IN THIS CHAPTER

- **Finding out what the process of statistics is all about**
- **Gaining success with statistics in your everyday life, your career, and in the classroom**

Chapter 1

Statistics in a Nutshell

The world today is overflowing with data to the point where anyone (even me!) can be overwhelmed. I wouldn't blame you if you were cynical right now about statistics you read about in the media — I am too at times. The good news is that while a great deal of misleading and incorrect information is lying out there waiting for you, a lot of great stuff is also being produced; for example, many studies and techniques involving data are helping improve the quality of our lives. Your job is to be able to sort out the good from the bad and be confident in your ability to do that. Through a strong understanding of statistics and statistical procedures, you gain power and confidence with numbers in your everyday life, in your job, and in the classroom. That's what this book is all about.

In this chapter, I give you an overview of the role statistics plays in today's data-packed society and what you can do to not only survive but thrive. You get a much broader view of statistics as a partner in the scientific method — designing effective studies, collecting good data, organizing and analyzing the information, interpreting the results, and making appropriate conclusions. (And you thought statistics was just number-crunching!)

Thriving in a Statistical World

It's hard to get a handle on the flood of statistics that affect your daily life in large and small ways. It begins the moment you wake up in the morning and check the news and listen to the meteorologist give you her predictions for the weather based on her statistical analyses of past data and present weather conditions. You pore over nutritional information on the side of your cereal box while you eat breakfast. At work you pull numbers from charts and tables, enter data into spreadsheets, run diagnostics, take measurements, perform calculations, estimate expenses, make decisions using statistical baselines, and order inventory based on past sales data.

At lunch you go to the No. 1 restaurant based on a survey of 500 people. You eat food that was priced based on marketing data. You go to your doctor's appointment where they take your blood pressure, temperature, weight, and do a blood test; after all the information is collected, you get a report showing your numbers and how you compare to the statistical norms.

You head home in your car that's been serviced by a computer running statistical diagnostics. When you get home, you turn on the news and hear the latest crime statistics, see how the stock market performed, and discover how many people visited the zoo last week.

At night, you brush your teeth with toothpaste that's been statistically proven to fight cavities, read a few pages of your *New York Times* Best-Seller (based on statistical sales estimates), and go to sleep — only to start it all over again the next morning. But how can you be sure that all those statistics you encounter and depend on each day are correct? In Chapter 2, I discuss in more depth a few examples of how statistics is involved in our lives and workplaces, what its impact is, and how you can raise your awareness of it.

WARNING

Some statistics are vague, inappropriate, or just plain wrong. You need to become more aware of the statistics you encounter each day and train your mind to stop and say "wait a minute!", sift through the information, ask questions, and raise red flags when something's not quite right. In Chapter 3, you see ways in which you can be misled by bad statistics and develop skills to think critically and identify problems before automatically believing results.

Like any other field, statistics has its own set of jargon, and I outline and explain some of the most commonly used statistical terms in Chapter 4. Knowing the language increases your ability to understand and communicate statistics at a higher level without being intimidated. It raises your credibility when you use precise terms to describe what's wrong with a statistical result (and why). And your presentations involving statistical tables, graphs, charts, and analyses will be informational and effective. (Heck, if nothing else, you need the jargon because I use it throughout this book; don't worry though, I always review it.)

In the next sections, you see how statistics is involved in each phase of the scientific method.

Designing Appropriate Studies

Everyone's asking questions, from drug companies to biologists; from marketing analysts to the U.S. government. And ultimately, everyone will use statistics to help them answer their questions. In particular, many medical and psychological studies are done because someone wants to know the answer to a question. For example,

- Will this vaccine be effective in preventing the flu?
- What do Americans think about the state of the economy?
- Does an increase in the use of social networking Web sites cause depression in teenagers?

The first step after a research question has been formed is to design an effective study to collect data that will help answer that question. This step amounts to figuring out what process you'll use to get the data you need. In this section, I give an overview of the two major types of studies — surveys and experiments — and explore why it's so important to evaluate how a study was designed before you believe the results.

Surveys

An *observational study* is one in which data is collected on individuals in a way that doesn't affect them. The most common observational study is the survey. *Surveys* are questionnaires that are presented to individuals who have been selected from a population of interest. Surveys take many different forms: paper surveys sent through the mail, questionnaires on Web sites, call-in polls conducted by TV networks, phone surveys, and so on.

WARNING

If conducted properly, surveys can be very useful tools for getting information. However, if not conducted properly, surveys can result in bogus information. Some problems include improper wording of questions, which can be misleading, lack of response by people who were selected to participate, or failure to include an entire group of the population. These potential problems mean a survey has to be well thought out before it's given.

REMEMBER

Many researchers spend a great deal of time and money to do good surveys, and you'll know (by the criteria I discuss in Chapter 16) that you can trust them. However, as you are besieged with so many different types of surveys found in the media, in the workplace, and in many of your classes, you need to be able to quickly examine and critique how a survey was designed and conducted and be able to point out specific problems in a well-informed way. The tools you need for sorting through surveys are found in Chapter 16.

Experiments

An *experiment* imposes one or more treatments on the participants in such a way that clear comparisons can be made. After the treatments are applied, the responses are recorded. For example, to study the effect of drug dosage on blood pressure, one group may take 10 mg of the drug, and another group may take 20 mg. Typically, a control group is also involved, in which subjects each receive a fake treatment (a sugar pill, for example), or a standard, nonexperimental treatment (like the existing drugs given to AIDS patients.)

REMEMBER

Good and credible experiments are designed to minimize bias, collect lots of good data, and make appropriate comparisons (treatment group versus control group). Some potential problems that occur with experiments include researchers and/or subjects who know which treatment they got, factors not controlled for in the study that affect the outcome (such as weight of the subject when studying drug dosage), or lack of a control group (leaving no baseline to compare the results with).

But when designed correctly, an experiment can help a researcher establish a cause-and-effect relationship if the difference in responses between the treatment group and the control group is statistically significant (unlikely to have occurred just by chance).

WARNING

Experiments are credited with helping to create and test drugs, determining best practices for making and preparing foods, and evaluating whether a new treatment can cure a disease, or at least reduce its impact. Our quality of life has certainly been improved through the use of well-designed experiments. However, not all experiments are well-designed, and your ability to determine which results are credible and which results are incredible (pun intended) is critical, especially when the findings are very important to you. All the info you need to know about experiments and how to evaluate them is found in Chapter 17.

Collecting Quality Data

After a study has been designed, be it a survey or an experiment, the individuals who will participate have to be selected, and a process must be in place to collect the data. This phase of the process is critical to producing credible data in the end, and this section hits the highlights.

Selecting a good sample

REMEMBER

Statisticians have a saying, "Garbage in equals garbage out." If you select your *subjects* (the individuals who will participate in your study) in a way that is *biased* — that is, favoring certain individuals or groups of individuals — then your results will also be biased. It's that simple.

Suppose Bob wants to know the opinions of people in your city regarding a proposed casino. Bob goes to the mall with his clipboard and asks people who walk by to give their opinions. What's wrong with that? Well, Bob is only going to get the opinions of a) people who shop at that mall; b) on that particular day; c) at that particular time; d) and who take the time to respond.

Those circumstances are too restrictive — those folks don't represent a cross section of the city. Similarly, Bob could put up a Web site survey and ask people to use it to vote. However, only people who know about the site, have Internet access, and want to respond will give him data, and typically only those with strong opinions will go to such trouble. In the end, all Bob has is a bunch of biased data on individuals that don't represent the city at all.

REMEMBER

To minimize bias in a survey, the key word is *random.* You need to select your sample of individuals *randomly* — that is, with some type of "draw names out of a hat" process. Scientists use a variety of methods to select individuals at random, and you see how they do it in Chapter 16.

Note that in designing an experiment, collecting a random sample of people and asking them to participate often isn't ethical because experiments impose a treatment on the subjects. What you do is send out requests for volunteers to come to you. Then you make sure the volunteers you select from the group represent the population of interest and that the data is well collected on those individuals so the results can be projected to a larger group. You see how that's done in Chapter 17.

After going through Chapters 16 and 17, you'll know how to dig down and analyze others' methods for selecting samples and even be able to design a plan you can use to select a sample. In the end, you'll know when to say "Garbage in equals garbage out."

Avoiding bias in your data

Bias is the systematic favoritism of certain individuals or certain responses. Bias is the nemesis of statisticians, and they do everything they can to minimize it. Want an example of bias? Say you're conducting a phone survey on job satisfaction of Americans; if you call people at home during the day between 9 a.m. and 5 p.m., you miss out on everyone who works during the day. Maybe day workers are more satisfied than night workers.

You have to watch for bias when collecting survey data. For instance: Some surveys are too long — what if someone stops answering questions halfway through? Or what if they give you misinformation and tell you they make $100,000 a year instead of $45,000? What if they give you answers that aren't on your list of possible answers? A host of problems can occur when collecting survey data, and you need to be able to pinpoint those problems.

WARNING

Experiments are sometimes even more challenging when it comes to bias and collecting data. Suppose you want to test blood pressure; what if the instrument you're using breaks during the experiment? What if someone quits the experiment halfway through? What if something happens during the experiment to distract the subjects or the researchers? Or they can't find a vein when they have to do a blood test exactly one hour after a dose of a drug is given? These problems are just some examples of what can go wrong in data collection for experiments, and you have to be ready to look for and find these problems.

After you go through Chapter 16 (on samples and surveys) and Chapter 17 (on experiments), you'll be able to select samples and collect data in an unbiased way, being sensitive to little things that can really influence the results. And you'll have the ability to evaluate the credibility of statistical results and to be heard, because you'll know what you're talking about.

Creating Effective Summaries

After good data have been collected, the next step is to summarize them to get a handle on the big picture. Statisticians describe data in two major ways: with numbers (called *descriptive statistics*) and with pictures (that is, charts and graphs).

Descriptive statistics

REMEMBER

Descriptive statistics are numbers that describe a data set in terms of its important features:

- If the data are *categorical* (where individuals are placed into groups, such as gender or political affiliation), they are typically summarized using the number of individuals in each group (called the *frequency*) or the percentage of individuals in each group (called the *relative frequency*).
- *Numerical data* represent measurements or counts, where the actual numbers have meaning (such as height and weight). With numerical data, more features can be summarized besides the number or percentage in each group. Some of these features include
  - Measures of center (in other words, where is the "middle" of the data?)
  - Measures of spread (how diverse or how concentrated are the data around the center?)
  - If appropriate, numbers that measure the relationship between two variables (such as height and weight)

WARNING

Some descriptive statistics are more appropriate than others in certain situations; for example, the average isn't always the best measure of the center of a data set; the median is often a better choice. And the standard deviation isn't the only measure of variability on the block; the interquartile range has excellent qualities too. You need to be able to discern, interpret, and evaluate the types of descriptive statistics being presented to you on a daily basis and to know when a more appropriate statistic is in order.

The descriptive statistics you see most often are calculated, interpreted, compared, and evaluated in Chapter 5. These commonly used descriptive statistics include frequencies and relative frequencies (counts and percents) for categorical data; and the mean, median, standard deviation, percentiles, and their combinations for numerical data.

Charts and graphs

Data is summarized in a visual way using charts and/or graphs. These are displays that are organized to give you a big picture of the data in a flash and/or to zoom in on a particular result that was found. In this world of quick information and mini-sound bites, graphs and charts are commonplace. Most graphs and charts make their points clearly, effectively, and fairly; however, they can leave room for too much poetic license, and as a result, can expose you to a high number of misleading and incorrect graphs and charts.

REMEMBER

In Chapters 6 and 7, I cover the major types of graphs and charts used to summarize both categorical and numerical data (see the preceding section for more about these types of data). You see how to make them, what their purposes are, and how to interpret the results. I also show you lots of ways that graphs and charts can be made to be misleading and how you can quickly spot the problems. It's a matter of being able to say "Wait a minute here! That's not right!" and knowing why not. Here are some highlights:

- Some of the basic graphs used for categorical data include pie charts and bar graphs, which break down variables, such as gender or which applications are used on teens' cellphones. A bar graph, for example, may display opinions on an issue using five bars labeled in order from "Strongly Disagree" up through "Strongly Agree." Chapter 6 gives you all the important info on making, interpreting, and, most importantly, evaluating these charts and graphs for fairness. You may be surprised to see how much can go wrong with a simple bar chart.
- For numerical data such as height, weight, time, or amount, a different type of graph is needed. Graphs called histograms and boxplots are used to summarize numerical data, and they can be very informative, providing excellent on-the-spot information about a data set. But of course they also can be misleading, either by chance or even by design. (See Chapter 7 for the scoop.)

WARNING

You're going to run across charts and graphs every day — you can open a newspaper and probably find several graphs without even looking hard. Having a statistician's magnifying glass to help you interpret the information is critical so that you can spot misleading graphs before you draw the wrong conclusions and possibly act on them. All the tools you need are ready for you in Chapter 6 (for categorical data) and Chapter 7 (for numerical data).

Determining Distributions

A *variable* is a characteristic that's being counted, measured, or categorized. Examples include gender, age, height, weight, or number of pets you own. A *distribution* is a listing of the possible values of a variable (or intervals of values), and how often (or at what density) they occur. For example, the distribution of gender at birth in the United States has been estimated at 52.4% male and 47.6% female.

REMEMBER

Different types of distributions exist for different variables. The following three distributions are the most commonly occurring distributions in an introductory statistics course, and they have many applications in the real world:

» If a variable is counting the number of successes in a certain number of trials (such as the number of people who got well by taking a certain drug), it has a *binomial* distribution.

» If the variable takes on values that occur according to a "bell-shaped curve," such as national achievement test scores, then that variable has a *normal* distribution.

» If the variable is based on sample averages and you have limited data, such as in a test of only ten subjects to see if a weight-loss program works, the *t*-distribution may be in order.

When it comes to distributions, you need to know how to decide which distribution a particular variable has, how to find probabilities for it, and how to figure out what the long-term average and standard deviation of the outcomes would be. To get you squared away on these issues, I've got three chapters for you, one dedicated to each distribution: Chapter 8 is all about the binomial, Chapter 9 handles the normal, and Chapter 10 focuses on the *t*-distribution.

TIP

For those of you taking an introductory statistics course (or any statistics course, for that matter), you know that one of the most difficult topics to understand is sampling distributions and the Central Limit Theorem (these two things go hand in hand). Chapter 11 walks you through these topics step by step so you understand what a sampling distribution is, what it's used for, and how it provides the foundation for data analyses like hypothesis tests and confidence intervals (see the next section for more about analyzing data). When you understand the Central Limit Theorem, it actually helps you solve difficult problems more easily, and all the keys to this information are there for you in Chapter 11.

Performing Proper Analyses

After the data have been collected and described using numbers and pictures, then comes the fun part: navigating through that black box called the *statistical analysis*. If the study has been designed properly, the original questions can be answered using the appropriate analysis — the operative word here being *appropriate*.

REMEMBER

Many types of analyses exist, and choosing the right analysis for the right situation is critical, as is interpreting results properly, being knowledgeable of the limitations, and being able to evaluate others' choice of analyses and the conclusions they make with them.

In this book, you get all the information and tools you need to analyze data using the most common methods in introductory statistics: confidence intervals, hypothesis tests, correlation and regression, and the analysis of two-way tables. This section gives you a basic overview of those methods.

Margin of error and confidence intervals

You often see statistics that try to estimate numbers pertaining to an entire population; in fact, you see them almost every day in the form of survey results. The media tells you what the average gas price is in the U.S., how Americans feel about the job the president is doing, or how many hours people spend on the Internet each week.

But no one can give you a single-number result and claim it's an accurate estimate of the entire population unless he collected data on every single member of the population. For example, you may hear that 60 percent of the American people support the president's approach to healthcare, but you know they didn't ask you, so how could they have asked everybody? And since they didn't ask everybody, you know that a one-number answer isn't going to cut it.

What's really happening is that data is collected on a sample from the population (for example, the Gallup Organization calls 2,500 people at random), the results from that sample are analyzed, and conclusions are made regarding the entire population (for example, all Americans) based on those sample results.

REMEMBER

The bottom line is, sample results vary from sample to sample, and this amount of variability needs to be reported (but it often isn't). The statistic used to measure and report the level of precision in someone's sample results is called the *margin of error.* In this context, the word *error* doesn't mean a mistake was made; it just means that because you didn't sample the entire population, a gap will exist between your results and the actual value you are trying to estimate for the population.

For example, someone finds that 60% of the 1,200 people surveyed support the president's approach to healthcare and reports the results with a margin of error of plus or minus 2%. This final result, in which you present your findings as a range of likely values between 58% and 62%, is called a *confidence interval.*

WARNING

Everyone is exposed to results including a margin of error and confidence intervals, and with today's data explosion, many people are also using them in the workplace. Be sure you know what factors affect margin of error (like sample size) and what the makings of a good confidence interval are and how to spot them. You should also be able to find your own confidence intervals when you need to.

In Chapter 12, you find out everything you need to know about the margin of error: All the components of it, what it does and doesn't measure, and how to calculate it for a number of situations. Chapter 13 takes you step by step through the formulas, calculations, and interpretations of confidence intervals for a population mean, population proportion, and the difference between two means and proportions.

Hypothesis tests

One main staple of research studies is called hypothesis testing. A *hypothesis test* is a technique for using data to validate or invalidate a claim about a population. For example, a politician may claim that 80% of the people in her state agree with her — is that really true? Or, a company may claim that they deliver pizzas in 30 minutes or less; is that really true? Medical researchers use hypothesis tests all the time to test whether or not a certain drug is effective, to compare a new drug to an existing drug in terms of its side effects, or to see which weight-loss program is most effective with a certain group of people.

REMEMBER

The elements about a population that are most often tested are

- The population mean (Is the average delivery time of 30 minutes really true?)
- The population proportion (Is it true that 80% of the voters support this candidate, or is it less than that?)
- The difference in two population means or proportions (Is it true that the average weight loss on this new program is 10 pounds more than the most popular program? Or, is it true that this drug decreases blood pressure by 10% more than the current drug?)

WARNING

Hypothesis tests are used in a host of areas that affect your everyday life, such as medical studies, advertisements, polling data, and virtually anywhere that comparisons are made based on averages or proportions. And in the workplace, hypothesis tests are used heavily in areas like marketing, where you want to determine whether a certain type of ad is effective or whether a certain group of individuals buys more or less of your product now compared to last year.

Often you only hear the conclusions of hypothesis tests (for example, this drug is significantly more effective and has fewer side effects than the drug you are using now); but you don't see the methods used to come to these conclusions. Chapter 14 goes through all the details and underpinnings of hypothesis tests so you can conduct and critique them with confidence. Chapter 15 cuts right to the chase of providing step-by-step instructions for setting up and carrying out hypothesis tests for a host of specific situations (one population mean, one population proportion, the difference of two population means, and so on).

After reading Chapters 14 and 15, you'll be much more empowered when you need to know things like which group you should be marketing a product to; which brand of tires will last the longest; whether a certain weight-loss program is effective; and bigger questions like which surgical procedure you should opt for.

Correlation, regression, and two-way tables

One of the most common goals of research is to find links between variables. For example,

- Which lifestyle behaviors increase or decrease the risk of cancer?
- What side effects are associated with this new drug?
- Can I lower my cholesterol by taking this new herbal supplement?
- Does spending a large amount of time on the Internet cause a person to gain weight?

Finding links between variables is what helps the medical world design better drugs and treatments, provides marketers with info on who is more likely to buy their products, and gives politicians information on which to build arguments for and against certain policies.

WARNING

In the mega-business of looking for relationships between variables, you find an incredible number of statistical results — but can you tell what's correct and what's not? Many important decisions are made based on these studies, and it's important to know what standards need to be met in order to deem the results credible, especially when a cause-and-effect relationship is being reported.

Chapter 18 breaks down all the details and nuances of plotting data from two numerical variables (such as dosage level and blood pressure), finding and interpreting *correlation* (the strength and direction of the linear relationship between *x* and *y*), finding the equation of a line that best fits the data (and when doing so is appropriate), and how to use these results to make predictions for one variable based on another (called *regression*). You also gain tools for investigating when a line fits the data well and when it doesn't, and what conclusions you can make (and shouldn't make) in the situations where a line does fit.

I cover methods used to look for and describe links between two categorical variables (such as the number of doses taken per day and the presence or absence of nausea) in detail in Chapter 19. I also provide info on collecting and organizing data into *two-way tables* (where the possible values of one variable make up the rows and the possible values for the other variable make up the columns), interpreting the results, analyzing the data from two-way tables to look for relationships, and checking for independence. And, as I do throughout this book, I give you strategies for critically examining results of these kinds of analyses for credibility.

Drawing Credible Conclusions

WARNING

To perform statistical analyses, researchers use statistical software that depends on formulas. But formulas don't know whether they are being used properly, and they don't warn you when your results are incorrect. At the end of the day, computers can't tell you what the results mean; you have to figure it out. Throughout this book you see what kinds of conclusions you can and can't make after the analysis has been done. The following sections provide an introduction to drawing appropriate conclusions.

Reeling in overstated results

Some of the most common mistakes made in conclusions are overstating the results or generalizing the results to a larger group than was actually represented by the study. For example, a professor wants to know which Super Bowl commercials viewers liked best. She gathers 100 students from her class on Super Bowl Sunday and asks them to rate each commercial as it is shown. A top-five list is formed, and she concludes that all Super Bowl viewers liked those five commercials the best. But she really only knows which ones *her students* liked best — she didn't study any other groups, so she can't draw conclusions about all viewers.

Questioning claims of cause and effect

One situation in which conclusions cross the line is when researchers find that two variables are related (through an analysis such as regression; see the earlier section "Correlation, regression, and two-way tables" for more info) and then automatically leap to the conclusion that those two variables have a cause-and-effect relationship.

For example, suppose a researcher conducted a health survey and found that people who took vitamin C every day reported having fewer colds than people who didn't take vitamin C every day. Upon finding these results, she wrote a paper and gave a press release saying vitamin C prevents colds, using this data as evidence.

Now, while it may be true that vitamin C does prevent colds, this researcher's study can't claim that. Her study was observational, which means she didn't control for any other factors that could be related to both vitamin C and colds. For example, people who take vitamin C every day may be more health conscious overall, washing their hands more often, exercising more, and eating better foods; all these behaviors may be helpful in reducing colds.

REMEMBER

Until you do a controlled experiment, you can't make a cause-and-effect conclusion based on relationships you find. (I discuss experiments in more detail earlier in this chapter.)

Becoming a Sleuth, Not a Skeptic

Statistics is about much more than numbers. To really "get" statistics, you need to understand how to make appropriate conclusions from studying data and be savvy enough to not believe everything you hear or read until you find out how the information came about, what was done with it, and how the conclusions were drawn. That's something I discuss throughout the book, but I really zoom in on it in Chapter 20, which gives you ten ways to be a statistically savvy sleuth by recognizing common mistakes made by researchers and the media.

TIP

For you students out there, Chapter 21 brings good statistical practice into the exam setting and gives you tips on increasing your scores. Much of my advice is based on understanding the big picture as well as the details of tackling statistical problems and coming out a winner on the other side.

REMEMBER

Becoming skeptical or cynical about statistics is very easy, especially after finding out what's going on behind the scenes; don't let that happen to you. You can find lots of good information out there that can affect your life in a positive way. Find a good channel for your skepticism by setting two personal goals:

- To become a well-informed consumer of the statistical information you see every day
- To establish job security by being the statistics "go-to" person who knows when and how to help others and when to find a statistician

Through reading and using the information in this book, you'll be confident in knowing you can make good decisions about statistical results. You'll conduct your own statistical studies in a credible way. And you'll be ready to tackle your next office project, critically evaluate that annoying political ad, or ace your next exam!