. From your population, take a sample of 20 beads.
3. Reject or do not reject H_0
4. Take 15 total samples, recording rejections and non-rejections.

Questions:

What was your Type I Error rate?
What was your Type II Error rate?
What was the power of your test?
How could you have improved the power of your test?

Errors and Power Practice Problem: (from BVD, 3e, p.503, #34)

A potter has a 40% breakage rate during kiln firing, so she buys more expensive clay. She fires 10 pieces and will decide to use the new clay if at most one of them breaks.

a) Suppose the new clay is no better. What is the probability she is convinced to switch?

b) Suppose the new clay can reduce breakage to 20%. WITP that her test will not detect the improvement?

c) How can she improve the power of her test?

Investigative Task

Name_____

A consumer organization was concerned that a refrigerator manufacturer was misleading customers by understating the average energy cost (measured in dollars per year) of a particular refrigerator model. The model was advertised to cost $68 per year. To investigate, researchers selected a random sample of 10 refrigerators of that model. Each refrigerator was then randomly assigned a different consumer. Each refrigerator was used for 5 years, and the total energy cost was used to compute the energy cost per year for that model.

(a) Define the parameter of interest and state the null and alternative hypotheses the consumer organization is interested in testing.

(b) One condition for conducting a one-sample t-test in this situation is that the energy costs for this model should be normally distributed. However, the boxplot and histogram shown below indicate that the distribution of the 10 sample values is skewed to the right.

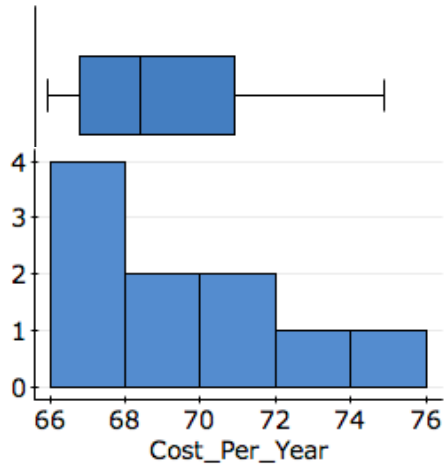

One possible statistic that measures skewness is the ratio:

$$\frac{\text{sample mean}}{\text{sample median}}$$

What values of that statistic sample median (small, large, close to one) might indicate that the population distribution of cost per year is skewed to the right? Explain.

(c) Even though the cost per year values in the sample were skewed to the right, it is still possible that the population distribution of cost per year values is normally distributed and that the skewness was due to sampling variability.

To investigate, 100 samples, each of size 10, were taken from a normal distribution with the same mean and standard deviation as the original sample. For each of those 100 samples, the statistic (sample mean)÷(sample median) was calculated. A dotplot of the 100 simulated statistics is shown below.

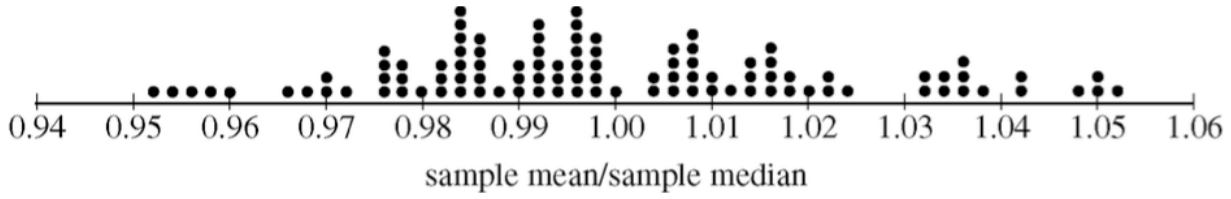

In the original sample, the value of the statistic $\frac{\text{sample mean}}{\text{sample median}}$ was 1.03.

Based on the value of 1.03 and the dotplot above, is it plausible that the original sample of 10 refrigerators came from a normal population, or do the simulated results suggest the original population is really skewed to the right? Explain.

(d) The table below shows summary statistics for cost per year measurements for the original sample of 10 refrigerators.

Minimum	Q1	Median	Q3	Maximum
66	67	68.5	71	75

Choosing only from the summary statistics in the table, define a formula for a different statistic that measures skewness.

What values of that statistic might indicate that the distribution is skewed to the right? Explain.

2015 #6:

6. Corn tortillas are made at a large facility that produces 100,000 tortillas per day on each of its two production lines. The distribution of the diameters of the tortillas produced on production line A is approximately normal with mean 5.9 inches, and the distribution of the diameters of the tortillas produced on production line B is approximately normal with mean 6.1 inches. The figure below shows the distributions of diameters for the two production lines.

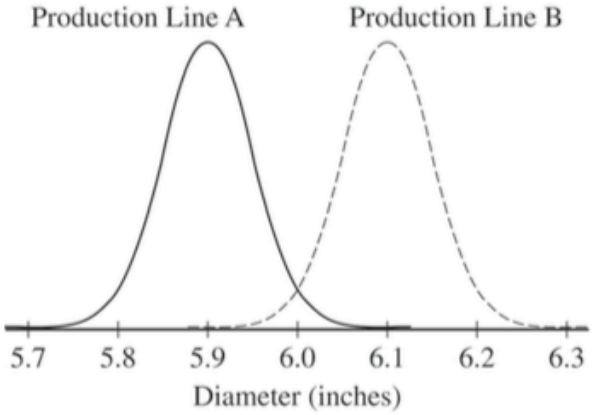

The tortillas produced at the factory are advertised as having a diameter of 6 inches. For the purpose of quality control, a sample of 200 tortillas is selected and the diameters are measured. From the sample of 200 tortillas, the manager of the facility wants to estimate the mean diameter, in inches, of the 200,000 tortillas produced on a given day. Two sampling methods have been proposed.

Method 1: Take a random sample of 200 tortillas from the 200,000 tortillas produced on a given day. Measure the diameter of each selected tortilla.

Method 2: Randomly select one of the two production lines on a given day. Take a random sample of 200 tortillas from the 100,000 tortillas produced by the selected production line. Measure the diameter of each selected tortilla.

(a) Will a sample obtained using Method 2 be representative of the population of all tortillas made that day, with respect to the diameters of the tortillas? Explain why or why not.

(b) The figure below is a histogram of 200 diameters obtained by using one of the two sampling methods described. Considering the shape of the histogram, explain which method, Method 1 or Method 2, was most likely used to obtain a such a sample.

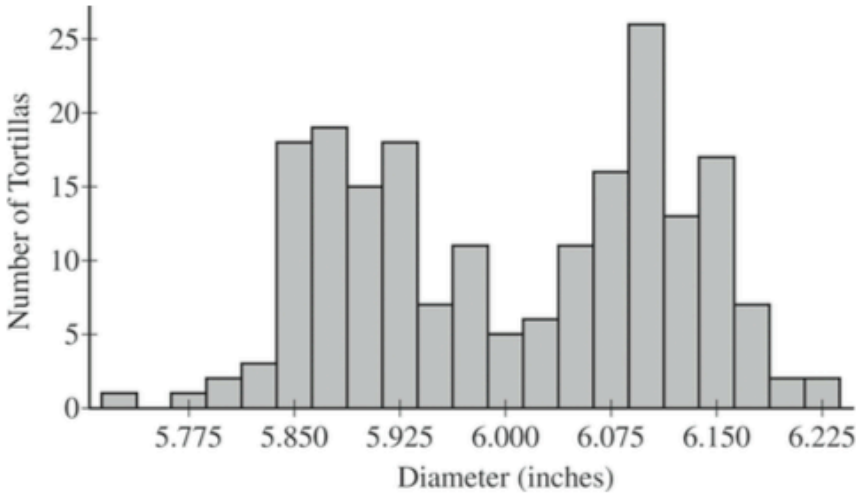

(c) Which of the two sampling methods, Method 1 or Method 2, will result in less variability in the diameters of the 200 tortillas in the sample on a given day? Explain.

Each day, the distribution of the 200,000 tortillas made that day has mean diameter 6 inches with standard deviation 0.11 inch.

(d) For samples of size 200 taken from one day's production, describe the sampling distribution of the sample mean diameter for samples that are obtained using Method 1.

(e) Suppose that one of the two sampling methods will be selected and used every day for one year (365 days). The sample mean of the 200 diameters will be recorded each day. Which of the two methods will result in less variability in the distribution of the 365 sample means? Explain.

(f) A government inspector will visit the facility on June 22 to observe the sampling and to determine if the factory is in compliance with the advertised mean diameter of 6 inches. The manager knows that, with both sampling methods, the sample mean is an unbiased estimator of the population mean. However, the manager is unsure which method is more likely to produce a sample mean that is close to 6 inches on the day of sampling. Based on your previous answers, which of the two sampling methods, Method 1 or Method 2, is more likely to produce a sample mean close to 6 inches? Explain.

Implications from Research

"Learning to think and reason with data and chance is a complex process involving many and varied issues."

©2006, NCTM's Sixty-Eighth Yearbook:
Thinking and Reasoning with Data and Chance
Gail Burrill, editor

Specific Suggestions:

- Transition students from "local view" to "global view" of data.
- Develop the idea of a sample.
- Encourage the need for a statistical term before formal introduction
- Progress from small to large data; then continuous distributions.

General Implications for Teachers:

- Sequence activities carefully
- Focus on the "big ideas"
- Use good, real data
- Encourage speculation and exploration
- Encourage making and testing conjectures
- Use appropriate technology
- Discuss and reflect in class

Research in the Statistics Classroom: Learning from Teaching Experiments
by Dani Ben-Zvi, Joan B. Garfield and Andrew Zieffler
©2006, NCTM's Sixty-Eighth Yearbook:
Thinking and Reasoning with Data and Chance

Engage Students in Activities Involving:

- the investigative process
- "Why?" questions
- meaning of variation
- moving beyond procedures
- new questions that "go beyond"

Chance, Beth, "Components of Statistical Thinking and Implications for Instruction and Assessment,"
Journal of Statistics Education, 10 (no 3), (2002)
http://www.amstat.org/publications/jse/v10n3/chance.html

"…researchers often find that few students properly interpret their results or draw valid conclusions based on the statistical results they obtain. They seem to complete procedures mechanically, without developing statistical thinking." (Ben-Zvi et al)

What is Statistical Thinking and How is It Developed?
by Sharon J. Lane-Getaz
©2006, NCTM's Sixty-Eighth Yearbook:
Thinking and Reasoning with Data and Chance

Fostering Classroom Discourse:

- Expose students to "messy" side of data analysis.
- "do not settle for simplistic solutions to problems."
- Ask lots of questions
- Choose tasks set in interesting, engaging contexts
- Continually assess the quality of students' understandings
  (Using AP Exam rubrics can help evaluate/score written responses.)

Engaging Students in Authentic Data Analysis by Randall E. Groth
©2006, NCTM's Sixty-Eighth Yearbook: Thinking and Reasoning with Data and Chance

How Students Learn Statistics (1999):

- Time spent on developing understanding leads to increased student performance on problem solving tests.
- Use of small groups leads to better group productivity, improved attitudes, and sometimes, increased achievement.
- Having students read through worked-out examples may be more effective than having them work through many of the conventional exercises assigned as homework.
- Students learn more from working on open-ended problems than from goal-specific problems where there is one right answer.

-Joan Garfield (see noblestatman.com for full article)

Garfield and Ben-Zvi's Eight Learning Principles:[2]

1. Students learn by constructing knowledge.
2. Students learn by active involvement in learning activities.
3. Students learn to do well only what they practice doing.
4. It is easy to underestimate the difficulty students have in understanding basic concepts of probability and statistics.
5. It is easy to overestimate how well students understand basic concepts.
6. Learning is enhanced by having students become aware of and confront their errors in reasoning.
7. Technological tools should be used to help students visualize and explore data, not just to follow algorithms to pre-determined ends.
8. Students learn better if they receive consistent and helpful feedback on their performance.

Marzano's Keys to Student Engagement--Students must answer in the affirmative:

How do I feel? Am I interested? [Getting their attention.] Is it important? Can I do it? [Getting them engaged.]

[2] Garfield, Joan, and Dani Ben-Zvi. "How Students Learn Statistics Revisited: A Current Review of Research on Teaching and Learning Statistics." *International Statistical Review* 75.3 (2007): 372-96

Make It Stick: Best Practices for Teachers

…students must be helped to understand such fundamental ideas as these:

- Some kinds of difficulties during learning help to make the learning stronger and better remembered.
- When learning is easy, it is often superficial and soon forgotten.
- Not all of our intellectual abilities are hardwired. In fact, when learning is effortful, it changes the brain, making new connections and increasing intellectual ability.
- You learn better when you wrestle with new problems before being shown the solution, rather than the other way around.
- To achieve excellence in any sphere, you must strive to surpass your current level of ability.
- Striving, by its nature, often results in setbacks, and setbacks are often what provide the essential information needed to adjust strategies to achieve mastery.

Brown, Peter C. (2014-04-14). Make It Stick (pp. 225-226). Harvard University Press. Kindle Edition.

General:

1. Teach students about learning.
2. Use frequent quizzing, and make them count toward students' grade, even if it is low-stakes.
3. Create study tools that incorporate retrieval practice, generation, and elaboration.
4. Make practice exercises count toward the course grade (studies show students learn better when practice exercises carry consequences).
5. Design quizzes and exercises that spiral through previously-learned material.
6. Space, interleave, and vary topics and problems.
7. Be transparent. Show students the ways you incorporate desirable difficulties into your lessons. Tell them the perils of "illusions of knowing."

Solutions:

Solutions for "The 12 Days of Christmas."

1. P(partridge and pear tree) = P(partridge)*P(pear tree)because of independence, (.58)(.76)= .4408
2. P(1 or 2 females) = 1 - binompdf(2,.53,0) = .7791
3. P(2 out of 3 French hens) = binompdf(3,.81,2) = .3740
4. P(first calling bird on trial 3) = geometpdf(.63,3) = .0862
5. P(3 or fewer golden rings) = binomcdf(5,.72,3) = .4303
6. P(laying goose by 4th) = geometcdf(.83,4) = .9992
7. P(4 out of 7 dead swans) = binompdf(7,.23,4) = .0447
8. E(x)- binomial = np = 8*.38 = 3.04
9. E(x)- geometric = 1/p = 1/.18 = 5.55555
10. P(1st lame lord after 6th) = Sum[geometpdf(.24,{7,8,9,10})= .1284
11. P(8 out of 11 frozen pipes) = 1-binomcdf(11,.63,7) = .3714
12. SD of binomial = SQRT (npq) = SQRT (12*.48*.52) = 1.7307

AP Statistics Exam General Info:

40MC, 6FR, 50%, 90min, #6, 20-25min, ¼, 1/8, big topics, new
Scout/scan/rank, easiest, stretch, calmly, think, ANY

FRAPPY! Student Samples Commentary

Comparing Distributions: 2015 #1

#1: Two Corporations (Comparing Boxplots, Interpreting in Context)

Comments from John Bennett: One issue that seemed to be repeated over and over was a "shoot oneself in the foot" syndrome. On both parts, students would often list several different statistics in sort of a "shotgun" approach, and if any of those sited statistics was wrong, we were instructed to drop that component one letter (from an E to a P, or P to an I). Another problem was correctly comparing the boxplots in A. Students often just listed the quartiles, but didn't compare them (using larger than, less than, etc.) when appropriate. This would NOT get them credit for the individual components for part A. Many students used the mean to comment on center but mean can't be determined from a box plot. Also, context needs to be about the DATA. Students often mentioned the Corporations, but never applied context to the data (salaries, dollars, etc) and lost credit.

C: EP(√√√X)=3 No explanation of relevance (“earn more money…”) on Part b-ii.
F: IE=2 The “$” is the only part scored correct on Part a (context). There are no explicit ***comparisons*** of center, spread or outliers. (Complete shape information cannot be determined from a boxplot, so any mention of shape is disregarded.)

Video comments: http://www.youtube.com/watch?v=YQ9kPdAonpU

The mean score on this problem was 2.24 (out of 4)

Floyd Bullard's Beads Activity (for showing power and errors)

"I hand each student (or group) an opaque bag of 200 colored beads, of which the proportion of blue beads is between 0.00 and 1.00, and they range by 0.05. (So there are 21 bags.) A card inside each bag tells what fraction of the beads are blue, but they aren't to take the card out and look at it, and they are not to look in the bags.

Then each student/group is to draw 20 beads at random and, using those beads, test the null hypothesis that their bag contains 50% blues against the null hypothesis that it doesn't. They don't have to do a write-up or anything, or check conditions, etc.; I actually want this part go to go pretty fast, so they're encouraged to just use their calculator to perform the test. I ask them to write down whether their test rejected the null hypothesis at the alpha=0.05 level.

Then they return their beads to their bags, and take another random sample, and do the same: write down whether the test rejects H_O or not. And they do this over and over again until they've collected 15 samples.

Analysis: several possibilities. Write down on the board how many times each student/group rejected H_O and how many times they did not reject H_O. Reveal the actual % blue in each bag , and discuss what type of errors were made as well as the frequency of those errors. Then discuss power. Power should be higher for bags where the null proportion was farther away from the actual proportion."

Age Guessing:

Barack Obama	8/4/61
Harrison Ford	7/13/42
Bill Gates	10/28/55
Johnny Depp	6/9/63
Bono	5/10/60
Alex Trebek	7/22/40
Oprah Winfrey	1/29/54
Miley Cyrus	11/23/92
Peyton Manning	3/24/76
J. K. Rowling	7/31/65
Mick Jagger	7/26/43
Mark Zuckerberg	5/14/84

Pottery Problem Answers:

a) $P(X=0) + P(X=1) = 0.0464$

b) $P(X \geq 2) = 0.6242$

c) Increase number of pieces fired or lower her standard of "proof" (saying 0, 1, or 2 broken will convince her the new clay is better).

Tommy John and P-Values

(from the book, *What is a P-value, Anyway?*)

In a scientific paper describing a clinical trial comparing a new pain drug with a placebo, the authors wrote something like this: "Although there was no difference in baseline age between the groups ($p = 0.458$), controls were significantly more likely to be male ($p = 0.000$)."

This statement is worse than Tommy John's worst day because there are actually four errors in this sentence (or maybe even 4½). See if you can find them.

(answers on succeeding page)

Solution to Investigative Task:

≈ 2009 #6: (data set was changed to refrigerators instead of mpg of cars)

Part (a):

The parameter of interest is $\mu =$ population mean cost per year (in dollars) of a particular model of refrigerator.

H_0: $\mu = \$68$

H_A: $\mu < \$68$

Part (b):

If the distribution is right-skewed, one would expect the mean to be greater than the median. Therefore the ratio $\frac{\text{sample mean}}{\text{sample median}}$ should be large (greater than 1).

Part (c):

Because we are testing for right-skewness, the estimated p-value will be the proportion of the simulated statistics that are greater than or equal to the observed value of 1.03. The dotplot shows that 14 of the 100 values are more than 1.03. Because this simulated p-value (0.14) is larger than any reasonable significance level, we do not have convincing evidence that the original population is skewed to the right and conclude that it is plausible that the original sample came from a normal population.

Part (d):

One possible statistic is $\frac{\text{maximum–median}}{\text{median–minimum}}$. If the distribution is right-skewed, one would expect the distance from the median to the maximum to be larger than the distance from the median to the minimum; thus the ratio should be greater than one.

Answers to Tommy John problem:

1. Accepting the null hypothesis
2. Giving a p-value for baseline differences between random groups (p-values test hypotheses).
3. Inappropriate levels of precision (what do the 5 and 8 tell us?)
4. Reporting a p-value = 0.

4½: Why were they measuring baseline ages anyway? All patients will be in the trial over the same time period.