

CHAPTER 5

Sampling

Sample Planning

Define Sample Components and the Population
Evaluate Generalizability
Assess the Diversity of the Population
Consider a Census

Sampling Methods

Probability Sampling Methods
Simple Random Sampling
Systematic Random Sampling
Stratified Random Sampling
Cluster Sampling
Probability Sampling Methods Compared

Nonprobability Sampling Methods

Availability Sampling
Quota Sampling
Purposive Sampling
Snowball Sampling
Lessons About Sample Quality
Generalizability in Qualitative Research

Sampling Distributions

Estimating Sampling Error
Determining Sample Size

Conclusions

A common technique in journalism is to put a “human face” on a story. For instance, a *Boston Globe* reporter (Abel 2008) interviewed a participant for a story about a housing program for chronically homeless people. “Burt” had worked as a welder, but alcoholism and both physical and mental health problems interfered with his plans. By the time he was 60, Burt had spent many years on the streets. Fortunately, he obtained an independent apartment through a new Massachusetts program, but even then “the lure of booze and friends from the street was strong” (Abel 2008:A14).

It is a sad story with an all-too-uncommon happy—although uncertain—ending. Together with one other such story and comments by several service staff, the article provides a persuasive rationale for the new housing program. However, we don’t know whether the two participants interviewed for the story are like most program participants, most homeless persons in Boston, or most homeless persons throughout the United States—or whether they

are just two people who caught the eye of this one reporter. In other words, we don't know how generalizable their stories are, and if we don't have confidence in generalizability, then the validity of this account of how the program participants became homeless is suspect. Because we don't know whether their situation is widely shared or unique, we cannot really judge what the account tells us about the social world.

In this chapter, you will learn about sampling methods, the procedures that primarily determine the generalizability of research findings. I first review the rationale for using sampling in social research and consider two circumstances when sampling is not necessary. The chapter then turns to specific sampling methods and when they are most appropriate, using examples from research on homelessness. This section is followed by a section on sampling distributions, which introduces you to the logic of statistical inference—that is, how to determine how likely it is that our sample statistics represent the population from which the sample was drawn. By the chapter's end, you should understand which questions you need to ask to evaluate the generalizability of a study as well as what choices you need to make when designing a sampling strategy. You should also realize that it is just as important to select the “right” people or objects to study as it is to ask participants the right questions.

SAMPLE PLANNING

You have encountered the problem of generalizability in each of the studies you have read about in this book. For example, Keith Hampton and Barry Wellman (1999) discussed their findings in Netville as though they could be generalized to residents of other communities; Norman Nie and Lutz Erbring (2000) generalized their Internet survey findings to the entire American adult population, and the National Geographic Society (2000) Web survey findings were generalized to the entire world. Whether we are designing a sampling strategy or evaluating someone else's findings, we have to understand how and why researchers decide to sample and what the consequences of these decisions are for the generalizability of the study's findings.

Define Sample Components and the Population

Let's say that we are designing a survey about adult homeless persons in one city. We don't have the time or resources to study the entire adult **population** of the city, even though it consists of the set of individuals or other entities to which we wish to be able to generalize our findings. Even the city of Boston, which conducts an annual census of homeless persons, does not have the resources to actually survey the homeless persons they count. So instead, we resolve to study a **sample**, a subset of this population. The individual members of this sample are called **elements**, or elementary units.

In many studies, we sample directly from the elements in the population of interest. We may survey a sample of the

Population The entire set of individuals or other entities to which study findings are to be generalized.

Sample A subset of a population that is used to study the population as a whole.

Elements The individual members of the population whose characteristics are to be measured.

150 INVESTIGATING THE SOCIAL WORLD

Sampling frame A list of all elements or other units containing the elements in a population.

Enumeration units Units that contain one or more elements and that are listed in a sampling frame.

Sampling units Units listed at each stage of a multistage sampling design.

entire population of students in a school, based on a list obtained from the registrar's office. This list, from which the elements of the population are selected, is termed the **sampling frame**. The students who are selected and interviewed from that list are the elements.

In some studies, the entities that can be reached easily are not the same as the elements from which we want information, but they include those elements. For example, we may have a list of households but not a list of the entire population of a town, even though the adults are the elements that we actually want to sample. In this situation, we could draw a sample of households so that we can identify the adult individuals in these households. The households are termed **enumeration units**, and the adults in the households are the elements (Levy & Lemeshow 1999:13–14).

Sometimes, the individuals or other entities from which we collect information are not actually the elements in our study. For example, a researcher might sample schools for

a survey about educational practices and then interview a sample of teachers in each sampled school to obtain the information about educational practices. Both the schools and the teachers are termed **sampling units**, because we sample from both (Levy & Lemeshow 1999:22). The schools are selected in the first stage of the sample, so they are the *primary sampling units* (in this case, they are also the elements in the study). The teachers are *secondary sampling units* (but they are not elements, because they are used to provide information about the entire school) (see Exhibit 5.1).

It is important to know exactly what population a sample can represent when you select or evaluate sample components. In a survey of “adult Americans,” the general population may reasonably be construed as all residents of the United States who are at least 21 years old. But always be alert to ways in which the population may have been narrowed by the sample selection procedures. For example, perhaps only English-speaking residents of the United States were surveyed. The population for a study is the aggregation of elements that we actually focus on and sample from, not some larger aggregation that we really wish we could have studied.

Some populations, such as the homeless, are not identified by a simple criterion such as a geographic boundary or an organizational membership. Clear definition of such a population is difficult but quite necessary. Anyone should be able to determine just what population was actually studied. However, studies of homeless persons in the early 1980s “did not propose definitions, did not use screening questions to be sure that the people they interviewed were indeed homeless, and did not make major efforts to cover the universe of homeless people” (Burt 1996:15). (Perhaps just homeless persons in one shelter were studied.) The result was a “collection of studies that could not be compared” (Burt 1996:15). Several studies of homeless persons in urban areas addressed the problem by employing a more explicit definition of the population: “people who had no home or permanent place to stay of their own (meaning they rented or owned it themselves) and no regular arrangement to stay at someone else’s place” (Burt 1996:18).

Even this more explicit definition still leaves some questions unanswered: What is a “regular arrangement”? How permanent does a “permanent place” have to be? In a study of

EXHIBIT 5.1 Sample Components in a Two-Stage Study

homeless persons in Chicago, Michael Sosin, Paul Colson, and Susan Grossman (1988) answered these questions in their definition of the population of interest:

We define the homeless as: those current[ly] residing for at least one day but for less than fourteen with a friend or relative, not paying rent, and not sure that the length of stay will surpass fourteen days; those currently residing in a shelter, whether overnight or transitional; those currently without normal, acceptable shelter arrangements and thus sleeping on the street, in doorways, in abandoned buildings, in cars, in subway or bus stations, in alleys, and so forth; those residing in a treatment center for the indigent who have lived at the facility for less than 90 days and who claim that they have no place to go, when released. (p. 22)

152 INVESTIGATING THE SOCIAL WORLD

This definition reflects accurately Sosin et al.'s concept of homelessness and allows researchers in other locations or at other times to develop procedures for studying a comparable population. The more complete and explicit the definition is of the population from which a sample was selected, the more precise our generalizations can be.

Evaluate Generalizability

Once we have defined clearly the population from which we will sample, we need to determine the scope of the generalizations we will make from our sample. Do you recall from Chapter 2 the two different meanings of generalizability?

Can the findings from a sample of the population be generalized to the population from which the sample was selected? Did Nie and Erbring's (2000) findings apply to the United States, National Geographic's (2000) to the entire world, or Wechsler et al.'s (2000) study of binge drinking to all U.S. college students? This type of generalizability was defined as *sample generalizability* in Chapter 2.

Can the findings from a study of one population be generalized to another, somewhat different population? Are e-mail users in Netville similar to those in other Ontario suburbs? In other provinces? In the United States? Are students similar to full-time employees, housewives, or other groups in their drinking patterns? Do findings from a laboratory study about alcohol effects at a small northeastern U.S. college differ from those that would be obtained at a college in the Midwest? What is the generalizability of the results from a survey of homeless persons in one city? This type of generalizability question was defined as *cross-population generalizability* in Chapter 2.

This chapter focuses attention primarily on the problem of sample generalizability: Can findings from a sample be generalized to the population from which the sample was drawn? This is really the most basic question to ask about a sample, and social research methods provide many tools with which to address it.

Sample generalizability depends on sample quality, which is determined by the amount of **sampling error**—the difference between the characteristics of a sample and the characteristics of the population from which it was selected. The larger the sampling error, the less representative the sample—and thus the less generalizable the findings. To assess sample quality when you are planning or evaluating a study, ask yourself these questions:

- From what population were the cases selected?
- What method was used to select cases from this population?
- Do the cases that were studied represent, in the aggregate, the population from which they were selected?

But researchers often project their theories onto groups or populations much larger than, or simply different from, those they have actually studied. The population to which generalizations are made in this way can be termed the **target population**—a set of elements larger than or different from the population that was sampled and to which the researcher would like to generalize any study findings. When we generalize findings to target populations, we must be somewhat speculative. We must carefully consider the validity of claims that the findings can be applied to other groups, geographic areas, cultures, or times.

Because the validity of cross-population generalizations cannot be tested empirically, except by conducting more research in other settings, I will not focus much attention on this problem here. But I'll return to the problem of cross-population generalizability in Chapter 7, which addresses experimental research, and in Chapter 12, which discusses methods for studying different societies.

Assess the Diversity of the Population

Sampling is unnecessary if all the units in the population are identical. Physicists don't need to select a representative sample of atomic particles to learn about basic physical processes. They can study a single atomic particle because it is identical to every other particle of its type. Similarly, biologists don't need to sample a particular type of plant to determine whether a given chemical has toxic effects on that particular type. The idea is "If you've seen one, you've seen 'em all."

What about people? Certainly, all people are not identical (nor are other animals, in many respects). Nonetheless, if we are studying physical or psychological processes that are the same among all people, sampling is not needed to achieve generalizable findings. Psychologists and social psychologists often conduct experiments on college students to learn about processes that they think are identical across individuals. They believe that most people would have the same reactions as the college students if they experienced the same experimental conditions. Field researchers who observed group processes in a small community sometimes make the same assumption.

There is a potential problem with this assumption, however: There's no way to know for sure if the processes being studied are identical across all people. In fact, experiments can give different results depending on the type of people who are studied or the conditions for the experiment. Stanley Milgram's (1965) classic experiments on obedience to authority, which you studied in Chapter 3, illustrate this point very well. You remember that the original Milgram experiments tested the willingness of male volunteers in New Haven, Connecticut, to comply with the instructions of an authority figure to give "electric shocks" to someone else, even when these shocks seemed to harm the person receiving them. In most cases, the volunteers complied. Milgram concluded that people are very obedient to authority.

Were these results generalizable to all men, to men in the United States, or to men in New Haven? The initial experiment was repeated many times to assess the generalizability of the findings. Similar results were obtained in many replications of the Milgram experiments, that is, when the experimental conditions and subjects were similar to those studied by Milgram. Other studies showed that some groups were less likely to react so obediently. Given certain conditions, such as another "subject" in the room who refused to administer the shocks, subjects were likely to resist authority.

So what do the initial experimental results tell us about how people will react to an authoritarian movement in the real world, when conditions are not so carefully controlled? In the real social world, people may be less likely to react obediently as well. Other individuals may argue against obedience to a particular leader's commands, or people may see on TV the consequences of their actions. But alternatively, people in the real world may be even more

Sampling error Any difference between the characteristics of a sample and the characteristics of a population. The larger the sampling error, the less representative the sample.

Target population A set of elements larger than or different from the population sampled and to which the researcher would like to generalize study findings.

154 INVESTIGATING THE SOCIAL WORLD

Representative sample A sample that “looks like” the population from which it was selected in all respects that are potentially relevant to the study. The distribution of characteristics among the elements of a representative sample is the same as the distribution of those characteristics among the total population. In an unrepresentative sample, some characteristics are overrepresented or underrepresented.

obedient to authority than were the experimental subjects, for example, when they get swept up in mobs or are captivated by ideological fervor. Milgram’s initial research and the many replications of it give us great insight into human behavior, in part, because they help identify the types of people and conditions to which the initial findings (lack of resistance to authority) can be generalized. But generalizing the results of single experiments is always risky, because such research often studies a small number of people who are not selected to represent any particular population.

The main point is that social scientists rarely can skirt the problem of demonstrating the generalizability of their findings. If a small sample has been studied in an experiment or a field research project, the study should be replicated in different settings or, preferably, with a **representative sample** of the population to which generalizations are sought (see Exhibit 5.2). The social world and the people in it are just too diverse to be considered “identical units.” Social psychological experiments and small field studies have produced good social science, but they need to be replicated in other settings, with other subjects, to claim any generalizability. Even when we believe that we have uncovered basic social processes in a laboratory experiment or field observation, we should be very concerned with seeking confirmation in other samples and in other research.

Census Research in which information is obtained through the responses that all available members of an entire population give to questions.

Consider a Census

In some circumstances, it may be feasible to skirt the issue of generalizability by conducting a **census**—studying the entire population of interest—rather than drawing a sample. This is what the federal government tries to do every 10 years with the U.S. Census. Censuses

also include studies of all the employees (or students) in small organizations, studies comparing all 50 states, and studies of the entire population of a particular type of organization in some area. However, in comparison with the U.S. Census and similar efforts in other countries, states, and cities, the population that is studied in these other censuses is relatively small.

The reason that social scientists don’t often attempt to collect data from all the members of some large population is simply that doing so would be too expensive and time-consuming—and they can do almost as well with a sample. Some social scientists conduct research with data from the U.S. Census, but it’s the government that collects the data and it’s your tax dollars that pay for the effort. To conduct the 2000 Census, the Congress and the president allocated almost \$4.5 billion (Prewitt 2000), and the U.S. Bureau of the Census spent 12 years doing the planning (U.S. Bureau of the Census, 2000a). For the 2010 Census, the Census Bureau is already testing new approaches, including an Internet-based response option (U.S. Bureau of the Census 2003).

EXHIBIT 5.2 Representative and Unrepresentative Samples

Even if the population of interest for a survey is a small town of 20,000 or students in a university of 10,000, researchers will have to sample. The costs of surveying “just” thousands of individuals exceed by far the budgets for most research projects. In fact, not even the U.S. Bureau of the Census can afford to have everyone answer all the questions that should be covered in the census. So it draws a sample. Every household must complete a short version of the census (it had seven basic questions in 2000), and a sample consisting of one in six households must complete a long form (with 53 additional questions) (Rosenbaum 2000).

The fact is that it is hard to get people to complete a survey is another reason why survey research can be costly. Even the U.S. Bureau of the Census (1999) must make multiple efforts

156 INVESTIGATING THE SOCIAL WORLD

to increase the rate of response in spite of the federal law requiring all citizens to complete their census questionnaire. After the Census Bureau spent \$167 million on publicity (Forero 2000b), two-thirds of the population returned their census questionnaire through the mail, ending a three-decade decline (U.S. Bureau of the Census 2000e). However, half a million temporary workers and up to six follow-ups were required to contact the rest of the households that did not respond by mail (U.S. Bureau of the Census 2000b, 2000c). As the U.S. 2000 Census progressed, concerns arose about underrepresentation of minority groups (Kershaw 2000), impoverished cities (Zielbauer 2000), well-to-do individuals in gated communities and luxury buildings (Langford 2000), and even college students (Abel 2000), so the Bureau conducted an even more intensive sample survey to learn about the characteristics of those who had still not responded (Anderson & Fienberg 1999; U.S. Bureau of the Census, 2000d). The number of persons missed in the census was still estimated to be between 3.2 and 6.4 million (U.S. Bureau of the Census 2001), and controversy continued over underrepresentation of some groups (Armas 2002; Holmes 2001a).

The average survey project has far less legal and financial backing, and so an adequate census is not likely to be possible. Even in Russia, which spent almost \$200 million to survey its population of about 145 million, resource shortages after the collapse of the Soviet Union prevented an adequate census (Myers 2002). The census had to be postponed from 1999 to 2002 due to insufficient funds and had to rely on voluntary participation. In spite of an \$8 million advertising campaign, many residents in impoverished regions refused to take part (Tavernise 2002). In Vladivostok, “many residents, angry about a recent rise in electricity prices, refused to take part. Residents on Russian Island . . . boycotted to protest dilapidated roads” (Tavernise 2002:A13).

In most survey situations, it is much better to survey only a limited number from the total population so that there are more resources for follow-up procedures that can overcome reluctance or indifference about participation. (I will give more attention to the problem of nonresponse in Chapter 8.)

SAMPLING METHODS

We can now study more systematically the features of samples that make them more or less likely to represent the population from which they are selected. The most important

Probability sampling method A sampling method that relies on a random, or chance, selection method so that the probability of selection of population elements is known.

Nonprobability sampling method Sampling method in which the probability of selection of population elements is unknown.

distinction that needs to be made about the samples is whether they are based on a probability or a nonprobability sampling method. Sampling methods that allow us to know in advance how likely it is that any element of a population will be selected for the sample are termed **probability sampling methods**. Sampling methods that do not let us know in advance the likelihood of selecting each element are termed **nonprobability sampling methods**.

Probability sampling methods rely on a random, or chance, selection procedure, which is, in principle, the same as flipping a coin to decide which of two people “wins” and which one “loses.” Heads and tails are equally

likely to turn up in a coin toss, so both persons have an equal chance of winning. That chance, their **probability of selection**, is 1 out of 2, or .5.

Flipping a coin is a fair way to select one of two people because the selection process harbors no systematic bias. You might win or lose the coin toss, but you know that the outcome was due simply to chance, not to bias. For the same reason, a roll of a six-sided die is a fair way to choose one of six possible outcomes (the odds of selection are 1 out of 6, or .17). Dealing out a hand after shuffling a deck of cards is a fair way to allocate sets of cards in a poker game (the odds of each person getting a particular outcome, such as a full house or a flush, are the same). Similarly, state lotteries use a random process to select winning numbers. Thus, the odds of winning a lottery, the probability of selection, are known, even though they are very much smaller (perhaps 1 out of 1 million) than the odds of winning a coin toss.

There is a natural tendency to confuse the concept of **random sampling**, in which cases are selected only on the basis of chance, with a haphazard method of sampling. On first impression, “leaving things up to chance” seems to imply not exerting any control over the sampling method. But to ensure that nothing but chance influences the selection of cases, the researcher must proceed very methodically, leaving nothing to chance except the selection of the cases themselves. The researcher must follow carefully controlled procedures if a purely random process is to occur. In fact, when reading about sampling methods, do not assume that a random sample was obtained just because the researcher used a random selection method at some point in the sampling process. Look for those two particular problems; selecting elements from an incomplete list of the total population and failing to obtain an adequate response rate.

If the sampling frame is incomplete, a sample selected randomly from that list will not really be a random sample of the population. You should always consider the adequacy of the sampling frame. Even for a simple population such as a university’s student body, the registrar’s list is likely to be at least a bit out-of-date at any given time. For example, some students will have dropped out, but their status will not yet be officially recorded. Although you may judge the amount of error introduced in this particular situation to be negligible, the problems are greatly compounded for a larger population. The sampling frame for a city, state, or nation is always likely to be incomplete because of constant migration into and out of the area. Even unavoidable omissions from the sampling frame can bias a sample against particular groups within the population.

A very inclusive sampling frame may still yield systematic bias if many sample members cannot be contacted or refuse to participate. Nonresponse is a major hazard in survey research because **nonrespondents** are likely to differ systematically from those who take the time to participate. You should not assume that findings from a randomly selected sample will be

Probability of selection The likelihood that an element will be selected from the population for inclusion in the sample. In a census of all elements of a population, the probability that any particular element will be selected is 1.0. If half of the elements in the population are sampled on the basis of chance (say, by tossing a coin), the probability of selection for each element is one half, or .5. As the size of the sample as a proportion of the population decreases, so does the probability of selection.

Random sampling A method of sampling that relies on a random, or chance, selection method so that every element of the sampling frame has a known probability of being selected.

158 INVESTIGATING THE SOCIAL WORLD

Nonrespondents People or other entities who do not participate in a study although they are selected for the sample.

Systematic bias Overrepresentation or underrepresentation of some population characteristics in a sample due to the method used to select the sample. A sample shaped by systematic sampling error is a biased sample.

generalizable to the population from which the sample was selected if the rate of nonresponse is considerable (certainly not if it is much above 30%).

Probability Sampling Methods

Probability sampling methods are those in which the probability of selection is known and is not zero (so there is some chance of selecting each element). These methods randomly select elements and therefore have no **systematic bias**; nothing but chance determines which elements are included in the sample. This feature of probability samples makes them much more desirable than nonprobability samples, when the goal is to generalize to a larger population.

Although a random sample has no systematic bias, it will certainly have some sampling error due to chance. The probability of selecting a head is .5 in a single toss of a coin and in 20, 30, or however many tosses of a coin you like. But it is perfectly possible to toss a coin twice and get a head both times. The random “sample” of the two sides of the coin is selected in an unbiased fashion, but it still is unrepresentative. Imagine selecting randomly a sample of 10 people from a population comprising 50 men and 50 women. Just by chance, can’t you imagine finding that these 10 people include 7 women and only 3 men? Fortunately, we can determine mathematically the likely degree of sampling error in an estimate based on a random sample (as we’ll discuss later in this chapter)—assuming that the sample’s randomness has not been destroyed by a high rate of nonresponse or by poor control over the selection process.

In general, both the size of the sample and the homogeneity (sameness) of the population affect the degree of error due to chance; the proportion of the population that the sample represents does not. To elaborate,

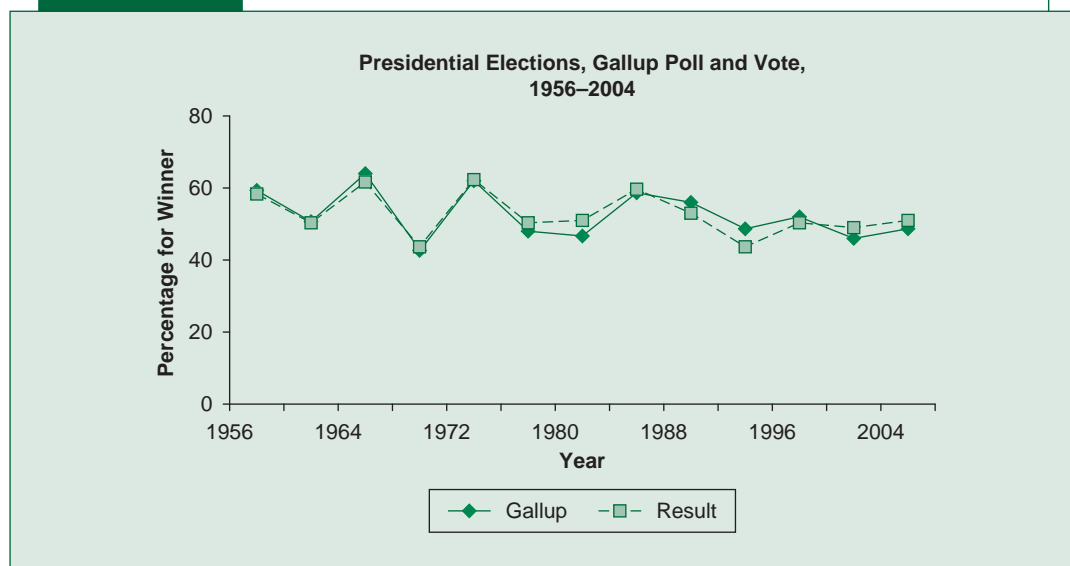
- *The larger the sample, the more confidence we can have in the sample’s representativeness.* If we randomly pick 5 people to represent the entire population of our city, our sample is unlikely to be very representative of the entire population in terms of age, gender, race, attitudes, and so on. But if we randomly pick 100 people, the odds of having a representative sample are much better; with a random sample of 1,000, the odds become very good indeed.
- *The more homogeneous the population, the more confidence we can have in the representativeness of a sample of any particular size.* Let’s say we plan to draw samples of 50 from each of two communities to estimate mean family income. One community is very diverse, with family incomes varying from \$12,000 to \$85,000. In the other, more homogeneous community, family incomes are concentrated in a narrow range, from \$41,000 to \$64,000. The estimated mean family income based on the sample from the homogeneous community is more likely to be representative than is the estimate based on the sample from the more heterogeneous community. With less variation to represent, fewer cases are needed to represent the homogeneous community.
- *The fraction of the total population that a sample contains does not affect the sample’s representativeness unless that fraction is large.* We can regard any sampling fraction

less than 2% with about the same degree of confidence (Sudman 1976:184). In fact, sample representativeness is not likely to increase much until the sampling fraction is quite a bit higher. Other things being equal, a sample of 1,000 from a population of 1 million (with a sampling fraction of 0.001, or 0.1%) is much better than a sample of 100 from a population of 10,000 (although the sampling fraction for this smaller sample is 0.01, or 1%, which is 10 times higher). The size of the samples is what makes representativeness more likely, not the proportion of the whole that the sample represents.

Polls to predict presidential election outcomes illustrate both the value of random sampling and the problems that it cannot overcome. In most presidential elections, pollsters have predicted accurately the outcomes of the actual votes by using random sampling and, these days, phone interviewing to learn for which candidate the likely voters intend to vote. Exhibit 5.3 shows how close these sample-based predictions have been in the last 13 contests. The exceptions were the 1980 and 1992 elections, when third-party candidates had an unpredicted effect. Otherwise, the small discrepancies between the votes predicted through random sampling and the actual votes can be attributed to random error.

The Gallup poll did quite well in predicting the result of the hotly contested 2000 presidential election. The final Gallup prediction was that George W. Bush would win with 48% (Al Gore was predicted to receive only 46%, while Green Party nominee Ralph Nader was predicted to secure 4%). Although the race turned out much closer, with Gore actually winning the popular vote (before losing in the electoral college), Gallup accurately noted that there appeared to have been a late-breaking trend in favor of Gore (Newport 2000). In 2004, the final Gallup prediction of 49% for Bush was within 2 percentage points of his winning total of 51% (actually, 50.77%); the “error” is partially due to the 1% of votes cast for third-party candidate Ralph Nader.

EXHIBIT 5.3 Presidential Election Outcomes: Predicted and Actual



160 INVESTIGATING THE SOCIAL WORLD

Nevertheless, election polls have produced some major errors in prediction. The reasons for these errors illustrate some of the ways in which unintentional systematic bias can influence sample results. In 1936, a *Literary Digest* poll predicted that Alfred M. Landon would defeat President Franklin Delano Roosevelt in a landslide, but instead Roosevelt took 63% of the popular vote. The problem? The *Digest* mailed out 10 million mock ballots to people listed in telephone directories, automobile registration records, voter lists, and so on. But in 1936, during the Great Depression, only relatively wealthy people had phones and cars, and they were more likely to be Republican. Furthermore, only 2,376,523 completed ballots were returned, and a response rate of only 24% leaves much room for error. Of course, this poll was not designed as a random sample, so the appearance of systematic bias is not surprising. Gallup was able to predict the 1936 election results accurately with a randomly selected sample of just 3,000 (Bainbridge 1989:43–44).

In 1948, pollsters mistakenly predicted that Thomas E. Dewey would beat Harry S. Truman, based on the random sampling method that George Gallup had used successfully since 1934. The problem? Pollsters stopped collecting data several weeks before the election, and in those weeks, many people changed their minds (Kenney, 1987). The sample was systematically biased by underrepresenting shifts in voter sentiment just before the election.

The fast-paced 2008 presidential primary elections were also challenging for the pollsters, primarily among Democratic Party voters. In the early New Hampshire primary, polls successfully predicted Republican John McCain's winning margin of 5.5% (the polls were off by only 0.2%, on average). However, all the polls predicted that Barack Obama would win the Democratic primary by a margin of about 8 percentage points, but he lost to Hillary Clinton by 12 points (47% to 35%). In a careful review of different explanations that have been proposed for that failure, the president of the Pew Research Center, Andrew Kohut (2008:A27), concluded that the problem was that voters who are poorer, less well educated, and white and who tend to refuse to respond to surveys tend to be less favorable to blacks than other voters. These voters, who were unrepresented in the polls, were more likely to favor Clinton over Obama.

Because they do not disproportionately exclude or include particular groups within the population, random samples that are successfully implemented avoid systematic bias in the selection process. However, when some types of people are more likely to refuse to participate in surveys or are less likely to be available for interviews, systematic bias can still creep into the sampling process. In addition, random error will still influence the specific results obtained from any random sample. Different types of random samples vary in their ability to minimize random error. The four most common methods for drawing random samples are simple random sampling, systematic random sampling, stratified random sampling, and cluster sampling.

Simple random sampling A method of sampling in which every sample element is selected only on the basis of chance, through a random process.

Random number table A table containing lists of numbers that are ordered solely on the basis of chance; it is used for drawing a random sample.

Simple Random Sampling

Simple random sampling requires some procedure that generates numbers or otherwise identifies cases strictly on the basis of chance. As you know, flipping a coin or rolling a die can be used to identify cases strictly on the basis of chance, but these procedures are not very efficient tools for drawing a sample. A **random number table**, such as the one in Appendix E, simplifies the process considerably. The

researcher numbers all the elements in the sampling frame and then uses a systematic procedure for picking corresponding numbers from the random number table. (Practice Exercise 1 at the end of this chapter explains the process step-by-step.) Alternatively, a researcher may use a lottery procedure. Each case number is written on a small card, and then the cards are mixed up and the sample is selected from the cards.

When a large sample must be generated, these procedures are very cumbersome. Fortunately, a computer program can easily generate a random sample of any size. The researcher must first number all the elements to be sampled (the sampling frame) and then run the computer program to generate a random selection of the numbers within the desired range. The elements represented by these numbers are the sample.

Organizations that conduct phone surveys often draw random samples using another automated procedure, called **random digit dialing**. A machine dials random numbers within the phone prefixes corresponding to the area in which the survey is to be conducted. Random digit dialing is particularly useful when a sampling frame is not available. The researcher simply replaces any inappropriate number (e.g., those that are no longer in service or that are for businesses) with the next randomly generated phone number.

The probability of selection in a true simple random sample is equal for each element. If a sample of 500 is selected from a population of 17,000 (i.e., a sampling frame of 17,000), then the probability of selection for each element is $500/17,000$, or .03. Every element has an equal chance of being selected, just like the odds in a toss of a coin ($1/2$) or a roll of a die ($1/6$). Thus, simple random sampling is an “equal probability of selection method,” or EPSEM.

Simple random sampling can be done either with or without replacement sampling. In **replacement sampling**, each element is returned to the sampling frame after it is selected so that it may be sampled again. In sampling without replacement, each element selected for the sample is then excluded from the sampling frame. In practice, it makes no difference whether sampled elements are replaced after selection as long as the population is large and the sample is to contain only a small fraction of the population. Random sampling with replacement is, in fact, rarely used.

In a study involving simple random sampling, Bruce Link and his associates (1996) used random digit dialing to contact adult household members in the continental United States for an investigation of public attitudes and beliefs about homeless people. Sixty-three percent of the potential interviewees responded. The sample actually obtained was not exactly comparable with the population sampled: Compared with U.S. Census figures, the sample overrepresented women, people aged 25 to 54, married people, and those with more than a high school education; it underrepresented Latinos.

How does this sample strike you? Let’s assess sample quality using the questions posed earlier in the chapter:

- *From what population were the cases selected?* There is a clearly defined population: the adult residents of the continental United States (who live in households with phones).

Random digit dialing The random dialing by a machine of numbers within designated phone prefixes, which creates a random sample for phone surveys.

Replacement sampling A method of sampling in which sample elements are returned to the sampling frame after being selected, so they may be sampled again. Random samples may be selected with or without replacement.

162 INVESTIGATING THE SOCIAL WORLD

- *What method was used to select cases from this population?* The case selection method is a random selection procedure, and there are no systematic biases in the sampling.
- *Do the cases that were studied represent, in the aggregate, the population from which they were selected?* The findings will very likely represent the population sampled because there were no biases in the sampling and a very large number of cases were selected. However, 37% of those selected for interviews could not be contacted or chose not to respond. This rate of nonresponse seems to create a small bias in the sample for several characteristics.

We must also consider the issue of cross-population generalizability: Do findings from this sample have implications for any larger group beyond the population from which the sample was selected? Because a representative sample of the entire U.S. adult population was drawn, this question has to do with cross-national generalizations. Link and his colleagues don't make any such generalizations. There's no telling what might occur in other countries with different histories of homelessness and different social policies.

Systematic random sampling A method of sampling in which sample elements are selected from a list or from sequential files.

Systematic Random Sampling

Systematic random sampling is a variant of simple random sampling. The first element is selected randomly from a list or from sequential files, and then every n th element is selected. This is a convenient method for drawing a random sample when the population elements

are arranged sequentially. It is particularly efficient when the elements are not actually printed (i.e., there is no sampling frame) but instead are represented by folders in filing cabinets.

Systematic random sampling requires the following three steps:

1. The total number of cases in the population is divided by the number of cases required for the sample. This division yields the sampling interval, the number of cases from one sampled case to another. If 50 cases are to be selected out of 1,000, the **sampling interval** is 20; every 20th case is selected.
2. A number from 1 to 20 (or whatever the sampling interval is) is selected randomly. This number identifies the first case to be sampled, counting from the first case on the list or in the files.
3. After the first case is selected, every n th case is selected for the sample, where n is the sampling interval. If the sampling interval is not a whole number, the size of the sampling interval is varied systematically to yield the proper number of cases for the sample. For example, if the sampling interval is 30.5, the sampling interval alternates between 30 and 31. In almost all sampling situations, systematic random sampling yields what is essentially a simple random sample.

Sampling interval The number of cases from one sampled case to another in a systematic random sample.

Periodicity A sequence of elements (in a list to be sampled) that varies in some regular, periodic pattern.

The exception is a situation in which the sequence of elements is affected by **periodicity**—that is, the sequence varies in some regular, periodic pattern. For example, the houses in a new development with the same number of

houses in each block (e.g., eight) may be listed by block, starting with the house in the northwest corner of each block and continuing clockwise. If the sampling interval is 8, the same as the periodic pattern, all the cases selected will be in the same position (see Exhibit 5.4). But in reality, periodicity and the sampling interval are rarely the same.

Stratified Random Sampling

Although all probability sampling methods use random sampling, some add steps to the sampling process to make sampling more efficient or easier. **Stratified random sampling** uses information known about the total population prior to sampling to make the sampling process more efficient. First, all elements in the population (i.e., in the sampling frame) are distinguished according to their value on some relevant characteristic. That characteristic forms the sampling strata. Next, elements are sampled randomly from within these strata. For example, race may be the basis for distinguishing individuals in some population of interest.

EXHIBIT 5.4 The Effect of Periodicity on Systematic Random Sampling



164 INVESTIGATING THE SOCIAL WORLD

Stratified random sampling A method of sampling in which sample elements are selected separately from population strata that are identified in advance by the researcher.

Proportionate stratified sampling Sampling method in which elements are selected from strata in exact proportion to their representation in the population.

Within each racial category, individuals are then sampled randomly. Of course, using this method requires more information prior to sampling than is the case with simple random sampling. It must be possible to categorize each element in one and only one stratum, and the size of each stratum in the population must be known.

This method is more efficient than drawing a simple random sample because it ensures appropriate representation of elements across strata. Imagine that you plan to draw a sample of 500 from an ethnically diverse neighborhood. The neighborhood population is 15% black, 10% Hispanic, 5% Asian, and 70% white. If you drew a simple random sample, you might end up with somewhat disproportionate numbers of each group. But if you created sampling strata based on race and ethnicity, you could randomly select cases from each

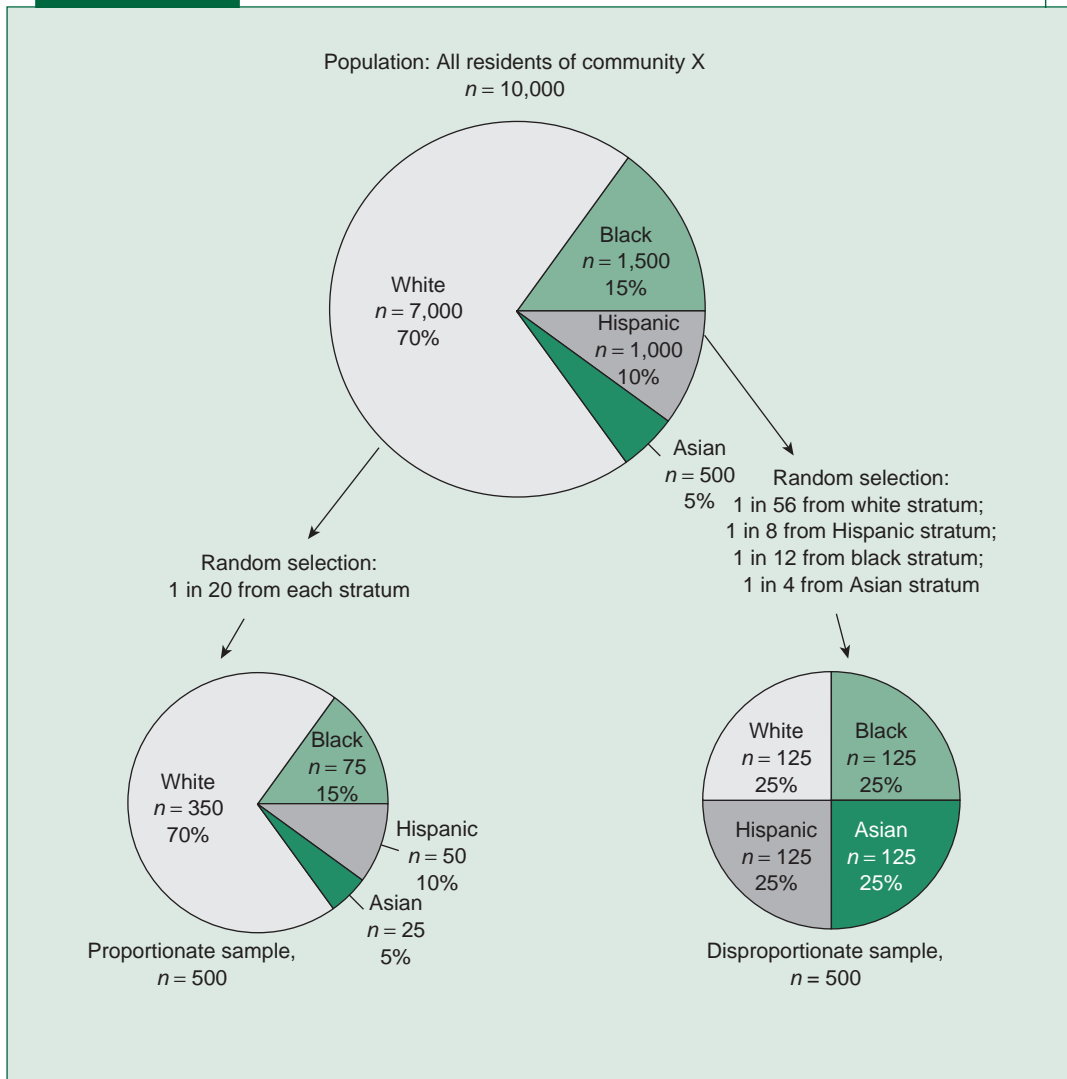
stratum: 75 blacks (15% of the sample), 50 Hispanics (10%), 25 Asians (5%), and 350 whites (70%). By using **proportionate stratified sampling**, you would eliminate any possibility of sampling error in the sample's distribution of ethnicity. Each stratum would be represented exactly in proportion to its size in the population from which the sample was drawn (see Exhibit 5.5).

This is the strategy used by Brenda Booth et al. (2002) in a study of homeless adults in two Los Angeles county sites with large homeless populations. Specifically, Booth et al. (2002:432) selected subjects at random from homeless shelters, meal facilities, and from literally homeless populations on the streets. Respondents were sampled proportionately to their numbers in the downtown and Westside areas, as determined by a one-night enumeration. They were also sampled proportionately to their distribution across three nested sampling strata: the population using shelter beds, the population using meal facilities, and the unsheltered population using neither.

In **disproportionate stratified sampling**, the proportion of each stratum that is included in the sample is intentionally varied from what it is in the population. In the case of the sample stratified by ethnicity, you might select equal numbers of cases from each racial or ethnic group: 125 blacks (25% of the sample), 125 Hispanics (25%), 125 Asians (25%), and 125 whites (25%). In this type of sample, the probability of selection of every case is known but unequal between strata. You know what the proportions are in the population, and so you can easily adjust your combined sample statistics to reflect these true proportions. For instance, if you want to combine the

Disproportionate stratified sampling Sampling in which elements are selected from strata in different proportions from those that appear in the population.

ethnic groups and estimate the average income of the total population, you would have to "weight" each case in the sample. The weight is a number you multiply by the value of each case based on the stratum it is in. For example, you would multiply the incomes of all blacks in the sample by 0.6 (75/125), the incomes of all Hispanics by 0.4 (50/125), and so on. Weighting in this way reduces the influence of the oversampled strata and increases the influence of the undersampled strata to what they would have been if pure probability sampling had been used.

EXHIBIT 5.5 Stratified Random Sampling


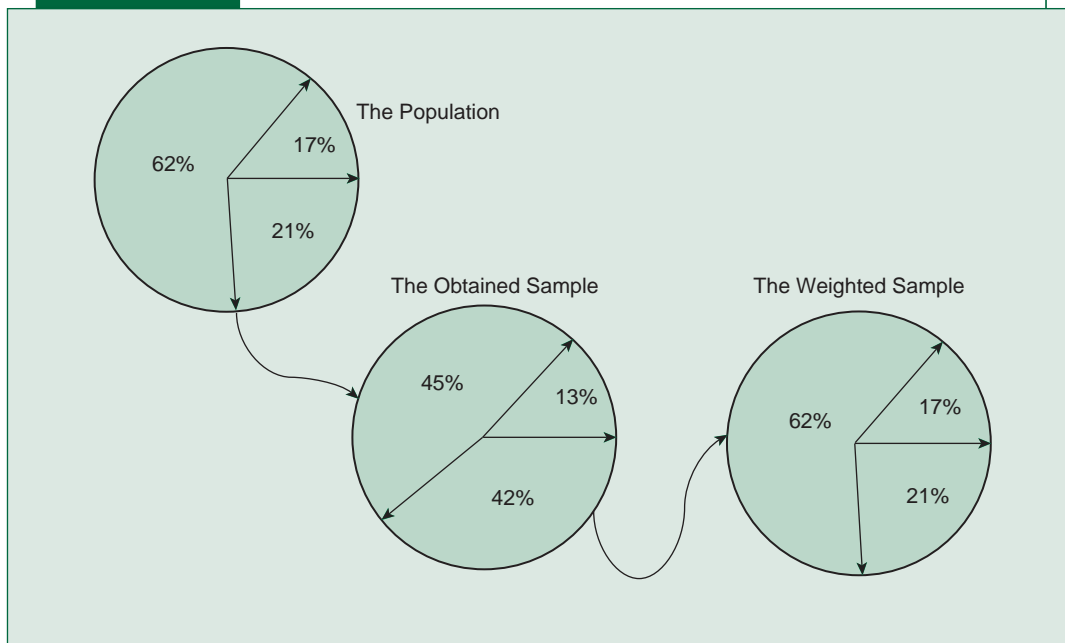
Booth et al. (2002:432) included one element of *disproportionate* random sampling in their otherwise proportionate random sampling strategy for homeless persons in Minneapolis: They oversampled women so that they comprised 26% of the sample, as compared with their actual percentage of 16% in the homeless population. Why would anyone select a sample that is so unrepresentative in the first place? The most common reason is to ensure that cases from smaller strata are included in the sample in sufficient numbers so as to allow separate statistical estimates and to facilitate comparisons between strata. Remember that one of the determinants

166 INVESTIGATING THE SOCIAL WORLD

of sample quality is sample size. The same is true for subgroups within samples. If a key concern in a research project is to describe and compare the incomes of people from different racial and ethnic groups, then it is important that the researchers base the mean income of each group on enough cases to be a valid representation. If few members of a particular minority group are in the population, they need to be oversampled. Such disproportionate sampling may also result in a more efficient sampling design if the costs of data collection differ markedly between the strata or if the variability (heterogeneity) of the strata differs.

Weighting is also sometimes used to reduce the lack of representativeness of a sample that occurs due to nonresponse. On finding that the obtained sample does not represent the population in terms of some known characteristics such as, perhaps, gender or education, the researcher weights the cases in the sample so that the sample has the same proportions of men and women, or high school graduates and college graduates, as the complete population (see Exhibit 5.6). Keep in mind, though, that this procedure does not solve the problems caused by an unrepresentative sample because you still don't know what the sample composition should have been in terms of the other variables in your study; all you have done is to reduce the sample's unrepresentativeness in terms of the variables used in weighting. This may, in turn, make it more likely that the sample is representative of the population in terms of other characteristics, but you don't really know.

EXHIBIT 5.6 Weighting an Obtained Sample to Match a Population Proportion



Cluster Sampling

Cluster sampling is useful when a sampling frame of elements is not available, as often is the case for large populations spread out across a wide geographic area or among many different organizations. A **cluster** is a naturally occurring, mixed aggregate of elements of the population, with each element appearing in one, and only one, cluster. Schools could serve as clusters for sampling students, blocks could serve as clusters for sampling city residents, counties could serve as clusters for sampling the general population, and businesses could serve as clusters for sampling employees.

Drawing a cluster sample is, at least, a two-stage procedure. First, the researcher draws a random sample of clusters. A list of clusters should be much easier to obtain than a list of all the individuals in each cluster in the population. Next, the researcher draws a random sample of elements within each selected cluster. Because only a fraction of the total clusters are involved, obtaining the sampling frame at this stage should be much easier.

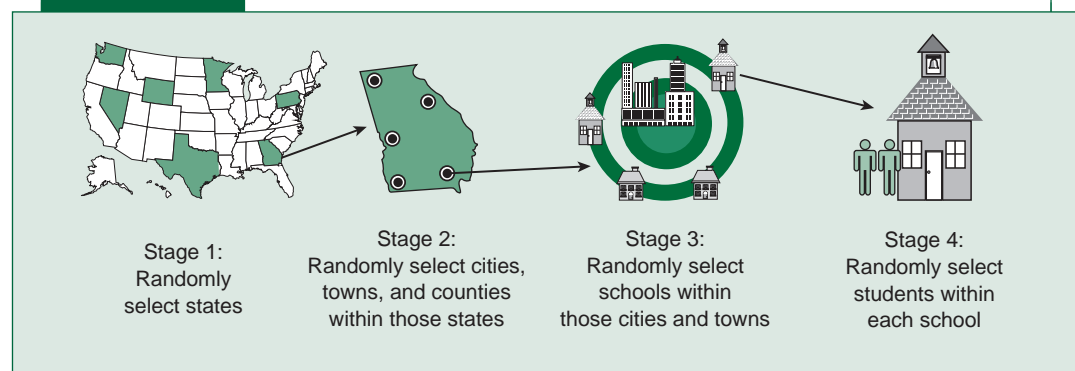
In a cluster sample of city residents, for example, blocks could be the first-stage clusters. A research assistant could walk around each selected block and record the addresses of all occupied dwelling units. Or, in a cluster sample of students, a researcher could contact the schools selected in the first stage and make arrangements with the registrar to obtain lists of students at each school. Cluster samples often involve multiple stages (see Exhibit 5.7), with clusters within clusters, as when a national sample of individuals might involve first sampling states, then geographic units within those states, then dwellings within those units, and finally, individuals within the dwellings. In multistage cluster sampling, the clusters at the first stage of sampling are termed the *primary sampling units* (Levy & Lemeshow 1999:228).

How many clusters should be selected, and how many individuals within each cluster should be selected? As a general rule, the sample will be more similar to the entire population

Cluster sampling Sampling in which elements are selected in two or more stages, with the first stage being the random selection of naturally occurring clusters and the last stage being the random selection of elements within clusters.

Cluster A naturally occurring, mixed aggregate of elements of the population.

EXHIBIT 5.7 Multistage Cluster Sampling



168 INVESTIGATING THE SOCIAL WORLD

if the researcher selects as many clusters as possible—even though this will mean the selection of fewer individuals within each cluster. Unfortunately, this strategy also maximizes the cost of the sample for studies using in-person interviews. The more clusters a researcher selects, the more time and money will have to be spent traveling to the different clusters to reach the individuals for interviews.

The calculation of how many clusters to sample and how many individuals are within the clusters is also affected by the degree of similarity of individuals within clusters: The more similar the individuals are within the clusters, the fewer the number of individuals needed to represent each cluster. So if you set out to draw a cluster sample, be sure to consider how similar individuals are within the clusters as well as how many clusters you can afford to include in your sample.

Cluster sampling is a very popular method among survey researchers, but it has one general drawback: Sampling error is greater in a cluster sample than in a simple random sample, because there are two steps involving random selection rather than just one. This sampling error increases as the number of clusters decreases, and it decreases as the homogeneity of cases per cluster increases. In sum, it's better to include as many clusters as possible in a sample, and it's more likely that a cluster sample will be representative of the population if cases are relatively similar within clusters.

Probability Sampling Methods Compared

Can you now see why researchers often prefer to draw a stratified random sample or a cluster sample rather than a simple random sample? Exhibit 5.8 should help you remember the key features of these different types of sample and to determine when each is most appropriate.

Many professionally designed surveys use combinations of clusters and stratified probability sampling methods. For example, Peter Rossi (1989) drew a disproportionate stratified cluster sample of shelter users for a homelessness study in Chicago (see Exhibit 5.9). The shelter sample was stratified by size, with smaller shelters having a smaller likelihood of selection than larger shelters. In fact, the larger shelters were all selected; they had a probability of selection of 1.0. Within the selected shelters, shelter users were then sampled using a systematic random selection procedure (except in the small shelters, in which all persons were interviewed). Homeless persons living on the streets were also sampled randomly. In the first stage, city blocks were classified into strata based on the likely concentration of homeless persons (estimated by several knowledgeable groups). Blocks were then picked randomly within these

EXHIBIT 5.8 Features of Probability Sampling Methods

<i>Feature</i>	<i>Simple</i>	<i>Systematic</i>	<i>Stratified</i>	<i>Cluster</i>
Unbiased selection of cases	Yes	Yes	Yes	Yes
Sampling frame required	Yes	No	Yes	No
Ensures representation of key strata	No	No	Yes	No
Uses natural grouping of cases	No	No	No	Yes
Reduces sampling costs	No	No	No	Yes
Sampling error compared with SRS	–	Same	Lower	Higher

strata and, on the survey night between 1 a.m. and 6 a.m., teams of interviewers screened each person found outside on that block for his or her homeless status. Persons identified as homeless were then interviewed (and given \$5 for their time). The rate of response for two different samples (fall and winter) in the shelters and on the streets was between 73% and 83%.

How would we evaluate the Chicago homeless sample, using the sample evaluation questions?

- *From what population were the cases selected?* The population was clearly defined for each cluster.
- *What method was used to select cases from this population?* The random selection method was carefully described.
- *Do the cases that were studied represent, in the aggregate, the population from which they were selected?* The unbiased selection procedures make us reasonably confident in the representativeness of the sample, although we know little about the nonrespondents and therefore may justifiably worry that some types of homeless persons were missed.

Cross-population generalization seems to be reasonable with this sample, because it seems likely that the findings reflect general processes involving homeless persons. Rossi (1989) clearly thought so, because his book's title refers to homelessness in America, not just in Chicago.

Nonprobability Sampling Methods

Nonprobability sampling methods are often used in qualitative research; they also are used in quantitative studies when researchers are unable to use probability selection methods. In qualitative research, a focus on one setting or a very small sample allows a more intensive portrait of activities and actors, but it also limits field researchers' ability to generalize and

EXHIBIT 5.9

Chicago Shelter Universe and Shelter Samples, Fall and Winter Surveys

A. Shelter Universe and Samples

	<i>Fall</i>	<i>Winter</i>
Eligible shelters in universe	28	45
Universe bed capacities	1,573	2,001
Shelters drawn in sample	22	27

B. Details of Winter Shelter Sample

<i>Shelter Size Classification</i>	<i>Number in Universe</i>	<i>Number in Sample</i>	<i>Occupant Sampling Ratio</i>
Large (37 or more beds)	17	17	0.25
Medium (18–33 beds)	12	6	0.50
Small (under 18 beds)	16	4	1.00

Note: Shelters were drawn with probabilities proportionate to size, with residents sampled disproportionately within shelters to form a self-weighting sample. Sampling ratios for the Phase 2 sample are given in Panel B.

170 INVESTIGATING THE SOCIAL WORLD

lowers the confidence that others can place in these generalizations. The use of nonprobability sampling methods in quantitative research too often reflects a lack of concern with generalizability or a lack of understanding of the importance of probability-based sampling.

There are four common nonprobability sampling methods: availability sampling, quota sampling, purposive sampling, and snowball sampling. Because these methods do not use a random selection procedure, we cannot expect a sample selected with any of these methods to yield a representative sample. They should not be used in quantitative studies if a probability-based method is feasible. Nonetheless, these methods are useful when random sampling is not possible, when a research question calls for an intensive investigation of a small population, or when a researcher is performing a preliminary, exploratory study.

Availability Sampling

Elements are selected for **availability sampling** because they're available or easy to find. Thus, this sampling method is also known as a haphazard, accidental, or convenience sample.

Availability sampling Sampling in which elements are selected on the basis of convenience.

There are many ways to select elements for an availability sample: Standing on street corners and talking to whoever walks by, asking questions of employees who have time to talk when they pick up their paycheck at a personnel office, or approaching particular individuals at opportune times while observing activities in a social setting. You

may find yourself interviewing available students at campus hangouts as part of a course assignment. To study sexual risk-taking among homeless youth in Minneapolis, Linda Halcón and Alan Lifson (2004:73) hired very experienced street youth outreach workers who approached youth known or suspected to be homeless and asked if they would be willing to take part in a 20- to 30-minute interview.

The interviewers then conducted the 44-question interview, after which they gave respondents some risk reduction and referral information and a \$20 voucher.

A participant observation study of a group may require no more sophisticated approach. When Philippe Bourgois, Mark Lettiere, and James Quesada (1997) studied homeless heroin addicts in San Francisco, they immersed themselves in a community of addicts living in a public park. These addicts became the availability sample.

An availability sample is often appropriate in social research—for example, when a field researcher is exploring a new setting and trying to get some sense of the prevailing attitudes or when a survey researcher conducts a preliminary test of a new set of questions.

Now I'd like you to use the sample evaluation questions to evaluate person-in-the-street interviews of the homeless. If your answers are something like "The population was unknown," "The method for selecting cases was haphazard," and "The cases studied do not represent the population," you're right! There is no clearly definable population from which the respondents were drawn, and no systematic technique was used to select the respondents. There certainly is not much likelihood that the interviewees represent the distribution of sentiment among homeless persons in the Boston area or of welfare mothers or of impoverished rural migrants or of whatever we imagine the relevant population is.

In a similar vein, perhaps person-in-the street comments to news reporters suggest something about what homeless persons think, or maybe they don't; we can't really be sure. But let's give reporters their due: If they just want to have a few quotes to make their story more

appealing, nothing is wrong with their sampling method. However, their approach gives us no basis for thinking that we have an overview of the community sentiment. The people who happen to be available in any situation are unlikely to be just like those who are unavailable. We can't be at all certain that what we learn can be generalized with any confidence to a larger population of concern.

Availability sampling often masquerades as a more rigorous form of research. Popular magazines periodically survey their readers by printing a questionnaire for readers to fill out and mail in. A follow-up article then appears in the magazine under a title such as "What You Think About Intimacy in Marriage." If the magazine's circulation is large, a large sample can be achieved in this way. The problem is that usually only a tiny fraction of readers return the questionnaire, and these respondents are probably unlike other readers who did not have the interest or time to participate. So the survey is based on an availability sample. Even though the follow-up article may be interesting, we have no basis for thinking that the results describe the readership as a whole—much less the population at large.

Do you see now why availability sampling differs so much from random sampling methods, which require that "nothing but chance" affects the actual selection of cases? What makes availability sampling "haphazard" is precisely that a great many things other than chance can affect the selection of cases, ranging from the prejudices of the research staff to the work schedules of potential respondents. To truly leave the selection of cases up to chance, we have to design the selection process very carefully so that other factors are not influential. There's nothing "haphazard" about selecting cases randomly.

Quota Sampling

Quota sampling is intended to overcome the most obvious flaw of availability sampling—that the sample will just consist of whoever or whatever is available, without any concern for its similarity to the population of interest. The distinguishing feature of a quota sample is that quotas are set to ensure that the sample represents certain characteristics in proportion to their prevalence in the population.

Suppose that you wish to sample adult residents of a town in a study of support for a tax increase to improve the town's schools. You know from the town's annual report what the proportions of town residents are in terms of gender, race, age, and number of children. You think that each of these characteristics might influence support for new school taxes, so you want to be sure that the sample includes men, women, whites, blacks, Hispanics, Asians, older people, younger people, big families, small families, and childless families in proportion to their numbers in the town population.

This is where quotas come in. Let's say that 48% of the town's adult residents are men and 52% are women, and that 60% are employed, 5% are unemployed, and 35% are out of the labor force. These percentages and the percentages corresponding to the other characteristics become the quotas for the sample. If you plan to include a total of 500 residents in your sample, 240 must be men (48% of 500), 260 must be women, 300 must be employed, and so on. You may even set more refined quotas, such as certain numbers of employed women, employed men, unemployed men, and so on. With the quota list in hand, you (or your research

Quota sampling A nonprobability sampling method in which elements are selected to ensure that the sample represents certain characteristics in proportion to their prevalence in the population.

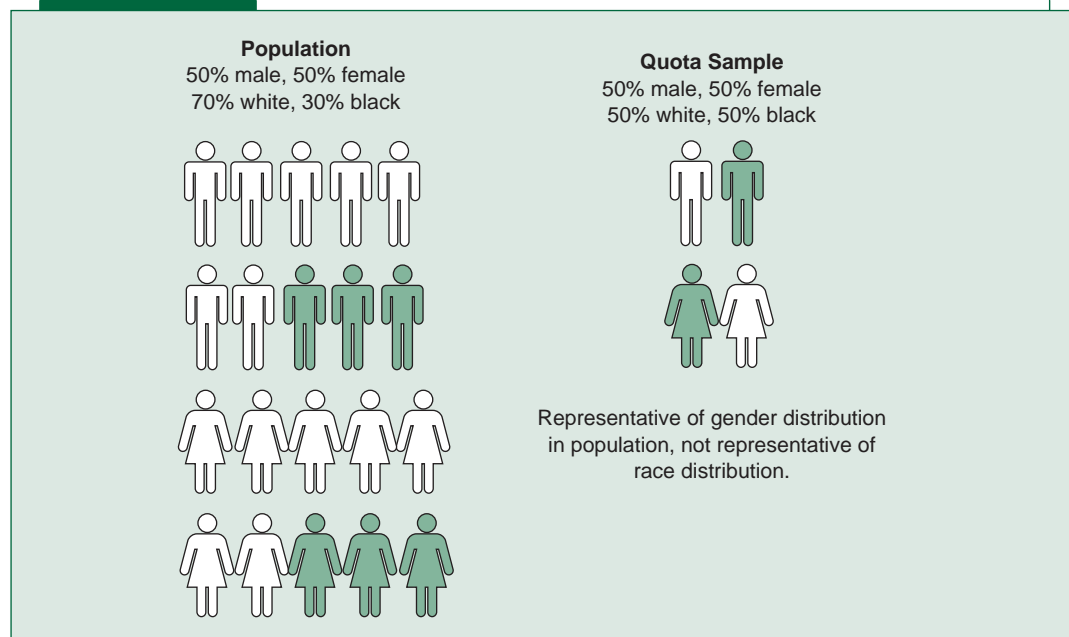
172 INVESTIGATING THE SOCIAL WORLD

staff) can now go out into the community looking for the right number of people in each quota category. You may go door to door, bar to bar, or just stand on a street corner until you have surveyed 240 men, 260 women, and so on.

The problem is that even when we know that a quota sample is representative of the particular characteristics for which quotas have been set, we have no way of knowing if the sample is representative in terms of any other characteristics. In Exhibit 5.10, for example, quotas have been set for gender only. Under these circumstances, it's no surprise that the sample is representative of the population only in terms of gender, not in terms of race. Interviewers are only human; they may avoid potential respondents with menacing dogs in the front yard, or they could seek out respondents who are physically attractive or who look like they'd be easy to interview. Realistically, researchers can set quotas for only a small fraction of the characteristics relevant to a study, so a quota sample is really not much better than an availability sample (although following careful, consistent procedures for selecting cases within the quota limits always helps).

This last point leads me to another limitation of quota sampling: You must know the characteristics of the entire population to set the right quotas. In most cases, researchers know what the population looks like in terms of no more than a few of the characteristics relevant to their concerns—and in some cases, they have no such information on the entire population.

If you're now feeling skeptical of quota sampling, you've gotten the drift of my remarks. Nonetheless, in some situations, establishing quotas can add rigor to sampling procedures. It's almost always better to maximize possibilities for comparison in research, and quota sampling techniques can help qualitative researchers do this. For instance, Doug Timmer, Stanley Eitzen, and Kathryn Talley (1993:7) interviewed homeless persons in several cities and other locations

EXHIBIT 5.10 Quota Sampling

for their book on the sources of homelessness. Persons who were available were interviewed, but the researchers paid some attention to generating a diverse sample. They interviewed 20 homeless men who lived on the streets without shelter and 20 mothers who were found in family shelters. About half of those the researchers selected in the street sample were black, and about half were white. Although the researchers did not use quotas to try to match the distribution of characteristics among the total homeless population, their informal quotas helped ensure some diversity in key characteristics.

Does quota sampling remind you of stratified sampling? It's easy to understand why, since they both select sample members based, in part, on the basis of one or more key characteristics. Exhibit 5.11 summarizes the differences between quota sampling and stratified random sampling. The key difference, of course, is quota sampling's lack of random selection.

Purposive Sampling

In **purposive sampling**, each sample element is selected for a purpose, usually because of the unique position of the sample elements. Purposive sampling may involve studying the entire population of some limited group (directors of shelters for homeless adults) or a subset of a population (mid-level managers with a reputation for efficiency). Or a purposive sample may be a "key informant survey," which targets individuals who are particularly knowledgeable about the issues under investigation.

Purposive sampling A nonprobability sampling method in which elements are selected for a purpose, usually because of their unique position.

Herbert Rubin and Irene Rubin (1995) suggest three guidelines for selecting informants when designing any purposive sampling strategy. Informants should be

- "knowledgeable about the cultural arena or situation or experience being studied,"
- "willing to talk," and
- "represent[ative of] the range of points of view." (p. 66)

In addition, Rubin and Rubin (1995) suggest continuing to select interviewees until you can pass two tests:

- *Completeness*: "What you hear provides an overall sense of the meaning of a concept, theme, or process." (p. 72)
- *Saturation*: "You gain confidence that you are learning little that is new from subsequent interview[s]." (p. 73)

EXHIBIT 5.11 Comparison of Stratified and Quota Sampling Methods

Feature	Stratified	Quota
Unbiased (random) selection of cases	Yes	No
Sampling frame required	Yes	No
Ensures representation of key strata	Yes	Yes

174 INVESTIGATING THE SOCIAL WORLD

Adhering to these guidelines will help ensure that a purposive sample adequately represents the setting or issues studied.

Of course, purposive sampling does not produce a sample that represents some larger population, but it can be exactly what is needed in a case study of an organization, community, or some other clearly defined and relatively limited group. In an intensive organizational case study, a purposive sample of organizational leaders might be complemented with a probability sample of organizational members. Before designing her probability samples of hospital patients and homeless persons, Dee Roth (1990:146–147) interviewed a purposive sample of 164 key informants from organizations that had contact with homeless people in each of the counties she studied.

Snowball Sampling

Snowball sampling is useful for hard-to-reach or hard-to-identify populations for which there is no sampling frame, but the members of which are somewhat interconnected (at least some

Snowball sampling A method of sampling in which sample elements are selected as they are identified by successive informants or interviewees.

members of the population know each other). It can be used to sample members of groups such as drug dealers, prostitutes, practicing criminals, participants in Alcoholics Anonymous groups, gang leaders, informal organizational leaders, and homeless persons. It may also be used for charting the relationships among members of some group (a sociometric study), for exploring the population of interest prior to developing a formal sampling plan, and for

developing what becomes a census of informal leaders of small organizations or communities. However, researchers using snowball sampling normally cannot be confident that their sample represents the total population of interest, so generalizations must be tentative.

Rob Rosenthal (1994) used snowball sampling to study homeless persons living in Santa Barbara, California:

I began this process by attending a meeting of homeless people I had heard about through my housing advocate contacts. . . . One homeless woman . . . invited me to . . . where she promised to introduce me around. Thus a process of snowballing began. I gained entree to a group through people I knew, came to know others, and through them gained entree to new circles. (pp. 178, 180)

One problem with this technique is that the initial contacts may shape the entire sample and foreclose access to some members of the population of interest:

Sat around with [my contact] at the Tree. Other people come by, are friendly, but some regulars, especially the tougher men, don't sit with her. Am I making a mistake by tying myself too closely to her? She lectures them a lot. (Rosenthal 1994:181)

More systematic versions of snowball sampling can reduce the potential for bias. For example, “respondent-driven sampling” gives financial incentives to respondents to recruit peers (Heckathorn 1997). Limitations on the number of incentives that any one respondent can receive increase the sample’s diversity. Targeted incentives can steer the sample to include

specific subgroups. When the sampling is repeated through several waves, with new respondents bringing in more peers, the composition of the sample converges on a more representative mix of characteristics than would occur with uncontrolled snowball sampling. Exhibit 5.12 shows how the sample spreads out through successive recruitment waves to an increasingly diverse pool (Heckathorn 1997:178). Exhibit 5.13 shows that even if the starting point were all white persons, respondent-driven sampling would result in an appropriate ethnic mix from an ethnically diverse population (Heckathorn 2002:17).

Lessons About Sample Quality

Some lessons are implicit in my evaluations of the samples in this chapter:

- We can't evaluate the quality of a sample if we don't know what population it is supposed to represent. If the population is unspecified because the researchers were never clear about the population they were trying to sample, then we can safely conclude that the sample itself is no good.

EXHIBIT 5.12 Respondent-Driven Sampling

