CHAPTER 15

Developing the Sampling Plan

Learning Objectives

1. Explain the difference between a census and a sample.
2. List the six steps researchers use to draw a sample of a population.
3. Explain the difference between a parameter and a statistic.
4. Explain the difference between a probability sample and a nonprobability sample.
5. Explain what a judgment sample is and describe its best use and its hazards.
6. Define *quota sample.*
7. Specify the two procedures that distinguish a stratified sample.
8. Cite two reasons researchers might choose to use a stratified sample rather than a simple random sample.
9. Explain the difference between a proportionate stratified sample and a disproportionate stratified sample.
10. List the steps followed in drawing a cluster sample.

Introduction

Once you've specified the problem, developed an appropriate research design, and carefully crafted your data collection instrument, the next step in a marketing research project is to develop the sampling plan. The sampling plan refers to the process of selecting the people or objects (that is, companies, products, etc.) to be surveyed, interviewed, or observed. One way to do this is to collect information from or about each member of the population—this is known as a **census**. Another way would be to collect information from a portion of the population by taking a **sample** of elements from the larger group and, on the basis of the information collected from the subset, to make projections about what would be true for the population. As we'll see, the ability to make inferences about the overall population based on information collected from only some of the population members depends on how we select the sample.

census
A type of sampling plan in which data are collected from or about each member of a population.

sample
Selection of a subset of elements from a larger group of objects.

Figure 15.1 lays out a six-step process for drawing a sample and collecting data. In this chapter, we focus on the first three steps of the process. In Chapter 16, we'll consider sample size issues.

There are several reasons why working with a sample is often better than trying to conduct a census. First, complete counts on populations of even moderate size are often costly and time consuming. Often, the information will be obsolete by the time the census is completed and the information processed. In some cases, a census is impossible. If, for example, researchers wanted to test the life of a company's electric light bulbs by leaving all its inventory of bulbs on until they burned out, they would have reliable data, but no product to sell.

Finally—and this one may surprise you—a sample might give more accurate results than a census. Conducting a census rather than a sample will often require larger field staffs, which, in turn, introduces greater potential for some kinds of error. As an example, just how excited is a U.S. census worker likely to be about visiting a household on the tenth floor of an old apartment building with no working elevators in a crime-ridden part of town? It's no surprise that the U.S. census tends to undercount people in difficult areas. Because of things like this, the U.S. Bureau of the Census uses sample surveys to check the accuracy of various censuses.

Learning Objective

1. Explain the difference between a census and a sample.

DEFINING THE TARGET POPULATION

Learning Objective

2. List the six steps researchers use to draw a sample of a population.

The first step in the process outlined in Figure 15.1 is to define the target population. By **population**, we mean all the individuals or objects that meet certain requirements for membership in the overall group. We often refer to those who qualify as *population elements*. For example, a study aimed at establishing a demographic profile of frozen pizza eaters requires specifying who is to be considered a frozen pizza eater. Anyone who has ever eaten a frozen pizza? Those who eat at least one frozen pizza a month? Those who eat a certain minimum number of frozen pizzas per month?

population
All cases that meet designated specifications for membership in the group.

You must be very clear and precise in defining the population. For example, does the population consist of individuals, households, business firms, other institutions, credit card transactions, light bulbs on an assembly line, or something else? In making these decisions, it sometimes helps to specify what units are *not* to be included. Geographic boundaries and a time period for the study must always be specified, although additional restrictions are often placed on the elements. When the elements are individuals, for example, the relevant target population may be defined as all those over 18 years of age, or females only, or those with a high school education only, or those who have visited a certain restaurant within the last 30 days. When establishing

FIGURE 15.1 Six-Step Procedure for Drawing a Sample

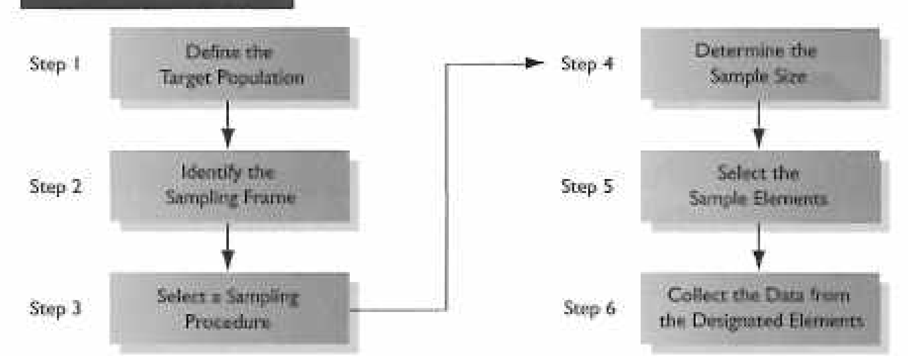


© TOM YANG/INDEX STOCK IMAGERY

Researchers must carefully define a population in order to get accurate results. Should anyone who has ever eaten a frozen pizza be included when determining the demographics of frozen pizza eaters?

criteria for population membership, remember to think broadly, as the following classic example illustrates:

> Take the case of the manufacturer of dog food ... who went out and did an intensive market study. He tested the demand for dog food; he tested the package size, the design, the whole advertising program. Then he launched the product with a big campaign, got the proper distribution channels, put it on the market, and had tremendous sales. But two months later, the bottom dropped out—no follow-up sales. So he called in an expert, who took the dog food out to the local pound, put it in front of the dogs—and they would not touch it. For all the big marketing study, no one had tried the product on the dogs.[1]

Somewhere along the way it would have been helpful to have included a study in which dogs comprised the population. That didn't happen, however, probably because it is people who buy dog food and not the dogs themselves. Nevertheless, the careless specification of population elements can lead to bad consequences.

The problem of specifying the geographic boundaries for the target population is sometimes more difficult in international marketing research studies because of the additional complexity an international perspective introduces. For example, urban versus rural areas may be significantly different from each other in various countries. Also, the composition of the population can vary depending on the location within the country. In Chile, for example, the north has a highly centralized Indian population, whereas the south has high concentrations of persons of European descent.

incidence
The percent of a general population or group that qualifies for inclusion in the population.

In general, the simpler the definition of the target population, the higher the incidence and the easier and less costly it is to find the sample. **Incidence** refers to the proportion of a general grouping of people or objects that meets the criteria to be a member of the defined population. For example, if we start with the general public in the United States (that is, everybody living in the United States) and we define our study population to be all females living in the United States, the incidence rate is about 51%. Incidence has a direct bearing on the time and cost it takes to complete studies. When incidence is high (i.e., most elements in the general population qualify for the study because only one or very few easily satisfied criteria are used to screen potential respondents), the cost and time to collect data are minimized. Alternatively, as the number of criteria for population membership increases, so do the cost and time necessary to find them.

Exhibit 15.1 shows the number of people ages seven years and older who are estimated to have participated in various sporting activities during 2007. The data in Exhibit 15.1 suggest that it would be more difficult and costly to focus a study on people who play ice hockey, only 2.1 million people, than people who walk for exercise, 89.8 million people. Here's the important thing: You must be precise in specifying exactly what elements are of interest and what elements are to be excluded. A clear statement of research purpose is very important for determining the appropriate elements of interest.

Learning Objective

3. Explain the difference between a parameter and a statistic.

Parameters versus Statistics

Before we go on, let's revisit why we are drawing a sample in the first place. Our goal with a sample is to determine what is likely to be true for a population based on data obtained from only a subset of that population. We typically work with a sample, rather than a census, because a sample is easier and less costly to obtain than is a census.

Exhibit 15.1 Number of U.S. Participants in Various Sports (individuals seven years and older who participated more than once)

Sport	Number (in millions)
Exercise Walking	89.8
Exercising with Equipment	52.8
Swimming	52.3
Camping (vacation/overnight)	47.5
Bowling	43.5
Bicycle Riding	37.4
Fishing	35.3
Workout at Club	33.8
Weight Lifting	33.2
Boating, Motor/Power	31.9
Running/Jogging	30.4
Aerobic Exercising	30.3
Billiards/Pool	29.5
Hiking	28.6
Basketball	24.1
Golf	22.7
Target Shooting	20.9
Hunting with Firearms	19.5
Baseball	14.0
Soccer	13.8
Backpack/Wilderness Camp	13.0
Tennis	12.3
Dart Throwing	12.1
Volleyball	12.0
In-Line Roller Skating	10.7
Yoga	10.7
Scooter Riding	10.6
Skateboarding	10.1
Softball	10.0
Football (tackle)	9.2
Paintball Games	7.4
Mountain Biking (off road)	7.4
Target Shooting—Airgun	6.6
Archery (target)	6.6
Kayaking	5.9
Hunting w/Bow & Arrow	5.7
Skiing (alpine)	5.5
Water Skiing	5.3
Snowboarding	5.1
Mtn/Rock Climbing	4.6
Muzzleloading	3.6
Scuba Diving (open water)	2.4
Wrestling	2.1
Hockey (ice)	2.1
Skiing (cross-country)	1.7
Lacrosse	1.2

Source: "2007 Participation—Ranked by Total Participation," National Sporting Goods Association, downloaded from http://www.nsga.org, August 4, 2008.

Manager's Focus

Generating a representative sample is essential to producing useful market research intelligence. The first step in this process—defining the population—is more your responsibility than the researcher's. The definition of the target population should emerge directly from your market information needs. It involves answering the question "From whom do we need information?" Often, the answer to this question involves more than one group.

Target populations can be your current target markets, possible new target markets, former or disgruntled customers, loyal or important customers, key customer groups for competitors, certain types of marketing channel members, boundary-spanning employees within your own organization, or any other group that can provide information you need. Defining the population(s), therefore, is not an abstract statistical process but is instead a practical matter that is based directly on the marketing situations you are facing and the information you need to address those situations. Once you have guided the research provider to a proper definition of the target population(s) (which really should be done before the questionnaire is designed because researchers need to know with whom they will be communicating), your role in the sampling process is largely completed. You should step back and allow the researcher to create an appropriate plan for generating a sample from the target population(s).

parameter
A characteristic or measure of a population.

Any population has certain characteristics; these characteristics are called parameters, and we assume that if we could take measurements of these characteristics from all population elements without any kind of error getting into our data that we would know what is true about the population on these parameters. For example, suppose the population for a study consists of all adults living in Phoenix, Arizona. We could describe this population on a number of parameters, including average age, proportion with a college degree, range of incomes, attitude toward a new service offering, awareness of a new retail store that has just opened, and so on. Note that within the population, there is a real quantity or value for each of these parameters, even though we'll never know for sure what these true values are (because as a practical matter, we can never measure something without error).

statistic
A characteristic or measure of a sample.

sampling error
The difference between results obtained from a sample and results that would have been obtained had information been gathered from or about every member of the population.

When we work with a sample drawn from a population, we are attempting to describe the population parameters based on the measures we take from the sample members. That is, we calculate the average age, range of income, or awareness level for the sample as a means of gaining insights into what likely would be true for the population. In short, we work with statistics, which are characteristics or measures of a sample, to draw inferences about the larger population's parameters. When we work with a sample instead of a census, it is likely that our results will be at least a little different than they would have been had we gathered information from or about every member of the population. This difference is known as sampling error.

How big a problem is sampling error? It certainly is something that a researcher must take into account. Fortunately, sampling error can be estimated with relative ease, provided that you've drawn the right kind of sample. Plus, if a researcher wants to decrease sampling error, it is possible to do so—just increase the sample size. Because it is expensive to increase sample size, however, researchers must learn to balance the need for decreased sampling error (which results in greater precision and/or confidence in a statistic—see Chapter 16) against the costs of increasing the sample size.

IDENTIFYING THE SAMPLING FRAME

Once the population has been carefully defined, the next step is to find an adequate sampling frame, a listing of population elements from which the actual sample will be drawn. Suppose that the target population for a particular study is all the households in the metropolitan Dallas area. At first glance, the Dallas phone book would seem an easily accessible sampling frame. However, on closer examination, it becomes clear that the telephone directory provides an inaccurate listing of Dallas households, omitting those with unlisted numbers (and those without regular phones—many people now rely solely on cell phones) and double-counting those with multiple listings. People who have recently moved into the area and are not yet listed are also omitted.

sampling frame
The list of population elements from which a sample will be drawn; the list could consist of geographic areas, institutions, individuals, or other units.

Unfortunately, perfect sampling frames usually don't exist except in unusual circumstances. That makes developing an acceptable sampling frame one of a researcher's most important and creative tasks. In Chapter 10, we discussed random-digit dialing as one means of overcoming the lack of a solid sampling frame for telephone interview studies. Sometimes researchers sample geographic areas and then subsample within these areas when, say, the target population is individuals but a current, accurate list of appropriate individuals is not available—we'll have more to say on this later in this chapter.

Much of the time, researchers work with sampling frames that have been developed by companies that specialize in compiling databases and then selling the names, addresses, phone numbers, and/or e-mail addresses. For example, infoUSA, a database company located in Omaha, Nebraska, employs over 600 data compilers who gather information from a broad range of sources. Its list of millions of businesses is updated regularly based on telephone directory listings, annual reports, government data, the business press, and other sources. The company calls each business in its database to verify the accuracy of the information. As a result, researchers could easily develop a sampling frame for a fairly specific population of businesses, because lists can be screened by any information in the database, including industry, geographic location, size of company, credit rating, and so on. Research Window 15.1 provides information about several of the approaches used by Survey Sampling International to develop samples.

MANAGER'S FOCUS

A quick note about terminology might be helpful here. Nonresearchers commonly confuse the term "population parameter" with the criteria used to define the target population. For example, the target population might be defined as all women between the ages of 21 and 30 in Los Angeles, California. The criteria for defining this population are gender, age, and geographic location, but these are *not* population parameters. The population parameters are other characteristics about the population that we want to estimate. For example, we may want to estimate things like (a) income level, (b) educational attainment level, (c) attitudes toward brands in a specific product category, or (d) frequency of purchasing our brand in the last three months. These are the population parameters among Los Angeles women between the ages of 21 and 30 that we would like to estimate by generating statistics from our sample.

We have actually discussed population parameters at previous points without ever using the term. The information needs you identify when formulating the marketing problem as well as the variables you specify when creating dummy tables are in reality the population parameters you want to estimate in your study. As a result, in the earlier research stages you will have provided the researcher with the necessary input to understand what the relevant population parameters are.

research window 15.1

Survey Sampling International (SSI)

SSI specializes in the sample selection process across several methods of administration and across numerous countries.

© WINSTON DAVIDIAN/ISTOCKPHOTO

Mail Samples

Using information collected from telephone white pages and other sources, SSI has developed the largest available national database of U.S. households. Samples can be drawn for specific geographic areas using the systematic sampling approach. The company notes that its deliverable rate is about 80% on mail surveys after accounting for wrong addresses.

Telephone Samples

SSI was one of the early innovators in telephone research products, including random-digit dialing (RDD). The company offers RDD in 21 countries for consumer and business samples as well as RDD for wireless/mobile telephones in nine countries in North America and Western Europe. Samples can be developed based on age, income, and ethnic background.

Online Samples

SSI manages proprietary online communities of research respondents in 22 countries including the United States (www.surveyspot.com) and China (www.opinionworld.cn). Very specific samples can be developed based on respondent purchases, age, or interests.

Source: Developed from information downloaded from http://www.surveysampling.com, August 6, 2008.

Learning Objective

4. Explain the difference between a probability sample and a nonprobability sample.

SELECTING A SAMPLING PROCEDURE

The third step in the procedure for drawing a sample involves selecting a particular sampling procedure—unless researchers have decided to attempt a census. For small populations of individuals, stores, or other objects, a census probably makes sense. Even if some individuals can't be reached or don't respond, if the number in the population is less than 400–500, it is usually reasonable to attempt to gather information from or about each population element. This is not a hard and fast rule, however. In some cases, it might be desirable to conduct a census when numbers are much higher, or to draw a sample when the population size is smaller. As is true for most aspects of the marketing research process, the needs of the particular research situation must be taken into account.

Most of the time, researchers find it necessary to draw a sample from the population. As we mentioned before, if managers and researchers are interested in what would likely be true for the whole population, as opposed to just the population elements that made it into the sample, it is important to draw the right kind of sample. Sampling techniques can be divided into two broad categories: probability and nonprobability samples. In a **probability sample**, each member of the target population has a *known, nonzero* chance of being included in the sample. The chances of each member of the target population being included in the sample may not be equal, but everyone has some chance of being included. Plus, there is a random component, not under the control of the researcher, in how population elements are selected for the sample.

probability sample
A sample in which each target population element has a known, nonzero chance of being included in the sample.

With **nonprobability samples**, on the other hand, there is no way of estimating the probability that any target population element will be included in the sample. Thus, there is no way of ensuring that the sample is representative of the population. All nonprobability samples rely on personal judgment somewhere in the sample-selection process rather than on a random process to select sample members. While these judgments may sometimes result in good estimates of a population characteristic, there is no way of determining objectively if the sample is adequate. It is only when the target population elements have been selected with known probabilities that it is possible to evaluate the precision of a sample result. Knowing the precision of a result based on a sample is what allows us to draw inferences about the population from which the sample was drawn. For this reason, probability sampling is usually considered to be better than nonprobability sampling.

nonprobability sample
A sample that relies on personal judgment in the element selection process.

Samples can also be categorized by whether they are fixed or sequential samples, each of which can apply to both probability and nonprobability samples. In **fixed samples**, the sample size is decided before the study begins, and all the necessary information is collected before the results are analyzed. In a **sequential sample**, the number of elements to be sampled is not decided in advance but is determined by a series of decisions as the data are collected. For example, if the evidence is not conclusive with an initial small sample, more observations will be made. If the results are still inconclusive, the size of the sample will be expanded further. At each stage, a decision is made as to whether more information should be collected or whether the evidence is now sufficient to permit a conclusion. The sequential sample allows trends in the data to be evaluated as the data are being collected. It also allows researchers to reduce costs when it is clear that additional data would add little information.

fixed sample
A sample for which size is determined in advance and needed information is collected from the designated elements.

sequential sample
A sample formed on the basis of a series of successive decisions.

Figure 15.2 shows the basic types of samples broken into two categories, nonprobability and probability samples. These basic sample types can be combined into more complex sampling plans when necessary. If you understand the basic types, though, you should be able to understand the more complex designs. In the following sections, we describe the different types of sampling plans.

FIGURE 15.2 Classification of Sampling Techniques

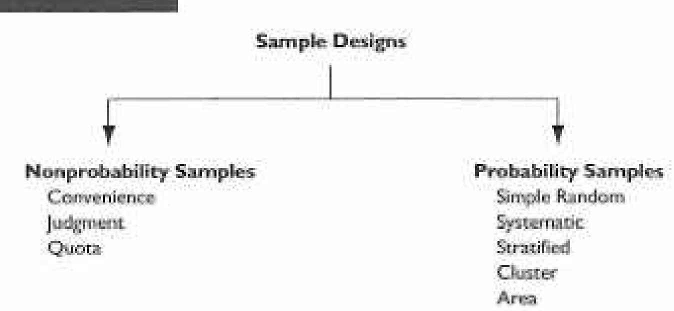

NONPROBABILITY SAMPLES

Nonprobability samples involve personal judgment somewhere in the selection process. Sometimes the researcher chooses who should be included in the sample; sometimes field workers choose who they should approach; and sometimes the choices involve when and where to collect data. With all nonprobability samples, however, objectivity about which population elements get into the sample is missing and not all population elements have an opportunity to be included. As a result, it is impossible to assess the degree of sampling error. Without knowing how much sampling error results from a particular sampling procedure, we can't gauge the accuracy of the estimates with any precision. As a result, we know the results for the particular elements that made it into our sample, but we can't say anything at all about what would have been true for the overall population. Managers may still choose to act on the basis of results from nonprobability samples, but they are taking risks when they do so.

Convenience Samples

convenience sample A nonprobability sample in which population elements are included in the sample because they were readily available.

With **convenience samples**, the name says it all: Inclusion in the sample is a matter of convenience. People or objects are selected for the sample because they happen to be in the right place at the right time to be included. Convenience samples are easy—just go out and find a location where lots of people who are likely to be members of the population are located and do interviews or pass out surveys. Lots of organizations put surveys on Web sites so that people who visit the sites can respond electronically—but what about people who don't visit the Web site? Sometimes radio and television programs invite their audiences to call a certain telephone number or to go to a particular Web site to respond to the "question of the day" or to vote for some option. All of these examples, and many more, amount to convenience samples. They are typically quick and easy ways of collecting data. If there is no need to ensure that the sample adequately represents the population, a convenience sample is probably the way to go. For instance, convenience samples are commonly used with exploratory research, where the goal is to generate insights or to develop hypotheses.

Problems arise, however, when people begin to draw important conclusions based on data from convenience samples. The main problem is that we have no way of knowing if those included in a convenience sample are representative of the larger target population. As a simple example, passing out surveys to passersby at the corner of Manvel Avenue and Tenth Street in a certain city during business hours on a Tuesday means that anyone who happened *not* to be at that corner during that time period had no chance of participating. It's very likely that important points of view may have been missed. Or maybe those points of view *were* represented—the problem is that we can never know for sure. As a result, we can't assume that the results of the study would apply to the whole population. And using bigger numbers of respondents doesn't necessarily make the sample any more representative, as the following example illustrates.

Some years ago, one of the local television stations in Madison, Wisconsin, featured a daily public opinion poll on topics of interest to the local community. The polls were labeled the "Pulse of Madison" and were conducted in the following way. During the 6 o'clock news every evening, the station would ask a question about some controversial issue to which people could reply with a yes or no. Persons in favor would call one number; persons opposed would call another. The number of viewers calling each number was recorded electronically. Percentages of those in favor and opposed would then be reported on the 10 o'clock news. With some 500 to 1,000 people calling in their opinions each night, the local television commentator seemed to interpret these results as reflecting the true state of opinion in the community.

On one 6 o'clock broadcast, the following question was posed: "Do you think the drinking age in Madison should be lowered to 18?" The existing legal limit was 21. Almost 4,000 people called in that night with 78% in favor of lowering the drinking

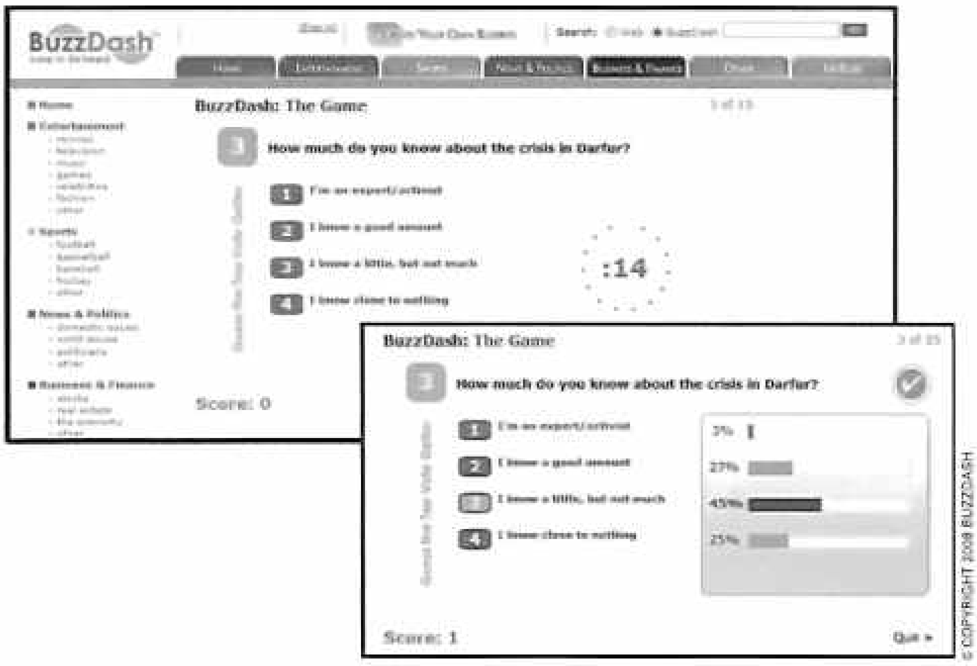

An example of a convenience sample could be asking a reader or viewer to go to a particular Web site to respond to the "question of the day" or to vote for some option.

Manager's Focus

Each of the sampling methods discussed in this chapter can be useful in the right circumstances. What you need to understand, then, are the situations in which each is appropriate. You have just learned that for each type of nonprobability sample there is no way to know or estimate the likelihood of selecting or not selecting specific population members for inclusion in the sample. As a result, you have no way of assessing sampling error. This does not mean nonprobability samples are unrepresentative. It is possible for them to be highly representative, but the likelihood of this occurring is lower and there is simply no way to assess this likelihood.

Consequently, nonprobability samples are most useful when it is not necessary to assume responses came from a representative sample. This is only the case when we are conducting exploratory research where our primary aim is to gain tentative insights about what might be happening or important in a market situation. Nonprobability samples should be avoided when the purpose is to describe or profile a population or to identify cause-and-effect relationships between variables. In order to be confident in our descriptions of a population or in our inferences about what marketing variables cause specific outcomes, we need information about the probability of population members being included in the sample. This means probability samples are more useful when conducting descriptive or causal studies.

Much too commonly, nonprobability samples, particularly convenience samples, are used in descriptive and causal studies in an effort to save time and money. If a sampling frame is available from which a probability sample can be drawn, it is unwise to trade off information about sample representativeness in favor of time or monetary savings. It is dangerous to make inferences about a population based on data from a nonprobability sample. If you choose to do so, you should understand that the little bit of money saved by employing a nonprobability sample may be offset by the costs of taking inappropriate marketing actions based on imprecise statistical estimates. In the long run, it is in your organization's best interest to invest in a quality probability sample to facilitate effective decision making. If you face a situation where it is very difficult or impossible to generate a probability sample, you should at least invest the time and money necessary to compare your nonprobability sample's characteristics against external information sources (such as census data) to determine how representative it is on key demographic or other characteristics.

age requirement. Surely 4,000 responses must be representative! Wrong. As you may have guessed, certain segments of the population were more interested in the issue than others. It was no surprise to later learn that university students had taken half-hour phone shifts on an arranged basis. Each person would call the yes number, hang up, call again, hang up, and so on, until it was the next person's turn. Thus, neither the size of the sample nor the proportion favoring the age change was surprising. The sample was simply not representative.

Learning Objective

5. Explain what a judgment sample is and describe its best use and its hazards.

judgment sample
A nonprobability sample in which the sample elements are handpicked because they are expected to serve the research purpose.

snowball sample
A judgment sample that relies on the researcher's ability to locate an initial set of respondents with the desired characteristics.

Judgment Samples

With judgment samples, the sample elements are handpicked by the researcher because it is expected that they can serve the research purpose. Procter & Gamble (P&G) used this method once when it advertised for "interns" ages 13 to 17 from the area around its Cincinnati headquarters. The company's food and beverage division hired this group of teenagers to serve as a kind of consumer panel. Working 10 hours a week in exchange for $1,000 and a trip to a concert, they reviewed television commercials, visited the mall with P&G managers to study retail displays, tested new products, and discussed their purchasing behavior. By selecting the panel members through a "hiring" process rather than randomly, the company could focus on traits it considered helpful—for example, the teenagers' ability to articulate their views clearly—at the risk that their views might not be representative of their age group.[2]

The snowball sample is a judgment sample that is sometimes used to sample special populations—populations with a low incidence rate within the general public, for example. This type of judgment sample relies on the researcher's ability to locate an initial set of respondents with the desired characteristics. These individuals are then used as informants to identify others with the desired characteristics.

Imagine, for example, that a company wanted to determine the desirability of a certain product that would enable deaf people to communicate over telephone lines. Researchers might begin by identifying some key people in the deaf community and asking them for names of other deaf people who might be used in the study. Those asked to participate would also be asked for names of others who might cooperate. In this way, the sample "snowballs" by getting larger as participants identify still other possible respondents.

As long as the researcher is at the early stages of research when ideas or insights are being sought—and when the researcher realizes its limitations—the judgment sample is perfectly appropriate. In some cases, it may be about the only way to develop a sample of people who meet specific criteria that don't occur frequently and/or cannot easily be observed. A judgment sample becomes dangerous, however, when it is used in descriptive or causal studies and its weaknesses are ignored.

Learning Objective

6. Define *quota sample*.

quota sample
A nonprobability sample chosen so that the proportion of sample elements with certain characteristics is about the same as the proportion of the elements with the characteristics in the target population.

Quota Samples

Researchers sometimes use another kind of nonprobability sample, the quota sample, in order to build a sample that mirrors the population on one or more important aspects. For example, imagine that you have been asked to draw a sample of 1,000 undergraduate students for short personal interviews on a college campus. If you used a quota sample, you'd make sure that your sample contained the right proportions of freshmen, sophomores, juniors, and seniors. You might also want to ensure that you had the right proportions of men and women in the sample, or maybe you'd want to have the right proportions of students from the various colleges or majors across campus. Again, the goal would be to build a sample that looked like the larger population of students.

Suppose that class and sex were the two most important variables and that 30% of the undergraduates on campus were freshmen, 25% were sophomores, 25% were juniors, and 20% were seniors. Furthermore, across all the classes, half of the students were men. Your quota sample would require that 300 respondents were first-year students (150 men, 150 women), 250 were sophomores (125 men, 125 women), 250 were juniors (again, 125 men, 125 women), and 200 were seniors (100 men, 100 women). If you had hired interviewers, you would complete the project by giving each interviewer a quota—thus, the name *quota sample*—specifying the types of undergraduates he or she is to contact. For example, one interviewer might be instructed to find and collect data from 50 students according to the following quota:

- 20 freshmen—5 male and 15 female
- 15 sophomores—10 male and 5 female
- 10 juniors—5 male and 5 female
- 5 seniors—2 male and 3 female

A quota sample of students on a college campus should be sure to include equal proportions of freshmen, sophomores, juniors, and seniors.

Across all interviewers, the total number of student respondents would add up to 1,000, and if your interviewers successfully reached their quotas, the proportions in the sample will match those in the population. Note, however, that the specific sample elements (i.e., students) to be used would be left to the discretion of the individual interviewers. That's what makes a quota sample a nonprobabilistic sampling plan. And even though the resulting sample *looks* like the overall population on two key aspects, it may not accurately reflect other aspects of the population.

The key problem with quota samples and other types of nonprobabilistic sampling plans is that biases usually influence the selection of population elements to be included in the sample. It's no surprise that people would rather interview their friends and neighbors, conduct in-home surveys in nicer parts of town, or collect data in a convenient location with lots of potential respondents (e.g., a shopping mall, airline terminal, or business district). These sorts of biases may or may not actually change the results of the study, but we don't know and can't correct for them when analyzing the data. When the sample elements are selected objectively, on the other hand, using a probabilistic sample, researchers have certain tools they can rely on to make the question of whether a particular sample is representative less difficult. With probability samples, we rely on the sampling procedure and not on the composition of the specific sample to solve the problem of representation.

ETHICAL DILEMMA 15.1

You are designing an experiment to compare the effectiveness of different types of commercials and need to recruit a large group of subjects of varying ages to watch television for an hour every night for a week. You approach your local church minister and tell him that you will make a donation to the church restoration fund for every member of the congregation who agrees to participate.

- When might incentives be coercive?
- Is it ethical to coerce people to participate in research?
- Will the quality of the data suffer from the coercive recruitment of participants?

PROBABILITY SAMPLES

In a probability sample, we can calculate the likelihood that any given population element will be included, because the final sample elements are selected objectively by a specific process and not according to the whims of the researcher or field-worker. Since the elements are selected objectively, we can make inferences to the larger population based on the results from the sample. With a probability sample, we can assess the likely amount of sampling error, the difference between the sample results and what would have been true had we taken measures from all population members. In this section, we'll introduce several different types of probability samples, starting with the most well known, the simple random sample.

Simple Random Samples

simple random sample
A probability sampling plan in which each unit included in the population has a known and equal chance of being selected for the sample.

Most people have had experience with simple random samples either in beginning statistics courses or in reading about results from such samples in newspapers or magazines. In a simple random sample, each unit included in the sample has a known and equal chance of being selected for study, and every combination of population elements is a sample possibility. For example, if we wanted a simple random sample of all students enrolled in a particular college, we might have a computer pick a sample randomly from a list of students in that college.

Simple random samples and other types of probability samples depend on the sampling distribution of the particular statistic being considered for the ability to draw inferences about the larger population (see Appendix 15A at the end of this chapter for more details). Even though any particular sample of population elements may include some unusual or extreme cases, the principles underlying the sampling distribution of the statistic allow us to establish the range within which the population parameter is likely to fall if all cases in the population were considered. This range is known as the *confidence interval;* we discuss the calculation of confidence intervals in Chapter 19.

It is important that you recognize the difference between "random" in a scientific sense and "random" in an everyday sense. For many people, a "random sample" would simply mean walking down the street and passing out surveys to people they don't know, or going through the telephone book, haphazardly calling numbers "at random." To a researcher, however, a "random sample" is one in which the particular population elements are selected by some objective process outside the control of the researcher. Believe it or not, researchers have been known to roll dice, flip coins, or draw numbers out of a hat in order to create the randomness necessary with probabilistic samples. For simple random samples, in which each sampling element is selected randomly, random number tables generated by computers were often used in the past by researchers to select the sample. Now it is more likely that the computer will use an internal random-number generator to select the simple random sample.

The ability to draw a simple random sample depends to a great extent on the availability of a good sampling frame. For some populations, this is no problem—for example, imagine that you needed to conduct a study among *Fortune* magazine's list of the 500 largest corporations in the United States. The list is readily available, and a simple random sample of these firms could easily be selected. For many other target populations of interest (for example, all families living in a particular city), a list of population elements simply doesn't exist, and researchers often resort to other sampling schemes.

If the population is moderate to large in size, you'll want the computer to randomly select the sample from the sampling frame, so having a digital version of the sampling frame is important. Fortunately, this is commonly the case. If a computer file containing the sampling frame is not available, researchers again will likely want to utilize another form of probabilistic sampling plan.

Systematic Samples

Suppose that you were asked to conduct telephone interviews with 250 college students at a particular school and that the university published a directory that contained the names and telephone numbers of all 5,000 of its students. If you had access to a computer file containing the information, it would be a relatively easy matter to draw a simple random sample from the list. If you don't have such a computer file, however, drawing a simple random sample isn't so simple. It is difficult to randomly select each sample member.

systematic sample
A probability sampling plan in which every *k*th element in the population is selected for the sample pool after a random start.

A systematic sample offers an easy, but very effective, solution. With a systematic sample, researchers randomly select the first population element to be included in the sample and then select every *k*th element following it in the sampling frame. In our example, let's assume for a moment that we'll be able to interview all 250 college students who are selected for the sample. We'll end up interviewing one out of every 20 students on campus (5,000/250 = 20). So, you would randomly select one of the first 20 names in the student directory, then count down 20 names on the list and select that name, count down 20 more names and select that name, and so on, until you have gone through the entire directory. (It may not sound like it, but this is *much* easier than trying to randomly select each member of the sample by hand.)

So what makes this a probabilistic sampling plan? It's because the first element is randomly selected, and every other element selected for the sample is a function of the first element, which makes them all randomly selected, in effect. And if there is any sort of natural order within the sampling frame, using a systematic sample may well produce a sample that is more representative of the larger population than would a simple random sample. Imagine, for example, that our student directory groups the students by class; All first-year students are presented in the first section, in alphabetical order, followed by sophomores, juniors, and seniors. A systematic sample will necessarily produce a sample with the right proportions of student names drawn from each class; that won't necessarily happen with a simple random sample.

Calculating the **sampling interval** (that is, k, the number of names to count when selecting the sample members) is easy—sort of. In general, we simply divide the number of population elements in the sampling frame by the number of elements that we need to draw to obtain the sample size we want. In the example above, $k = 5,000/250$, so our sampling interval was 20. Here's where it gets a little tricky, though. Remember how we assumed that we could conduct telephone interviews with all 250 students selected for the sample? For lots of reasons, it almost never works out that way.

sampling interval
The number of population elements to count (k) when selecting the sample members in a systematic sample.

As we discuss in the following chapter, it's a relatively straightforward matter for researchers to decide how large or small a sample they need in a given research situation. Someone has decided that in the current situation a sample size of 250 students is sufficient. If we select only 250 students for our sample, however, it is almost a certainty that we'll end up with fewer than 250 respondents—maybe a whole lot fewer. Why? Some people won't be home to answer their telephones, even if we try multiple times to reach them. Others will have changed telephone numbers since the directory was published. And some people will refuse to answer our questions because they are too busy or just don't care to help. As a result, in almost all cases we need to start with a larger number of population elements in our initial sample pool in order to end up with the desired sample size. We refer to the total number of elements to be selected for inclusion in the initial sample pool as **total sampling elements (TSE)**.

total sampling elements (TSE)
The number of population elements that must be drawn from the population and included in the initial sample pool in order to end up with the desired sample size.

The notion of TSE is general and applies to any type of sample, not just systematic samples. Anytime it is necessary to select a larger initial sample in order to reach the necessary sample size, the calculation of TSE becomes important. Calculating TSE typically requires making predictions about the proportion of sample elements that (a) have incorrect contact information (telephone number, e-mail address, or mailing address); (b) are ineligible because they don't meet criteria for inclusion in the sample; (c) refuse to participate; and (d) cannot be contacted, even after multiple tries.

The formula for TSE looks like this:

$$\text{total sampling elements}(TSE) = \frac{\text{sample size}}{(1 - BCI)(1 - I)(1 - R)(1 - NC)}$$

where BCI = estimated proportion of bad contact information (wrong telephone numbers, mailing or e-mail addresses), I = estimated proportion of ineligible elements in the sampling frame (that is, people or entities that don't meet the criteria to be population members but were included in the sampling frame), R = estimated proportion of refusals, and NC = estimated proportion of elements that cannot be contacted after repeated attempts. The challenge, as you might imagine, is to estimate in advance the proportions needed for the formula in a given situation.

Returning to the current problem, we need a sample of 250 respondents from the 5,000 students in the directory. Even if the directory is updated annually, we should assume that some of the telephone numbers won't be working, because some people may have left school and others will have changed telephone numbers. Let's assume that percentage is 15%; thus, $BCI = 0.15$. Because some of the people who have left school, and therefore are no longer eligible to be included in the population, might still have working telephone numbers, we also need to include an ineligibility proportion. That proportion is likely to be low, however, so we'll set it at 2% ($I = 0.02$). Refusal rates aren't typically all that high, but we want to be conservative to ensure that we end up with at least 250 respondents, so we'll set the refusal rate at 20% ($R = 0.20$). Finally, and this is often the biggest issue of all—we'll assume that we won't be able to reach 30% of the people selected for the sample, even after trying three to four times at different times of the day and evening; thus, $NC = 0.30$. Putting it all together, we need to draw a total of 536 students from the population in order to obtain a sample size of 250:

$$\text{total sampling elements (TSE)} = \frac{250}{(1 - 0.15)(1 - 0.02)(1 - 0.20)(1 - 0.30)} = 536$$

Once we know how many elements we need to draw from the population, it's a simple matter to determine the sampling interval:

$$\text{sampling interval} = \frac{\text{number of elements in the sampling frame}}{\text{total sampling elements}} = \frac{5,000}{536} = 9.3$$

To draw the sample, you would randomly select one of the first 9 names in the directory, perhaps using a random-number generator on a computer or calculator or even something as straightforward as pulling a number out of a hat. Once that name is selected, every ninth name following the starting name will be drawn. Since drawing every ninth name will result in a list of 556 students instead of the 536 that you want, you might choose to count down 9 names to get the second sample element, another 9 names to get the third sample element—and then count down 10 names to get the fourth. It doesn't matter, provided that you follow the same pattern throughout the whole sampling frame (that is, down 9, down 9, down 10, down 9, down 9, down 10, and so on); each name after the first is still a function of the position of the randomly selected first sample element.

Learning Objective

7. Specify the two procedures that distinguish a stratified sample.

Stratified Samples

Our goal in drawing a probabilistic sample from a population is to describe the population's characteristics, or parameters, based on statistics calculated from the sample. Stratified samples sometimes allow researchers to do this more efficiently and/or more effectively. A stratified sample is a probability sample in which (1) the population is divided into mutually exclusive and exhaustive subgroups (that is, each population element fits into one—and only one—subgroup) and (2) samples are chosen from each of the subgroups.

stratified sample
A probability sample in which (1) the population is divided into mutually exclusive and exhaustive subsets, and (2) a simple random sample of elements is chosen independently from each group or subset.

Suppose that there are two major universities in your particular geographic region. The schools compete in many different areas: for budget dollars from the state, for students, for recognition, for victories across a host of different sporting events. Imagine that you have been hired to measure the image of one of the schools among residents in the region who are 18 years or older and have lived in the region for at least one year. You have developed a telephone survey with appropriate measures and now must draw a sample. One possibility for accomplishing this is to draw a simple random sample from the population of residents, collect the necessary data, and calculate sample statistics. In our randomly chosen set of respondents, we'll have some who went to the school in question, some who went to the rival school, and some who didn't go to college at all or went somewhere else. As a result, we'll probably have some who hold very positive perceptions of the school, some who are likely to hold more negative perceptions of the school, and others who fall in between the extremes, respectively. In short, the variation in the measured image of the school is likely to be quite large.

As we'll see in the following chapter, as the variance of a population parameter gets larger, the required sample size for reaching a desired margin of sampling error also gets larger, all else equal. In our case, to get the same level of precision, we'll have to include more residents in our sample than we would if the variation in the image of the university was smaller. A stratified sample, however, provides a way for us to achieve the same precision in our results with fewer respondents, which can reduce costs significantly.

Learning Objective

8. Cite two reasons researchers might choose to use a stratified sample rather than a simple random sample.

Here's how it works. We would first divide the population (as defined above) into three groups: (1) those who attended the target university, (2) those who attended the rival university, and (3) those who attended neither university. It is reasonable to expect that the vast majority of those who attended the target university will hold positive perceptions of the university and that the variation in their responses will be quite low (that is, everybody feels about the same way about the school). Similarly, it is reasonable to expect that the vast majority of those who attended the rival university will hold more negative perceptions of the target university and that variation in their scores will also be reasonably low (again, there would be little difference in the scores within the group). Scores within the third group might be all over the place, which would make the variation within that group potentially much larger. This is the important part: Because the variances of image scores within the first two groups are much lower than for the overall population, the necessary sample size within those groups goes down considerably, making the total sample size decrease, and reducing costs.

If you are having trouble grasping why this works, think about it like this: If you had a room full of people who absolutely loved the target university across every dimension and each would provide the highest image score possible if asked, how many people in the room would you need to talk with to understand how everyone in the room felt about the school? *Only one*. That's right, you'd only need feedback from one person in that group. Stratified samples work on the principle of building subgroups of population elements that are as similar as possible within the groups on the attribute being assessed so that efficiencies can be gained.

There's one other key reason that stratified samples might be used, and this has more to do with effectiveness than efficiency. Sometimes it is necessary to work with stratified samples as a means of ensuring that particular categories of respondents are included in the final sample. Suppose, for example, that a manufacturer of diamond rings wants to conduct a study of sales of the product by social class. Unless special precautions are taken, it is likely that the upper class—which represents only about 3% of the total population—will not be represented at all, or will be represented by too few cases. Yet this may be an extremely important segment to the ring manufacturer. It is often true in marketing that a small subset of the population of interest will account for a large proportion of the behavior of interest—for example, consumption of the product. It then becomes critical that this subgroup be adequately represented in the sample. Stratified sampling is one way of ensuring adequate representation from each subgroup of interest. This is a common approach and is especially useful when an important subgroup makes up only a small portion of the population.

Learning Objective

9. Explain the difference between a proportionate stratified sample and a disproportionate stratified sample.

Proportionate and Disproportionate Stratified Samples With a *proportionate stratified sample*, the number of observations in the total sample is allocated among the subgroups in proportion to the *relative* number of elements in each subgroup in the population. If a subgroup contained one-fifth of all the population elements, then one-fifth of the overall sample should come from that subgroup, and so on. With a *disproportionate stratified sample*, however, greater efficiencies with stratified samples can be achieved. As noted, smaller samples are needed from subgroups with less variation on the parameter being estimated, and costs are lowered because the overall sample size is reduced. On the other hand, if the budget allows for, say, 1,000 telephone interviews, and the researcher wants to use all possible respondents, subgroups with greater variation can be sampled more than proportionately to their relative size, which increases the precision of results for that subgroup and for the overall sample. Thus, the efficiencies associated with stratified samples can be taken in the form of reduced costs or increased precision in the results. Either way, the outcomes can be better with a stratified sample than with a simple random sample. It is important to note, however, that these efficiencies are possible only when meaningful subgroups can be identified and when variability is substantially reduced within one or more subgroups.

Learning Objective

10. List the steps followed in drawing a cluster sample.

Cluster Samples

Cluster samples are another probability sampling technique often used by researchers. Cluster sampling is similar to stratified sampling in that the population is divided into mutually exclusive and exhaustive subgroups, but the similarities stop there. With cluster sampling, one or more subgroups are randomly selected, and either all the elements included in those subgroups are selected for the sample (*one-stage cluster sampling*) or a sample of elements is selected probabilistically from the randomly selected subgroups (*two-stage cluster sampling*). Clusters can be set up based on a variety of approaches as illustrated in Exhibit 15.2.

cluster sample
A probability sampling plan in which (1) the parent population is divided into mutually exclusive and exhaustive subsets and (2) a random sample of one or more subsets (clusters) is selected.

With stratified sampling, a sample of elements is selected from each subgroup. Not so with cluster sampling. This has important implications for the composition of the subgroups, or clusters. Since only some clusters will be randomly selected, it is important that the full range of the key variables being studied be represented in each cluster. Each cluster should reflect the diversity of the whole population. The goal with cluster sampling is thus to have clusters that are as heterogeneous as possible on the key issues. That way, no matter which cluster(s) are randomly selected, the full range is represented. (Recall that we wanted the subgroups to be as homogeneous within each subgroup as possible with stratified sampling.)

Exhibit 15.2 Possible Clusters to Use to Sample Various Types of Population Elements

Population Elements	Possible Clusters
College students	Colleges
Elementary school students	Schools
Manufacturing firms	Counties
	Localities
Airline travelers	Airports
	Planes
Hospital patients	Hospitals
Government workers	Government buildings

Systematic sampling, which we discussed earlier, is technically a form of cluster sampling. Once the sampling interval (k) is calculated, the population is effectively broken into k clusters, and one of these clusters is selected when the initial element is randomly selected. For example, if the sampling interval $k = 2$, we've split the population into two clusters (that is, cluster 1 = first, third, fifth, seventh, etc., name in the list; cluster 2 = second, fourth, sixth, eighth, etc., name in the list), and we'll end up with one or the other cluster based on whether the first or second name is randomly selected, perhaps by a coin flip.

Area Samples In every probability sampling plan discussed so far, we need a list of population elements in order to draw the sample. What can a researcher do when such a list isn't available? What about when a sampling frame is available, but personal interviews are to be conducted and the costs of traveling from one interview to the next would be prohibitive? In both of these situations, an area sample might be the best choice.

Suppose, for example, that you needed to do a series of 1,000 personal interviews with the residents of a town or city near you. Directories of all those living in the city at that particular moment in time simply do not exist: People move, others die, new households are constantly being formed. Even though accurate lists of families most likely won't be available, you will probably have access to a tool essential for area sampling, a map of the city. An area sample is a form of cluster sample in which geographic areas (city blocks, neighborhoods, housing additions, etc.) serve as primary sampling units. In our situation, we would begin by identifying on the map the geographic areas in which families reside, breaking them into a number of clusters. It is best if the clusters contain approximately the same number of families, but it doesn't always work out that way. Once the clusters (again, geographic areas) have been identified, one or more areas are randomly selected.

area sample
A form of cluster sampling in which areas (for example, census tracts, blocks) serve as the primary sampling units. The population is divided into mutually exclusive and exhaustive areas using maps, and a random sample of areas is selected.

With a *one-stage area sample*, researchers would attempt to contact each family living in the selected areas, perhaps by knocking on doors or by collecting addresses and attempting to match names to the addresses for initial contact by telephone or mail. With *two-stage area sampling*, researchers draw a probabilistic sample from each of the randomly selected areas.

How well do area samples meet the criterion of having heterogeneity on the important parameters within each area cluster? Often, not too well. People tend to live near other people who are similar to themselves, which means that any particular area in a city probably won't represent the whole city very well. This problem can be addressed somewhat by drawing multiple clusters/areas, but it doesn't make the problem go away. Still, the cost efficiencies of working with only a few clusters rather than a simple random sample of the whole population may outweigh representativeness issues. As with many other issues in marketing research, the key is to use common sense about the best approach in a specific situation.

COMBINING SAMPLE TYPES

As you can probably tell by now, sample design is a very detailed subject. Our discussion has concentrated on only the basic types of samples. You should be aware that the basic types can be, and are, combined in large-scale field studies to produce some very complex designs. For example, it is possible to have several levels of stratification—such as by geographic area and density of population—precede several stages of cluster sampling. Research Window 15.2 discusses how Nielsen uses multiple sampling approaches in its people meter research to produce ratings for television programs.

You won't be a sampling expert based on the limited information we've provided in this chapter, but you should have a basic understanding of some key issues and the basic types of sampling plans.