

Statistics

Like most people, you probably feel that it is important to "take control of your life." But what does this mean? Partly it means being able to properly evaluate the data and claims that bombard you every day. If you cannot distinguish good from faulty reasoning, then you are vulnerable to manipulation and to decisions that are not in your best interest. Statistics provides tools that you need in order to react intelligently to information you hear or read. In this sense, Statistics is one of the most important things that you can study.

To be more specific, here are some claims that we have heard on several occasions. (We are *not* saying that each one of these claims is true!)

- 4 out of 5 dentists recommend Dentyne.
- Almost 85% of lung cancers in men and 45% in women are tobacco-related.
- Condoms are effective 94% of the time.
- Native Americans are significantly more likely to be hit crossing the streets than are people of other ethnicities.
- People tend to be more persuasive when they look others directly in the eye and speak loudly and quickly.
- Women make 75 cents to every dollar a man makes when they work the same job.
- A surprising new study shows that eating egg whites can increase one's life span.
- People predict that it is very unlikely there will ever be another baseball player with a batting average over 400.
- There is an 80% chance that in a room full of 30 people that at least two people will share the same birthday.
- 79.48% of all statistics are made up on the spot.

All of these claims are statistical in character. We suspect that some of them sound familiar; if not, we bet that you have heard other claims like them. Notice how diverse the examples are. They come from psychology, health, law, sports, business, etc. Indeed, data and data-interpretation show up in discourse from virtually every facet of contemporary life.

Statistics are often presented in an effort to add credibility to an argument or advice. You can see this by paying attention to television advertisements. Many of the numbers thrown about in this way do not represent careful statistical analysis. They can be misleading, and push you into decisions that you might find cause to regret. For these reasons, learning about statistics is a long step towards taking control of your life. (It is not, of course, the only step needed for this purpose.) These chapters will help you learn statistical essentials. It will make you into an intelligent consumer of statistical claims.

You can take the first step right away. To be an intelligent consumer of statistics, your first reflex must be to question the statistics that you encounter. The British Prime Minister Benjamin Disraeli famously said, "There are three kinds of lies -- lies, damned lies, and statistics." This quote reminds us why it is so important to understand statistics. So let us invite you to reform your statistical habits from now on. No longer will you blindly accept numbers or findings. Instead, you will begin to think about the numbers, their sources, and most importantly, the procedures used to generate them.

We have put the emphasis on defending ourselves against fraudulent claims wrapped up as statistics. Just as important as detecting the deceptive use of statistics is the appreciation of the proper use of statistics. You must also learn to recognize statistical evidence that supports a stated conclusion. When a research team is testing a new treatment for a disease, statistics allows them to conclude based on a relatively small trial that there is good evidence their drug is effective. Statistics allowed prosecutors in the 1950's and 60's to demonstrate racial bias existed in jury panels. Statistics are all around you, sometimes used well, sometimes not. We must learn how to distinguish the two cases.

Populations and samples

Before we begin gathering and analyzing data we need to characterize the **population** we are studying. If we want to study the amount of money spent on textbooks by a typical first-year college student, our population might be all first-year students at your college. Or it might be:

- All first-year community college students in the state of Washington.
- All first-year students at public colleges and universities in the state of Washington.
- All first-year students at all colleges and universities in the state of Washington.
- All first-year students at all colleges and universities in the entire United States.
- And so on.

Why is it important to specify the population? We might get different answers to our question as we vary the population we are studying. First-year students at the University of Washington might take slightly more diverse courses than those at your college, and some of these courses may require less popular textbooks that cost more; or, on the other hand, the University Bookstore might have a larger pool of used textbooks, reducing the cost of these books to the students. Whichever the case (and it is likely that some combination of these and other factors are in play), the data we gather from your college will probably not be the same as that from the University of Washington. Particularly when conveying our results to others, we want to be clear about the population we are describing with our data.

Example: A newspaper website contains a poll asking people their opinion on a recent news article. In this case, the population of the survey is all readers of the website.

If we were able to gather data on every member of our population, say the average (we will define "average" more carefully in a subsequent section) amount of money spent on textbooks by each first-year student at your college during the 2009-2010 academic year, the resulting number would be called a **parameter**. We seldom see parameters, however, since surveying an entire population is usually very time-consuming and expensive, unless the population is very small. A survey of an entire population is called a **census**.

You are probably familiar with two common censuses: the official government Census that attempts to count the population of the U.S. every ten years, and voting, which asks the opinion of all eligible voters in a district. The first of these demonstrates one additional problem with a census: the difficulty in finding and getting participation from everyone in a large population, which can bias, or skew, the results.

There are occasionally times when a census is appropriate, usually when the population is fairly small. For example, if the manager of Starbucks wanted to know the average number of hours her employees worked last week, she should be able to pull up payroll records or ask each employee directly.

Since surveying an entire population is often impractical, we usually select a **sample** to study; that is, a smaller subset of the entire population. We will discuss sampling methods in greater detail in a later section. For now, let us assume that samples are chosen in an appropriate manner. If we survey a sample, say 100 first-year students at your college, and find the average amount of money spent by these students on textbooks, the resulting number is called a **statistic**. The **sample size** was 100.

Example: A researcher wanted to know how citizens of Tacoma felt about a voter initiative. To study this, she goes to the Tacoma Mall and randomly selects 500 shoppers and asks them their opinion. What is the sample and population?

The sample is the 500 shoppers questioned. The population is less clear. While the intended population of this survey was Tacoma citizens, the effective population was mall shoppers. There is no reason to assume that the 500 shoppers questioned would be representative of all Tacoma citizens.

Categorizing data

Once we have gathered data, we might wish to classify it. Roughly speaking, data can be classified as categorical data or quantitative data.

Also known as qualitative data, categorical data, not surprisingly, are pieces of information that allow us to classify the objects under investigation into various categories. For example, we might conduct a survey to determine the name of the favorite movie that each person in a math class saw in a movie theater. When we conduct such a survey, the responses would look like: *Finding Nemo*, *The Hulk*, or *Terminator 3: Rise of the Machines*. We might count the number of people who give each answer, but the answers themselves do not have any numerical values: we cannot perform computations with an answer like "*Finding Nemo*."

Quantitative data, on the other hand, are responses that are numerical in nature and with which we can perform meaningful arithmetic calculations. Examples of survey questions that would elicit quantitative answers are: the number of movies you have seen in a movie theater in the past 12 months (0, 1, 2, 3, 4, ...); the running time of the movie you saw most recently (104 minutes, 137 minutes, 104 minutes, ...); the amount of money you paid for a movie ticket the last time you went to a movie theater (\$5.50, \$7.75, \$9, ...).

Sometimes, determining whether or not data is categorical or quantitative can be a bit trickier. Suppose we gather respondents' ZIP codes in a survey to track their geographical location. ZIP codes are numbers, but we can't do any meaningful mathematical calculations with them (it doesn't make sense to say that 98036 is "twice" 49018 — that's like saying that Lynnwood, WA is "twice" Battle Creek, MI, which doesn't make sense at all), so ZIP codes are really categorical data.

Similarly, a survey about the movie you most recently attended might include a question such as "How would you rate the movie you just saw?" with these possible answers:

- 1 - it was awful
- 2 - it was just OK
- 3 - I liked it
- 4 - it was great
- 5 - best movie ever!

Again, there are numbers associated with the responses, but we can't really do any calculations with them: a movie that rates a 4 is not necessarily twice as good as a movie that rates a 2, whatever that means; if two people see the movie and one of them thinks it stinks and the other thinks it's the best ever it doesn't necessarily make sense to say that "on average they liked it."

As we study movie-going habits and preferences, we shouldn't forget to specify the population under consideration. If we survey 3-7 year-olds the runaway favorite might be *Finding Nemo*. 13-17 year-olds might prefer *Terminator 3*. And 33-37 year-olds might prefer...well, *Finding Nemo*.

Sampling methods

As we mentioned in a previous section, the first thing we should do before conducting a survey is to identify the population that we want to study. Suppose we are hired by a politician to determine the amount of support he has among the electorate should he decide to run for another term. What population should we study? Every person in the district? Not every person is eligible to vote, and regardless of how strongly someone likes or dislikes the candidate, they don't have much to do with him being re-elected if they are not able to vote.

What about eligible voters in the district? That might be better, but if someone is eligible to vote but does not register by the deadline, they won't have any say in the election either. What about registered voters? Many people are registered but choose not to vote. What about "likely voters?"

This is the criteria used in much political polling, but it is sometimes difficult to define a "likely voter." Is it someone who voted in the last election? In the last general election? In the last presidential election? Should we consider someone who just turned 18 a "likely voter?" They weren't eligible to vote in the past, so how do we judge the likelihood that they will vote in the next election?

In November 1998 former professional wrestler Jesse "The Body" Ventura was elected governor of Minnesota. Up until right before the election, most polls showed he had little chance of winning. There were several contributing factors to the polls not reflecting the actual intent of the electorate:

- Ventura was running on a third-party ticket and most polling methods are better suited to a two-candidate race.
- Many respondents to polls may have been embarrassed to tell pollsters that they were planning to vote for a professional wrestler.
- The mere fact that the polls showed Ventura had little chance of winning might have prompted some people to vote for him in protest to send a message to the major-party candidates.

But one of the major contributing factors was that Ventura recruited a substantial amount of support from young people, particularly college students, who had never voted before and who registered specifically to vote in the gubernatorial election. The polls did not deem these young people likely voters (since in most cases young people have a lower rate of voter registration and a turnout rate for elections) and so the polling samples were subject to **sampling bias**: they omitted a portion of the electorate that was weighted in favor of the winning candidate.

So even identifying the population can be a difficult job, but once we have identified the population, how do we choose an appropriate sample? Remember, although we would prefer to survey all members of the population, this is usually impractical unless the population is very small, so we choose a sample. There are many ways to sample a population, but there is one goal we need to keep in mind: we would like the sample to be *representative of the population*.

Returning to our hypothetical job as a political pollster, we would not anticipate very accurate results if we drew all of our samples from among the customers at a Starbucks, nor would we expect that a sample drawn entirely from the membership list of the local Elks club would provide a useful picture of district-wide support for our candidate.

One way to ensure that the sample has a reasonable chance of mirroring the population is to employ *randomness*. In a truly **simple random sample**, each member of the population has an equal probability of being chosen. If we could somehow identify all likely voters in the state, put each of their names on a piece of paper, toss the slips into a (very large) hat and draw 1000 slips out of the hat, we would have a simple random sample. In practice, computers are better suited for this sort of endeavor than millions of slips of paper and extremely large headgear.

It is always possible, however, that even a random sample might end up being biased or skewed. If we repeatedly take samples of 1000 people from among the population of likely voters in the state of Washington, some of these samples might tend to have a slightly higher percentage of Democrats (or Republicans) than does the general population; some samples might include more older people and some samples might include more younger people; etc. In most cases, this **sampling variability** is not significant.

To help account for this, pollsters might instead use a **stratified sample**. Suppose in a particular state that previous data indicated that the electorate was comprised of 39% Democrats, 37% Republicans and 24% independents. In a sample of 1000 people, they would then expect to get about 390 Democrats, 370 Republicans and 240 independents. To accomplish this, they could randomly select 390 people from among those voters known to

be Democrats, 370 from those known to be Republicans, and 240 from those with no party affiliation. Stratified sampling can also be used to select a sample with people in desired age groups, a specified mix ratio of males and females, etc.

A variation on this technique is called **quota sampling**, in which pollsters call people at random, but once they have met their quota of 390 Democrats, say, they only gather people who do not identify themselves as a Democrat. You may have had the experience of being called by a telephone pollster who started by asking you your age, income, etc. and then thanked you for your time and hung up before asking any "real" questions. Most likely, they already had contacted enough people in your demographic group and were looking for people who were older or younger, richer or poorer, etc. Quota sampling is usually a bit easier than stratified sampling, but also is not as random as a simple random sample.

Other sampling methods include **systematic sampling**, in which a sample is chosen in a systematic way, such as selecting every 100th name in the phone book. Systematic sampling is not as random as a simple random sample (if your name is Albert Aardvark and your sister Alexis Aardvark is right after you in the phone book, there is no way you could both end up in the sample) but it can yield acceptable samples.

Perhaps the worst types of sampling methods are **convenience samples** in which samples are chosen in a way that is convenient for the pollster, such as standing on a street corner and interviewing the first 100 people who agree to speak with you. One type of convenience sample is a **self-selected sample**, or **voluntary response sample**, in which respondents volunteer to participate. Usually such samples are skewed towards people who have a particularly strong opinion about the subject of the survey or who just have way too much time on their hands and enjoy taking surveys. A survey on a website is an example of a self-selected sample.

How to mess things up before you start

There are number of ways that a study can be ruined before you even start collecting data. The first we have already explored – **sampling** or **selection bias**, which is when the sample is not representative of the population. One example of this is **voluntary response bias**, which is bias introduced by only collecting data from those who volunteer to participate. This is not the only potential source of bias.

Consider a recent study which found that chewing gum may raise math grades in teenagers¹. This study was conducted by the Wrigley Science Institute, a branch of the Wrigley chewing gum company. This is an example of a **self-interest study**; one in which the researches have a vested interest in the outcome of the study. While this does not necessarily ensure that the study was biased, it certainly suggests that we should subject the study to extra scrutiny.

Another source of bias is **response bias**, in which something causes the responder to give inaccurate responses. This can be from sources as innocent as bad memory when answering a question like “when was the last time you visited your doctor”, or as intentional as

¹ Reuters. http://news.yahoo.com/s/nm/20090423/od_uk_nm/oukoe_uk_gum_learning. Retrieved 4/27/09

pressuring by the pollster. Consider, for example, how many voting initiative petitions people sign without even reading them.

Another source is **perceived lack of anonymity**, which can influence people. For example, in a poll about race relations, the respondent might not want to be perceived as racist even if they are, and give an untruthful answer. In other cases, answering truthfully might have consequences. For example, if an employer puts out a survey asking their employees if they have a drug abuse problem and need treatment help, the responses might not be accurate if the employees do not feel their responses are anonymous.

Another example of response bias is **loaded** or **leading questions**. For example, consider how people would be likely to respond to the question “Do you support the invasion and occupation of Iraq?” compared to “Do you support the liberation of Iraq from dictator rule?” These are leading questions – questions whose wording lead the respondent towards an answer - and can occur unintentionally. Also a concern is **question order**, where the order of questions changes the results. A psychology researcher provides an example²:

“My favorite finding is this: we did a study where we asked students, 'How satisfied are you with your life? How often do you have a date?' The two answers were not statistically related - you would conclude that there is no relationship between dating frequency and life satisfaction. But when we reversed the order and asked, 'How often do you have a date? How satisfied are you with your life?' the statistical relationship was a strong one. You would now conclude that there is nothing as important in a student's life as dating frequency.”

In addition to response bias, there is also **non-response bias**, introduced by people refusing to participate in a study or dropping out of an experiment. When people refuse to participate, we can no longer be so certain that our sample is representative of the population. Suppose, for example, we used a telephone poll to ask the question “Do you often have time to relax and read a book?”, and 50% of the people we called refused to answer the survey. It is unlikely that our results will be representative of the entire population.

Experiments

So far, we have primarily discussed **observational studies** – studies in which conclusions would be drawn from observations of a sample or the population. In some cases these observations might be unsolicited, such as studying the percentage of cars that turn right at a red light even when there is a “no turn on red” sign. In other cases the observations are solicited, like in a survey or a poll.

In contrast, it is common to use **experiments** when exploring how subjects react to an outside influence. In an experiment, some kind of **treatment** is applied to the subjects and the results are measured and recorded. Here are some examples of experiments:

² Swartz, Norbert. <http://www.umich.edu/~newsinfo/MT/01/Fal01/mt6f01.html>. Retrieved 3/31/2009

- 1) A pharmaceutical company tests a new medicine for treating Alzheimer's disease by administering the drug to 50 elderly patients with recent diagnoses. The treatment here is the new drug.
- 2) A gym tests out a new weight loss program by enlisting 30 volunteers to try out the program. The treatment here is the new program.
- 3) You test a new kitchen cleaner by buying a bottle and cleaning your kitchen. The new cleaner is the treatment.
- 4) A psychology researcher explores the effect of music on temperament by measuring people's temperament while listening to different types of music. The music is the treatment.

When conducting experiments, it is essential to isolate the treatment being tested. For example, suppose a middle school (junior high) finds that their students are not scoring well on the state's standardized math test. They decide to run an experiment to see if an alternate curriculum would improve scores. To run the test, they hire a math specialist to come in and teach a class using the new curriculum. To their delight, they see an improvement in test scores.

The difficulty with this scenario is that it is not clear whether the curriculum is responsible for the improvement, or whether the improvement is due to a math specialist teaching the class. This is called **confounding** – when it is not clear which factor or factors caused the observed effect. Confounding is the downfall of many experiments, though sometimes it is hidden.

For example, a drug company study about a weight loss pill might report that people lost an average of 8 pounds while using their new drug. However, in the fine print you find a statement saying that participants were encouraged to also diet and exercise. It is not clear in this case whether the weight loss is due to the pill, to diet and exercise, or a combination of both. In this case confounding has occurred.

There are a number of measures that can be introduced to help reduce the likelihood of confounding. The primary measure is to use a **control group**. When using a control group, the participants are divided into two or more groups, typically a control group and a treatment group. The treatment group receives the treatment being tested; the control group does not receive the treatment. Ideally, the groups are otherwise as similar as possible, isolating the treatment as the only potential source of difference between the groups. For this reason, the method of dividing groups is important. Some researchers attempt to ensure that the groups have similar characteristics (same number of females, same number of people over 50, etc.), but it is nearly impossible to control for every characteristic. Because of this, random assignment is very commonly used.

For example, to determine if a two day prep course would help high school students improve their scores on the SAT test, a group of students was randomly divided into two subgroups. The first group, the treatment group, was given a two day prep course. The second group, the control group, was not given the prep course. Afterwards, both groups were given the SAT. In another example a company testing a new plant food could grow two crops of plants, the treatment group receiving the new plant food and the control group not. The crop yield would then be compared.

Sometimes not giving the control group anything does not completely control for confounding variables. For example, suppose a medicine study is testing a new headache pill by giving the treatment group the pill and the control group nothing. If the treatment group showed improvement, we would not know whether it was due to the medicine in the pill, or a response to have taken any pill. This is called a **placebo effect** – when the effectiveness of a treatment is influenced by the patient’s perception of how effective they think the treatment will be. For example, a study found that when doing painful dental tooth extractions, patients told they were receiving a strong painkiller while actually receiving a saltwater injection found as much pain relief as patients receiving a dose of morphine.³

To control for the placebo effect, a **placebo**, or dummy treatment, is often given to the control group. This way, both groups are truly identical except for the specific treatment given. For the study of a medicine, a sugar pill could be used as a placebo. In a study on the effect of alcohol on memory, a non-alcoholic beer might be given to the control group as a placebo. In a study of a frozen meal diet plan, the treatment group would receive the diet food, and the control could be given standard frozen meals stripped of their original packaging. An experiment that gives the control group a placebo is called a **placebo controlled experiment**.

In some cases, it is more appropriate to compare to a conventional treatment than a placebo. For example, in a cancer research study, it would not be ethical to deny any treatment to the control group or to give a placebo treatment. In this case, the currently acceptable medicine would be given to the second group, called a **comparison group** in this case. In our SAT test example, the non-treatment group would most likely be encouraged to study on their own, rather than be asked to not study at all, to provide a meaningful comparison.

When using a placebo, it would defeat the purpose if the participant knew they were receiving the placebo. A **blind study** is one in which the participant does not know whether or not they are receiving the treatment or a placebo. In a study about anti-depression medicine, you would not want the psychological evaluator to know whether the patient is in the treatment or control group either, as it might influence their evaluation. A **double-blind study** is one in which those interacting with the participants don’t know who is in the treatment group and who is in the control group.

It should be noted that not every experiment needs a control group. For example, if a researcher is testing whether a new fabric can withstand fire, she simply needs to torch multiple samples of the fabric – there is no need for a control group.

³ Levine JD, Gordon NC, Smith R, Fields HL. (1981) Analgesic responses to morphine and placebo in individuals with postoperative pain. *Pain*. 10:379-89.

Exercises

Skills

1. A political scientist surveys 28 of the current 106 representatives in a state's congress. Of them, 14 said they were supporting a new education bill, 12 said there were not supporting the bill, and 2 were undecided.
 - a. What is the population of this survey?
 - b. What is the size of the population?
 - c. What is the size of the sample?
 - d. Give the sample statistic for the proportion of voters surveyed who said they were supporting the education bill.
 - e. Based on this sample, we might expect how many of the representatives to support the education bill?
2. The city of Raleigh has 9500 registered voters. There are two candidates for city council in an upcoming election: Brown and Feliz. The day before the election, a telephone poll of 350 randomly selected registered voters was conducted. 112 said they'd vote for Brown, 238 said they'd vote for Feliz, and 31 were undecided.
 - a. What is the population of this survey?
 - b. What is the size of the population?
 - c. What is the size of the sample?
 - d. Give the sample statistic for the proportion of voters surveyed who said they'd vote for Brown.
 - e. Based on this sample, we might expect how many of the 9500 voters to vote for Brown?
3. Identify the most relevant source of bias in this situation: A survey asks the following: Should the mall prohibit loud and annoying rock music in clothing stores catering to teenagers?
4. Identify the most relevant source of bias in this situation: To determine opinions on voter support for a downtown renovation project, a surveyor randomly questions people working in downtown businesses.
5. Identify the most relevant source of bias in this situation: A survey asks people to report their actual income and the income they reported on their IRS tax form.
6. Identify the most relevant source of bias in this situation: A survey randomly calls people from the phone book and asks them to answer a long series of questions.
7. Identify the most relevant source of bias in this situation: A survey asks the following: Should the death penalty be permitted if innocent people might die?
8. Identify the most relevant source of bias in this situation: A study seeks to investigate whether a new pain medication is safe to market to the public. They test by randomly selecting 300 men from a set of volunteers.

9. In a study, you ask the subjects their age in years. Is this data qualitative or quantitative?
10. In a study, you ask the subjects their gender. Is this data qualitative or quantitative?
11. Does this describe an observational study or an experiment: The temperature on randomly selected days throughout the year was measured.
12. Does this describe an observational study or an experiment? A group of students are told to listen to music while taking a test and their results are compared to a group not listening to music.
13. In a study, the sample is chosen by separating all cars by size, and selecting 10 of each size grouping. What is the sampling method?
14. In a study, the sample is chosen by writing everyone's name on a playing card, shuffling the deck, then choosing the top 20 cards. What is the sampling method?
15. A team of researchers is testing the effectiveness of a new HPV vaccine. They randomly divide the subjects into two groups. Group 1 receives new HPV vaccine, and Group 2 receives the existing HPV vaccine. The patients in the study do not know which group they are in.
 - a. Which is the treatment group?
 - b. Which is the control group (if there is one)?
 - c. Is this study blind, double-blind, or neither?
 - d. Is this best described as an experiment, a controlled experiment, or a placebo controlled experiment?
16. For the clinical trials of a weight loss drug containing *Garcinia cambogia* the subjects were randomly divided into two groups. The first received an inert pill along with an exercise and diet plan, while the second received the test medicine along with the same exercise and diet plan. The patients do not know which group they are in, nor do the fitness and nutrition advisors.
 - a. Which is the treatment group?
 - b. Which is the control group (if there is one)?
 - c. Is this study blind, double-blind, or neither?
 - d. Is this best described as an experiment, a controlled experiment, or a placebo controlled experiment?

Concepts

17. A teacher wishes to know whether the males in his/her class have more conservative attitudes than the females. A questionnaire is distributed assessing attitudes.
 - a. Is this a sampling or a census?
 - b. Is this an observational study or an experiment?
 - c. Are there any possible sources of bias in this study?

18. A study is conducted to determine whether people learn better with spaced or massed practice. Subjects volunteer from an introductory psychology class. At the beginning of the semester 12 subjects volunteer and are assigned to the massed-practice group. At the end of the semester 12 subjects volunteer and are assigned to the spaced-practice condition.
- Is this a sampling or a census?
 - Is this an observational study or an experiment?
 - This study involves two kinds of non-random sampling: (1) Subjects are not randomly sampled from some specified population and (2) Subjects are not randomly assigned to groups. Which problem is more serious? What affect on the results does each have?
19. A farmer believes that playing Barry Manilow songs to his peas will increase their yield. Describe a controlled experiment the farmer could use to test his theory.
20. A sports psychologist believes that people are more likely to be extroverted as adults if they played team sports as children. Describe two possible studies to test this theory. Design one as an observational study and the other as an experiment. Which is more practical?

Exploration

21. Studies are often done by pharmaceutical companies to determine the effectiveness of a treatment program. Suppose that a new AIDS antibody drug is currently under study. It is given to patients once the AIDS symptoms have revealed themselves. Of interest is the average length of time in months patients live once starting the treatment. Two researchers each follow a different set of 50 AIDS patients from the start of treatment until their deaths.
- What is the population of this study?
 - List two reasons why the data may differ.
 - Can you tell if one researcher is correct and the other one is incorrect? Why?
 - Would you expect the data to be identical? Why or why not?
 - If the first researcher collected her data by randomly selecting 40 states, then selecting 1 person from each of those states. What sampling method is that?
 - If the second researcher collected his data by choosing 40 patients he knew. What sampling method would that researcher have used? What concerns would you have about this data set, based upon the data collection method?
22. Find a newspaper or magazine article, or the online equivalent, describing the results of a recent study (the results of a poll are not sufficient). Give a summary of the study's findings, then analyze whether the article provided enough information to determine the validity of the conclusions. If not, produce a list of things that are missing from the article that would help you determine the validity of the study. Look for the things discussed in the text: population, sample, randomness, blind, control, placebos, etc.