CHAPTER 2

The Scientific Approach

It [science] is not perfect. It is only a tool. But it is by far the best tool we have, self-correcting, ongoing, [and] applicable to everything. It has two rules. First: there are no sacred truths; all assumptions must be critically examined; arguments from authority are worthless. Second: whatever is inconsistent with the facts must be discarded or revised.

—Carl Sagan (1980, p. 333)

Overview

This chapter contains a description of the scientific approach as it applies to the theory and practice of research. You will learn why science, despite being the best approach to research, is not subject to proof from outside its own logical system. Scientific knowledge and its growth are a function of agreement, and you will learn how agreement is facilitated by the use of inductive reasoning. You will also learn about distinctions between scientific and nonscientific research, various misconceptions about science, and the importance of theory in the research process. You will learn how to use theory and other resources to facilitate your understanding, critical evaluation, and application of research. Finally, you will learn about the ethics of consuming research.

INTRODUCTION

Many people think of scientific research as something done by intelligent-but-absent minded people wearing white coats while surrounded by strange-looking equipment with blinking lights. Some may think of scientists as despoilers of a simple, nontechnical lifestyle. Others may think of scientists as the harbingers of an idyllic age. None of these views is correct; one of the goals of this chapter is to dispel these and other myths about science. Science is not something one does; rather, it is an approach toward doing things, and one of the most important things scientists do is research. Scientists certainly do not all wear white laboratory coats, nor do we all use strange equipment, with or without blinking lights. Some scientists may be extremely intelligent or absent minded, but these qualities do not make a person a scientist; neither does adopting a scientific approach necessarily make someone intelligent or absent minded.

I noted in Chapter 1 that everyone, not just scientists, does research. What distinguishes scientific from other kinds of research is not the activity itself but the approach to the activity. Scientific research is, among other things, systematic. There are other guidelines about what is and what is not scientific research, as well as guidelines about what to do with scientific research once we have it, including ethical guidelines. Scientists know what these guidelines are, agree about them, and attempt to adhere to them. Nonscientists either do not know them or do not consistently use them. It is not research that distinguishes scientists from nonscientists; it is the approach one takes toward research. As we learned in Chapter 1, science is a systematic approach to the discovery of knowledge based on a set of rules that defines what is acceptable knowledge. Just as there are rules for such things as tennis or international diplomacy, there are rules for science. And just like tennis or international diplomacy, not everyone necessarily operates according to the same set of rules.

A PHILOSOPHY OF SCIENCE

Years ago I was discussing religion with a friend. We disagreed about a lot of things, but we were calmly discussing the relative merits of our personal beliefs. At one point I asked my friend to explain why she believes what she does. She replied very simply, "I believe it because I know it's true." Then I asked how she knew it was true, and she said, "I know in my heart it's true." She could not explain why she believes what she believes, any more than I could explain why I believe what I do about religion. We both thought we were correct, but neither of us could logically prove we were correct in any absolute sense. At best, we could point out we were not alone in our beliefs. Of course, most people accept the notion that there is no absolute proof when the topic is religious beliefs. What many people do not understand is that the same is true about science.

Any set of rules that defines what is acceptable, empirical knowledge may be called a **philosophy of science**. Among philosophers of science and among scientists, however, there is more than one accepted philosophy. This is partly because philosophers, like members of any other discipline, are developing, changing, and assessing new ideas and formulations in an attempt to improve upon what we know. Whatever their differences, however, all philosophers of science need to address the same four basic questions: (1) When is something true? (2) If we have more than one explanation, how can we tell which one is better? (3) How can we put what we know into practice? and (4) Why do we do it the way we do it?

In this chapter we will concentrate on a particular philosophy of science called nonjustificationism (Weimer, 1979). The name of this viewpoint is derived from the position that a scientific approach cannot be justified—proven valid—except through unproven assumptions; **nonjustificationism** is a *philosophy of science for which the major premise is that we cannot logically prove that the way we go about doing research is correct in any absolute sense*. While this conclusion may seem outlandish right now, the remaining discussion should help you understand why this outlandish conclusion is quite logical and not at all inconsistent with a scientific approach to understanding the world.

When Is Something True?

This first question to be answered through any philosophy of science is usually called the question of **rational inference**—*the difficulty inherent in supporting any claim about the existence of a universal truth*. Just as when my friend and I were discussing religious beliefs there was more than one truth, there is more than one solution to the problem of rational inference. In order to be scientific, whatever we accept as our answer to the question of when something is true, i.e., our interim solution to the rational inference problem, must be based on **facts**—*phenomena or characteristics available to anyone who knows how to observe them*. Recall Sagan's (1980) second rule of science: whatever does not agree with the facts is wrong and must be changed or rejected completely. Although the statement is simple, deciding how to go about the process is a little more complex.

Behavioral scientists, for example, are interested in understanding how people interact with each other at a variety of different levels. We want to understand as much about people and human phenomena as possible. No matter how many facts we have, however, we cannot understand them until we have a way to summarize those facts. Summarizing facts—making them comprehensible—is what theories are all about.

But anyone can make up a theory about human behavior. Given enough time, just about everyone in the world could articulate some sort of theory for any given phenomenon. Thus, we have the equivalent of a very large warehouse that is full of theories. This imaginary warehouse contains as many different theories about people as there are people in the world, multiplied by the number of different theories each of those individuals has for each of the various phenomena that make up human behavior. Clearly, we need to imagine a very large (and probably quite disorganized) warehouse. Of course, each discipline has its own warehouse of theories, so deciding what to do with all of the theories in all of the sciences can be somewhat daunting, but it is not impossible.

At a very simple level, all we have to do is compare each theory in the warehouse to the facts: if the theory does not fit the facts, we change it or throw it out of the warehouse. This process may sound good, but it just does not work that way. Theories are made up of **concepts**—*abstract words that are used to represent concrete phenomena*. We can point to concrete examples of concepts, but the concepts themselves are abstract. For example, conflict, as a theoretical concept, is not the same thing as a family argument or a revolution. Family argument and revolution are, of course, concrete examples of conflict, but they are only examples and not complete definitions. No matter how compellingly practical a concept may be, it is only an approximation of reality, and any given concrete phenomenon is only an approximation of a concept (Wartofsky, 1968). Theories symbolize or represent the real world in which we live and behave, but the concepts within the theories are not the same thing as the real world. Because concepts are abstract and the facts we rely on to test them are concrete, deciding whether or not a theory fits the facts is rather difficult.

This difficulty arises because we must rely on inductive reasoning when fitting facts to a theory. **Inductive reasoning** is *a process of generalization; it involves applying specific information to a general situation or future events*. Let me illustrate with a story about a college instructor of mine who consistently arrived 10 minutes late for class. About three weeks into the semester, I came to the conclusion that he would continue to do so, which meant I could sleep an extra 10 minutes on those days and still arrive "on time" for class. This conclusion was a generalization, an inductive inference. Based on the instructor's specific behavior—arriving late during the first three weeks—I attributed to him a general or abstract quality—tardiness—and used that abstract concept to predict his behavior in the future. Unfortunately, it never occurred to me he would show up on time for the midterm exam, and I developed cramps in my hand trying to write fast enough to make up for the time I lost by arriving late. It was a rather painful way to learn that inductive reasoning does not necessarily lead to absolute truth.

Despite the inability of inductive reasoning to lead us to absolute truth, we must rely on it in any scientific approach to research. We simply cannot let all those theories pile up in the warehouse until we have all of the facts; nor can we wait for all the facts before we begin to construct theories to put in the warehouse. Instead, we simply accept the notion that inductive reasoning is the best process of generalization we have until something better comes along.

Had I waited until after the midterm exam before attributing tardiness to my instructor, I could have saved myself some writing cramps (and perhaps gotten a better grade on the exam). But even then I could not have been sure that he would be on time for the final exam, nor could I be certain that he would not begin arriving on time after the midterm. Of course, I could have just arrived on time myself every day, but that would have meant missing out on hours of extra sleep accumulated across the entire semester. I weighed the alternatives and constructed my theory about his behavior. After he showed up late the first day after the midterm, I reverted to sleeping an extra 10 minutes, but I showed up on time for the final exam. I adjusted my initial theory to fit the new facts, but I did not wait until I had all of the facts before constructing my new theory.

I have simplified the arguments involved in this issue, but the basic point of the rational inference problem is rather simple: inductive inferences cannot be proved true, but we need to use them to construct theories until we have evidence to the contrary. If we have enough contrary evidence, we can throw a theory out of our warehouse, but that does not mean that any of the theories remaining in the warehouse is true. We are left with no choice but to provide support for a theory by trying to show that alternative, competing theories are not true. If we make a prediction from a theory and test the prediction, and if the prediction fits the facts, then we have *not* proved the theory to be true; instead, we have failed to prove that the theory is false. It is difficult to think in terms of double negatives—Theory X is not not-true—but that is the logic forced on us by the rational inference problem. Thus, research in which we test between two competing theories is better than research in which we test only one theory, because comparing theories is one way to deal with the rational inference problem.

How Can We Tell Which Theory Is Better?

The absence of absolute truth does not limit what we can learn in a scientific approach, but we are faced with a particular path in our quest to learn about behavior and other real-world phenomena. We can, as I mentioned above, test between two different theories and decide which one is better. Testing between theories is like a grand tournament in which every theory is pitted against every other theory; the theory with the best win-loss record at the end of the tournament is the winner. Of course, that does not mean that the winning theory is true—only that it is the best theory we have until another, better theory is entered in the tournament. Like all tournaments, the tournament of scientific theories has some rules about which theories are entered and how many times a theory has to lose in order to be eliminated.

The rules of the grand tournament of science bring us to the problem called **criteria for growth**—*finding standards that can be used to decide that one explanation is better than another*. We all know, for example, that as an explanation of the apparent movement of the sun across the sky, current theories of astronomy are more accurate (but less poetic) than the myth of Helios, the sun god, waking every morning and driving his fiery chariot across the sky. We would scoff at anyone who seriously believed the Helios explanation, just as any ancient Greek would have scoffed at our current theories. How we came to decide that astronomy is better than mythology involves our criteria for growth: paradigms and facts.

Theories, whether in or out of our imaginary warehouse, do not exist in a vacuum. Every theory is related to at least one other theory through shared concepts or propositions. Kuhn (1962) was the first to use the term paradigm (pronounced "pair-a-dime") to describe such groups of related theories. A **paradigm** is *a logical system that encompasses theories, concepts, models, procedures, and techniques*. The earth-centered solar system, for example, was once a paradigm in physics, just as instinct was once a paradigm in psychology (McDougall, 1908). At the time McDougall was theorizing about human behavior, the predominant explanations included some notion about instinctual processes; there was an instinct for survival, one for aggression, and so on. New observations about behavior were interpreted in terms of existing instincts, and if new observations did not fit, then new instincts were invented to account for the observations.

During a period of time in which a particular paradigm is accepted, which Kuhn referred to as a period of normal science, research is directed toward solving problems related to matching current theories with observations. At such times research tends to be directed toward refining theories, toward trying to make them better, such as inventing new instincts to fit research observations. New research and the refinements of theories add to the strength of the paradigm, which in turn leads to the perception that the paradigm, including its associated theories and procedures, is the best way to explain what goes on in the world.

Eventually, however, problems with the paradigm emerge as more and more information cannot be fit into the existing theories. I note "eventually" because no matter how reasonable or useful a paradigm may be, it, too, is based on inductive reasoning and thus cannot be considered to be universal truth. When enough problems emerge and an alternative paradigm, complete with its own theories and procedures, arises that fits the observations better, then the old paradigm gives way to a new one during what Kuhn calls a scientific revolution. Thus Galileo started a scientific revolution with his notion of a sun-centered solar system, although it took years before the followers of the earth-centered paradigm accepted the new paradigm. Then the new paradigm becomes *the* paradigm and the field returns again to normal science, until the next paradigm shift occurs.

Underlying all of normal and revolutionary science is reliance on facts. Observations are considered to be facts when people can point to concrete examples of the observation. Although it may seem **tautological** to require facts to be observable, that very requirement is one of the reasons McDougall's instinct theories eventually gave way to modern explanations of behavior; there was no way to observe—to be able to point to concrete examples of—the processes by which instincts influence behavior. Today, of course, we have some evidence for instinctual processes as one of several possible explanations for some behaviors (see, for example, Lea & Webley, 2006; Snyder, 1987), but we do not use instinct as the primary explanatory concept for all behavior.

In addition to being observable, facts must also be **objective**. Within a scientific approach, **objectivity** means *that an observation can be replicated, observed by more than one person under a variety of different conditions*. If I am the only researcher who can demonstrate a particular effect, it is not objective. If, however, several others note the same effect under different conditions, then we have a fact, an objective observation that needs to be incorporated into existing theories.

Thus, during normal science, theories are compared on the basis of their fit into the existing paradigm as well as our ability to use them to account for the existing facts. During revolutionary science, comparisons occur between old and new paradigms, but the basis for such comparisons remains the existing facts. Then, upon return to normal science, theories within the current paradigm are again evaluated in terms of their fit with the facts. It is important to note, however, that because a new paradigm may redefine what is an acceptable fact, the facts may change from time to time (Fleck, 1979).

How Can We Put What We Know Into Practice?

By now you may be having some serious doubts about how a scientific approach can be a path to anything except confusion. There are no absolute truths, and sometimes what were once considered to be facts are no longer considered to be so. We have arrived at the problem of **pragmatic action**—*determining how we should go about putting a scientific approach into practice*. Essentially, those who adopt a scientific approach must get together and decide how they are going to use that approach. The solution to the problem of pragmatic action, the answer to the question of how we put what we know into practice, lies in agreement.

Just as legal theorists assume that a decision made by 12 jurors is better than a decision made by one juror, scientists agree that evidence obtained by a number of different researchers is better than evidence obtained by one researcher; that is, objective data—repeatable observations—are agreed to be better than subjective data. The greater the number of researchers who produce the same research results, the more we consider those results to be facts to which we must fit our theories; notice that the theories must fit the facts, not the other way around. A variety of reasonable arguments support this agreement about objectivity, but no one can prove, in any absolute sense, that the consensus is correct. As Sagan suggested, it is not perfect, but it is the best we have.

One of the problems inherent in the use of objectivity is the variety of different research methods available to study any particular phenomenon (see, e.g., Watson, 1967). When researchers use different methods to study the same phenomena, they often come up with different observations. Consensus, then, must extend into agreement about which research methods are appropriate for which research questions, as well as agreement about whether or not a particular method was used properly. Essentially, that is what this book and the course you are taking are all about. You cannot rely solely on the assumption that the experts have used the correct research method to answer their question; you must be able to determine yourself whether the methods used by the researchers fit the way in which you want to use the research results.

For example, in the early years of research about differences between men and women, one of the more common methods was to select a group of men and a group of women, have both groups do something such as solve math problems, and then compare the performances of the two groups. If the performances of the groups were different, then the researchers concluded that the results reflected basic differences between the two sexes. Deaux and Major (1987), however, presented convincing, empirical arguments that such things as the context of the situation, self-presentation strategies, researchers' and participants' beliefs about whether or not the sexes ought to be different, and a variety of other factors can change the results obtained from such methods. Therefore, the potential influence of such factors must be considered before we conclude that gender differences reflect basic differences between men and women.

We now know that simply comparing a group of men to a group of women is not an effective way to examine gender differences. Then again, everyone "knew" back in the old days that such simple comparisons were the best way to study gender differences. Even though we rely on consensus for such purposes as fitting theories to facts and even for deciding what is a fact, we must keep in mind that a new consensus might emerge after we have obtained more information. Still, there can be no scientific approach without consensus.

Why Do We Do It the Way We Do?

Every time I discuss consensus as the basis for a scientific approach, I can hear my mother saying, "Would you jump off a cliff just because everyone else is doing it?" That was her response, for example, to my wanting to stay out late because my friends' parents allowed them to stay out late; I am sure you have heard the same response when you have tried to use similar reasoning, or you have provided the same response when your children used that reasoning. What we have come to, then, is the problem of **intellectual honesty**—*the individual scientist's ability to justify the use of science itself*. If we can never prove that theories are true, if paradigms are only temporary, and if facts and methods for gathering them may change, then why would we ever accept a scientific approach as a valid way to learn anything?

Consider a simple survey of students' attitudes about current grading practices. In order to understand and apply that study, we must rely on a great deal of background information. We must accept research about students' reading levels when examining the questionnaire, accept research that suggests that a survey is a reasonable way to measure attitudes, accept research concerning the best way to format the questions on the survey, accept research about which statistics are appropriate to analyze the data, and so on. All of that research comes from within a scientific approach, and we are using that information to add more facts to the same scientific approach. Where does it all end?

The solution to the intellectual honesty problem—the answer to why we do it the way we do—can again be found in Sagan's quote at the beginning of this chapter: it is "by far the best tool we have." We do it the way we do it because we have not found a better way. Very simply, we adopt a scientific approach because we have a certain amount of faith in it because it works; or as my grandfather used to say, "If it ain't broke, don't fix it." Note, however, that the faith is placed in the approach itself, not in any particular theory that comes from the approach.

Recall the debate between evolutionists and creationists mentioned in Chapter 1. Theistic intervention (active causation by a god, not necessarily the existence of a god) serves as an explanation for creationists because they have faith in their approach. Similarly, evolution serves as an explanation for evolutionists because scientists have faith in their approach. Neither side will be able to convince the other because they do not have any common ground of agreement; they place their intellectual honesty on two entirely different points of view (Bobkowski, 2007). If you refuse to place your faith in a scientific approach, no amount of argument on my part will convince you to do so. On the other hand, if you can accept the limitations of a scientific approach and still remain convinced it is the best tool we have for extending our knowledge about the world, then we can go on to discuss some of the differences between scientific and nonscientific approaches and we can begin to deal with the rules and guidelines of scientific research. But before moving on, see Table 2.1 for a summary of our discussion of the philosophy of science.

Table 2.1 Justificationist and Nonjustificationist Approaches to the Four Basic Questions Inherent in Any Philosophy of Science

The Questions	Justificationist Approach to Science	Nonjustificationist Approach to Science
The rational inference problem: When is something true?	Facts → One correct theory	Facts → Many incorrect theories
Criteria for growth: How can we tell which theory is better?	Better fit with paradigm and facts	Better fit with paradigm and facts
Pragmatic action: How can we put what we know into practice?	Consensus → Correct paradigm	Consensus → Better, but not correct paradigm
Intellectual honesty: Why do we do science the way we do?	Science → Absolute truth	Believe science is the best way to obtain knowledge

SCIENCE AND NONSCIENCE COMPARED

I keep bringing up the notion that we all conduct research all of the time. We are all, in one way or another, gathering new information to increase our knowledge about our world. Such everyday research is not necessarily scientific, but it does provide us with a way to satisfy our curiosity. In addition to the points noted above, the differences between scientific and nonscientific research generally revolve around avoiding mistakes. Mistakes can occur when we make observations, when we interpret observations, or when we accept various misconceptions about what is included in a scientific approach toward research.

Observation

Whenever we observe something, we make errors; period, no exceptions, ever. The errors, which researchers generally call bias, come from selecting what to observe and interpreting what we observe, as well as from the act of observation itself. We cannot avoid bias entirely, but we can attempt to reduce error to a minimum and be aware of error that we have not been able to eliminate.

For example, what we decide to observe creates a form of bias because it prevents us from making other observations at the same time. This is an error of omission that results simply because we cannot be in two places at the same time. That does not mean that what we do observe is wrong or incorrect, but rather that it is incomplete. Essentially, we need to keep in mind that what we have been observing is not all that could be observed. For example, researchers before Jacobs's (1967) study on suicide notes had concentrated mainly on suicide statistics. Durkheim's (1951/1897) original work had led them in that direction, and that approach was certainly adding to our understanding of suicide. But suicide statistics would never have given us the kind of information that Jacobs was able to obtain from his analysis of the contents of suicide notes. The bias in this example involved looking in only one place, a bias that Jacobs corrected with his study. This is one of the ways in which science is self-correcting; new observations correct the errors of previous observations.

Of course, objectivity is another way to reduce, but not eliminate, the bias inherent in observation. When more people observe the same thing, under the same or different conditions, then the collection of observations becomes more accurate (less biased, more complete). Different observers, different situations, different locations, and different definitions of what to observe all contribute to the objectivity of data, and all reduce observation error. Realizing that all observation contains some amount of error or bias is an important part of a scientific approach to research, for it prevents anyone from saying, "Your results are wrong and mine are correct." If we accept the notion that everyone's data are a little bit wrong (contain some error, some bias), then we can concentrate on trying to figure out why our observations do not agree; that is, we can begin to refine our theories so that they more closely fit the existing facts.

Logical Analysis

The quality of observations is one distinction between scientific and nonscientific research, but it is far from being the only one. Once observations are made, we must interpret them and draw conclusions about them. We have already discussed the scientific reliance on inductive reasoning, so it should come as little surprise that induction plays an important role in data interpretation.

Suppose I look out my window and observe 90° displayed on the scale of a thermometer. I could, of course, reasonably conclude that the temperature outside my office is 90°, assuming I had reason to believe that my thermometer was accurate. Anyone else could also look out the same window and note the same reading, and they would probably come to the same conclusion. Inductive reasoning enters the interpretation process when we attempt to move our conclusions beyond the immediate area outside my window, beyond the immediate confines of the data collection environment. Beyond my window is the remainder of the campus, the town, the county, the state, the country, and so on. How far beyond our immediate observations we can reasonably interpret those observations is both a matter of inductive reasoning and yet another distinction between scientific and nonscientific research.

Given our general knowledge about meteorology, we could reasonably conclude that the temperature around campus and town is about 90°. I would be reluctant to speculate about temperature across the state, as would most people. The same reluctance applies to interpreting data collected in a research project: how far we **generalize**, *relate findings gathered from the research situation to other situations*, is limited by common sense and background information about the research topic. If college students participated in a study about jury decision-making processes, I would feel comfortable generalizing the results to actual jurors by claiming actual jurors may use the same decision-making *processes* that the students used. However, I would not feel comfortable claiming that actual jurors would make the same *decisions* that the students did. The way in which students and actual jurors go about making decisions may be the same, but the decisions produced by that same process may be quite different because they may pay attention to different information (have different biases) and may have different life experiences with which to interpret the information they receive. **Overgeneralization**—*drawing conclusions too far beyond the scope of immediate observations*—brings scientific research into the realm of nonscientific research.

Research Reports

From time to time you may find yourself reading a research article in which it appears as though the researchers designed their study to test a theory, collected data, and supported the theory discussed in the introduction of the article. You should know, and the researchers should know, that logic does not enable us to support a theory. Yet they write such phrases as "research supports the theory of. . . ." or "the theory of X has received a great deal of empirical support." In such cases the language of scientific research appears to conflict with scientific philosophy.

Keep in mind that the reason that research cannot support a theory is that "support" for a theory comes not from finding results consistent with a theory but from failing to find results that do not fit the theory. Remember the double negative logic of science: Failing to disconfirm a theory is the only empirical way to provide support for a theory. But support for a theory does not mean the same thing as proof that a theory is correct. It is a little too verbose to write "a number of researchers have attempted to disconfirm theory X and have failed to do so" continually, and so we sometimes write "theory X has received empirical support."

Most authors of research articles create the impression that the researchers knew, from the start, exactly how the major results of the study would come out. Instead, research is often conducted with extremely little certainty about how the results will turn out. The researchers are not trying to hide their inability to predict the results accurately; rather, they are succinctly providing a theoretical context for their results. No matter how unexpected the results of research may be, they cannot contribute to what we already know unless they can be placed into a theoretical context.

For example, when I was asked by a defense attorney to consult on the voir dire (often called "jury selection") process in a criminal trial, I took advantage of an opportunity to interview the jurors after the trial. What I noticed from these interviews was that jurors who seemed very different on the basis of such characteristics as age and socioeconomic status also ended up on opposite sides of the hung jury resulting from the trial (Whitman & Dane, 1980). If you read the publication that resulted from those interviews, it appears as though I had a very logical theoretical framework before I conducted the interviews and that I found support for the theory in the results of my study. Not so.

I did have an initial theoretical framework before conducting the research, but it was nowhere near as logically organized as it was presented in the article. It was more like a hunch based on what I knew about jury behavior and social psychology. However, a scientific approach to research does not involve telling everyone about hunches; it involves presenting research results that help evaluate theories. So after interpreting the results, I wrote a very logical, organized, theoretical introduction to the article so that my presentation of the research results made sense in the context of the introduction. Kerr (1998) refers to such writing as HARKing (Hypothesizing After the Results are Known) and notes that there are costs associated with such writing. As a general rule, HARKing ought not to be done without an explicit statement that one has done so. As a consumer of research, however, you will find it difficult to determine whether or not an author has engaged in HARKing.

Definitive Studies

Although any study may satisfy someone's curiosity about a particular issue, no study ever satisfies scientific interest in an issue. That is, despite the fact that one often hears the phrase used in one or another context, there is no such thing as a **definitive study**—*a research project that completely answers a question*. Because any particular phenomenon is extremely complex, someone will always ask, "But what if. . . ?" Such questions point out the need for additional research. Proposing that a definitive study can exist produces premature closure of activity; as Yogi Berra is supposed to have said about a baseball game, "It ain't over 'til it's over." It is, of course, difficult to argue with such logic. Within a scientific approach to research, it is not over until it is no longer possible to ask "What if?"

Although definitive studies may not exist, there are highly influential studies that set an entire research program, or series of programs, in motion. These studies have a great deal of **heuristic value**—*they stimulate additional research activity*. Milgram's (1963) research on obedience is one example of a study with high heuristic value. It not only generated a great deal of controversy concerning research ethics; it also stimulated extensive research on compliance of individuals and groups. Munsterberg's (1913) studies of the accuracy of eyewitnesses' recollections, many of which were demonstrations conducted in the classroom, were also highly heuristic. Many examples of current research on eyewitnesses can be traced to one or another of his demonstrations.

As an administrator, you should neither look for nor believe you have found a study that conclusively proves whether or not a program is effective; you won't find such a study because they simply don't exist. You will, however, find many claims that others have found such a study. I recently searched for the phrase "science proves" on Google and turned up about 93,700 Web sites, not all of which were quack sites, which merely demonstrates that there are many individuals who don't understand the limitations of science. In case you are interested, another search for the phrase "research proves" resulted in about 158,000 Web sites; a quick scan of some of those sites convinced me that many writers confuse *prove* with *demonstrate*, a confusion that could lead to erroneous conclusions about the value of a program or policy.

Determinism

Perhaps the most misunderstood concept in a scientific approach to research is **determinism,** the *assumption that every event has at least one discoverable cause*. As defined here, determinism means nothing more than "events do not happen by themselves." We assume that there is always a causal agent and that the agent can be discovered through a scientific approach to research. If you think about it at all, you will realize that there could not be science without determinism. The purpose of psychology, for example, is to understand the causes of human behavior; if we did not assume that every human behavior had at least one cause, then there would be no point to trying to understand the causes of human behavior.

Many people, however, incorrectly mistake determinism for **predestination,** the *assumption that events are unalterable*. The two assumptions clearly are not at all similar. Indeed, there is some notion in determinism that once we are able to discover the cause of an event, we can alter the cause and thereby alter the event. There may, of course, be theories that include the assumption of predestination, and some of those theories may be tested through scientific research, but predestination is an aspect of a specific theory and not an assumption inherent in science.

Table 2.2 contains a summary of the differences between what is and what is not included in a scientific approach to research. Although there may be many other comparisons that could be drawn, you should have enough background in philosophy of science to begin putting it into practice.

Table 2.2 Comparisons Between Science and Nonscience

Science Is	Science Is Not
A way to obtain new information	An activity per se
Described by a philosophy	Defined by only one philosophy
Generalizing from facts	A way to prove theories true
Grounded in paradigms	Blind acceptance of tradition
Based on consensus	Relying on personal authority
A matter of faith	Uncritical faith
Deterministic	Predestination
The best approach we have	Refusing to search for a better approach

ETHICS OF CONSUMING RESEARCH

I continue to be amazed that codes of ethics for researchers, particularly those who do research on human beings, did not emerge until the middle of the 20th century. To that point, the public, including researchers, assumed that scientists had sufficient integrity so as to make formal guidelines and regulations unnecessary, but the torture and other inhumane treatment of concentration camp inmates by the Nazis during World War II convinced the world community that guidelines were necessary, and The Nuremberg Code was enacted (Nuremberg Military Tribunal, 1949). The code emphasized the importance of **informed consent**, *the process of ensuring prospective participants have all the knowledge they need to make a reasonable decision concerning their participation*, and a balance of risks and benefits such that the latter outweighed the former (Gorman & Dane, 1994). Expansion of the code for physicians and other medical researchers resulted in the Declaration of Helsinki (World Medical Association, 1964), which further clarified the relationship between research and treatment.

In the United States, breaches of research ethics, again primarily in biomedical research (see, e.g., Jones, 1993), led to the passage of the 1974 National Research Act. The act included creation of the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1979; Seiler & Murtha, 1981), the purpose of which was to recommend an overall policy for research with human participants. The commission provided a report that led to a number of changes in the ways in which biomedical and behavioral research is conducted. Eventually, a formal set of regulations was adopted in the United States ("Final regulations amending basic HHS policy for the protection of human research subjects," 1981; "Protection of Human Subjects," 2001) that included establishment of institutional review boards (IRB) to conduct prior review and continuing oversight of all human subjects research conducted within the purview of any institution or organization that received federal funds. Similar developments in other countries have led to nearly global adoption of some form of guidelines or regulations concerning the conduct of human subjects research (Office for Human Research Protections, 2007).

The existing regulations and guidelines provide considerable direction as researchers attempt to balance the mutual obligations of developing new knowledge (Cook, 1981; Mindick, 1982) and treating individuals involved in our research with proper consideration (Dane, 1990). If you find yourself in a position to conduct research, or direct others to conduct research, you should become familiar with these guidelines and regulations. For the purpose of consuming research, however, it is more important to understand the principles that are used to guide the development of such regulations and to recognize the ways in which these principles apply to research consumption.

Ethical Principles of Research

In the Belmont Report, the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (1979) established three principles by which all research should be guided—respect for persons, beneficence, and justice—and to these have been added trust and scientific integrity (Dane, 2006, 2007b; Dane & Parish, 2006). These five principles will form the basis of our discussion of the ethics of research consumption.

Respect for Persons

Respect for persons is a principle derived from the ethical theory proposed by Immanuel Kant (1788/1997), from which we obtain the admonition never to use another human being merely as a means to an end; that is, this principle involves maintaining others' **autonomy**, *the ability to direct oneself, particularly through the exercise of independent processing of information*. We see this implemented, among other ways, in the presumption of informed consent for research participants. In the context of research consumption, respect for persons involves giving researchers proper credit, representing others' work accurately, and providing comprehensive information to those who will benefit from our consumption of research.

At this point in your academic experience you are well familiar with the requirement to use quotation marks or other, similar conventions whenever you use someone else's words; to do otherwise is plagiarism. The purpose of this convention is to ensure that the individual who wrote the words receives credit for having done so and to prevent readers from thinking that the words are ours instead of the original author's. Similarly, whenever we obtain information from another source or use an idea obtained from someone else's work, we give credit by citing the source from which we obtained the information or idea; again, this prevents readers from thinking that the information or idea is original to us. Giving researchers proper credit, however, involves going beyond the usual conventions to avoid plagiarism; we also do so through the manner in which we write about the research. Thus, we should refer to the authors, not to unnamed "researchers" or "research" (see, e.g., American Psychological Association, 2001). For example, when describing the research mentioned in Chapter 1, we should write "Sales (1972) demonstrated an increase in membership in authoritarian religions during difficult economic times" instead of writing "Research has shown an increase in membership in authoritarian religions during difficult economic times (Sales, 1972)." While both sentences provide credit to Sales for the information, the former makes it clear that Sales conducted the research; the latter refers only to Sales as the source of the information and could mean that anyone conducted the research about which we learned from reading Sales.

Just as we respect others by giving them proper credit as the sources of the information we use, we also respect others by representing their work accurately. Obviously, we want to convey correctly information we obtain from others, but accuracy goes beyond "getting it right" in the sense that we must avoid oversimplifying research results. Rector (2002), for example, reported that an evaluation of the Not Me, Not Now abstinence-only advertising campaign (Doniger, Riley, Utter, & Adams, 2001) included a reduction in pregnancy rates among 15-year-old teens but did not report that Doniger et al. also found no change in pregnancy rates among women aged 17. The failure to report the additional results oversimplifies the results, and could mislead those reading the report, leaving them to think that the change in pregnancy rates was longer lasting than it actually was. Conveying the complications sometimes demonstrated in research can be difficult, but our ethical obligation to respect individuals, those who reported the research and those who will read our review of that research, requires that we overcome such difficulty. Accurately reporting results also involves making sure that our readers understand what was measured in the research we review. The Union of Concerned Scientists (2006), for example, reported numerous examples in which government reports concerning abstinence-only sex-education programs labeled various programs as effective without noting that *effective* was defined in terms of attendance or changes in attitudes about sexual behavior; actual sexual behavior was not measured in the research included in these reports. There is, of course, nothing wrong with considering attitude change to be a desired outcome, but it is inaccurate, and therefore undermines autonomy, when one implies that changes in attitudes toward sexual activity are synonymous with changes in sexual behavior, per se.

Finally, we demonstrate respect for persons when our reports about research are comprehensive, when they include all relevant research and not just those studies that conform to our preferences or those of our intended audience. Imagine your reaction upon reading about a researcher who reported only part of the data, only those data that were consistent with a conclusion drawn by the researcher even before data were collected. I hope your reaction would include outrage and a general conclusion akin to "that's not right" or "that's dishonest." Indeed, such behavior would be evidence of a lack of respect for one's audience through undermining autonomy by misleading the readers; it would be dishonest. A failure to present all relevant studies in a review of research, too, would be similarly dishonest. The key word in the previous sentence is "relevant." As of this writing, a quick search of the ProQuest database of publications yields 157 articles in scholarly journals on abstinence-only sexual education, 21 of which also include the key term *evaluation*. Thus, a review of the effectiveness of such programs would not necessarily need to include all 157 articles but most probably should include the 21 articles involving evaluations. Similarly, the same search conducted through Google yielded about 99,900 sites for "abstinence-only sex education," which would be impractical to review, but only 55 sites when the phrase "evaluation report" was added to the search box, a much more manageable number of sites to examine. Between the 21 ProQuest articles and the 55 sites identified through Google, one would have a very good start to meeting this ethical obligation to include relevant material.

In summary, respect for persons, in the context of consuming research, involves giving credit where it is due, accurately reporting research procedures and results, and comprehensively reporting the available, relevant research. To do anything less undermines the autonomy of our audience; it reduces the accuracy or amount of information for decision making available to our audience members.

Beneficence

While researchers have an ethical obligation to generate new knowledge (Cook, 1981; Mindick, 1982), they also have an ethical obligation to do so in a manner that promotes the public good (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1979). The notion that knowledge is beneficial in and of itself has a long tradition in Western thought (see, e.g., Plato, 2005), and so, to some extent, researchers satisfy both ethical obligations simply by producing new knowledge. As research consumers in the role of administrator, policy analyst, or policy implementer, we also have an ethical obligation to promote the public good, to engage in beneficence (Cooper, 1998; Svara, 2007). We, too, partially fulfill that obligation simply by producing new knowledge.

It may seem strange to think that a reviewer of others' work produces new knowledge—after all, the information already exists in the original research reports we review—but a good review goes beyond simply noting what others have found in their research. Yet even simply noting what others have done can be beneficent in the sense that a list and description of sources in a single document makes information available in a more convenient format. Nevertheless, a good review involves using the information in existing sources to make a particular point, to draw a conclusion that would not have been obvious from reading only one or two of the original sources.

Of course, one can argue that adopting an empirical approach to policy analysis is itself another way in which to promote beneficence. As noted by Mead (1969), increases in knowledge often precipitate fear about how that knowledge will be used. In the context of policy implementation and analysis, a careful review of research knowledge about the effects of the program can be used to reduce such fear. Reviews can be used to identify ways to improve programs, as opposed to recommend only "yes" or "no," in much the same way that research is used to improve scientific theories. The absence (or insufficiency) of empirical research can be used to promote additional research on a specific program or a social problem in general. Many people also tend to be afraid of empirical program or policy evaluation as a result of experience or hearsay regarding misuse of such efforts (see, e.g., Posavac, 1994), but a careful, comprehensive review of research can be used to allay such fears. Thus, the ethical obligation of beneficence, promoting the public good, is relevant to reviewing research for program or policy analysis; indeed, the obligation can be met by ensuring widespread dissemination of research reviews.

Justice

In the context of conducting research, the ethical obligation for justice refers generally to ensuring that risks and benefits associated with research are distributed equally throughout the population (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1979). For example, toward the end of the 20th century a great deal of attention was focused on the underrepresentation of women included as subjects in biomedical research (Mastroianni, Faden, & Federman, 1999). This meant that men were much more likely to benefit from the results of such research and that too little attention was paid to women's health issues, which eventually came to be perceived as patently unfair or unjust. More recently, we see that the risks associated with stem cell research are being inequitably borne by those governmental regions that allow such research (Dane, 2007a).

It may seem puzzling to think in terms of risks associated with reviewing existing research: After all, how likely is it that someone could be harmed by your writing about research that has already been done? But reviews can have considerable impact, particularly when they are used to evaluate social programs and policies. Jensen (1969a), for example, reviewed research on intelligence, particularly twin studies, to inform efforts related to **compensatory education**—*attempting to raise the educational level of individuals whose education has been disadvantaged for one reason or another*. Among other questionable aspects of his review was the combination of his noting that "all of the major heritability studies reported in the literature are based on samples of white European and North American populations, and our knowledge of the heritability of intelligence in different racial and cultural groups within these populations is nil" (1969a, p. 64) and, 20 pages later, using heritability data to support the claim that genetic factors were the primary explanation of the differences in intelligence-test scores between Whites and Blacks in the United States. The erroneous conclusion, however, was not merely an example of "could be better" writing; Brazzill (1969) noted that only five days after Jensen's research made headlines in Virginia newspapers, defense attorneys in a school integration suit were quoting from Jensen's article to support their claim that school systems should remain segregated. The defense attorneys and the newspaper article about Jensen's research appeared to have completely ignored the qualifications that Jensen included in his article and a subsequently published reply to critics (Jensen, 1969b). Jensen, fulfilling his ethical obligation, wrote a letter to the editor in response to the misapplication of his results, but by that point the damage was done.

Again, we return to the notion that careful writing is important, but meeting one's ethical obligation of justice involves going beyond careful writing. You have already read about the difficulties inherent in generalizing from research results, and this becomes particularly relevant when reviewing research for program evaluation or policy analysis. It becomes quite important to consider the range of research subjects included in the relevant research and compare that range to the diversity of individuals who are to benefit from a program or will be affected by a policy. If most evaluation research involved teenagers, for example, fulfilling the obligation of justice would include an explicit, comprehensively presented caution about applying the research to a program for younger children, as well as an explicit call for research that included younger children.

At the heart of the concept of justice with respect to empirical policy analysis is the notion that Campbell (1969, 1971) referred to as the experimenting society, which involves both a willingness to evaluate the outcomes of programs and policies and the willingness to try new programs without guaranteeing that the outcomes will be positive. If policy makers and others present a new program as some sort of guaranteed treatment or cure for a particular social problem, then they become committed to the program and very reluctant to reverse their positions and admit that the program was less effective than required or to alter the program substantially in light of empirical research. Once reputations are on the line, there is a strong tendency to find ways to describe the program as effective or valuable even in the face of empirical information to the contrary. Despite considerable research on the lack of effectiveness of Drug Abuse Resistance Education, also known as Project DARE (Lynam et al., 1999), the commitments, one could argue "guarantees," of those who initially and continue to tout the program (see, e.g., Dillon, 2006) have created considerable pressure to maintain the program. It is, of course, not possible to calculate how much time and money could have been saved if, instead, those involved in the program had adopted an "experimenting" approach and incorporated all relevant empirical research into subsequent planning and implementation. Few would argue with the notion that societal resources should be allocated equitably, fairly, or otherwise justly, and this ethical obligation applies equally to an equitable allocation of resources to programs proportionally to their effectiveness. Reviewing research and applying the results of that research is one way to fulfill this obligation.

Trust

As we learned earlier, the enterprise of science relies heavily on trust. Although a scientific approach is self-correcting, scientists have to rely on each other, trust each other, because it is logistically impossible to double-check everything, to look over each scientist's shoulders to ensure that he or she is doing exactly what was presented in a research report. Thus, those reading about our research on the effects of training nurses in advanced cardiac life support (ACLS) (Dane et al., 2000) have to trust that we correctly identified which nurses were trained and which were not, that we correctly identified which patients survived a resuscitation attempt and which did not, that we actually conducted the statistical analyses that we wrote about, and so on. The self-correcting nature of science will determine whether or not the sample we used was unique—if no one else can replicate the results we obtained, then our sample was unique—but ultimately, the entire scientific community must trust that we did the research the way we described doing the research.

So, too, those reading our reviews of research must trust that we actually read and understand the research reports we include in our reviews; they must trust that we made a good-faith effort to identify all relevant research reports, and so on. Our willingness to enter the scientific community as consumers of research thusly engenders an ethical obligation to be as trustworthy as those who produced the research being consumed. Beyond simple honesty, those who will read our reports and those who will be otherwise affected by the content of our reports are owed a high degree of competency in the construction of those reports.

Particularly in the areas of public policy, all those who are affected by such policy must place a considerable amount of trust in those who develop, analyze, and implement such policy. Indeed, some authors (e.g., Cooper, 1998; Svara, 2007) argue that maintaining public trust may be the most important ethical obligation of program administrators and others who work in areas affected by public policy. As with other ethical obligations, honest, comprehensive, careful, and expert review and application of research brings us quite far along the path of fulfilling the obligation of trust. Independently evaluating the quality of the research, for example, meets this obligation much better than merely accepting researchers' conclusions because they're the experts. We know that the self-correcting nature of research means that, indeed, errors are made; thus, we become part of the self-correcting process by critically examining research we review. Even though we may have biases about what conclusions we would like to include in our review, either for or against a particular program or policy, we have an obligation to set aside those biases and ensure that our conclusions are based on sound, methodological principles rather than on personal preferences.

Scientific Integrity

The ethical obligation for scientific integrity involves having respect for the scientific process itself and acting accordingly. For researchers, this involves, quite simply, "doing good research" in all that the phrase entails (Committee on Assessing Integrity in Research Environments, 2002). For consumers of research, this involves adopting a scientific approach to reading, understanding, and reporting on others' research, even if one is not actually a scientist. Thus, it involves behaving like a scientist when using science as the basis for commenting on programs or policies.

In addition to intellectual honesty, accuracy, fairness, and respect for those involved directly and indirectly, scientific integrity involves giving careful consideration to actual and potential conflicts and, if relevant, explicitly declaring those conflicts and comprehensively trying to overcome such conflicts. Just as in Chapter 1 we learned that one important part of critically reviewing research involves assessing the "who" of research, an important part of reviewing research ethically involves making potential readers aware of who we, the reviewers, are and what biases might have influenced our review. Although it is best to avoid conflicts of interest, it is probably not possible to do so entirely (Adams, 2007; Kimmelman, 2007; Pachter, Fox, Zimbardo, & Antonuccio, 2007). A program or policy you are reviewing may be the *raison d'être* for your department and an unfavorable review may put your position in jeopardy, or the program or policy may be a favorite of influential policy makers. In such circumstances, even the best attempt to be objective may not be fully successful. In a very well known study, for example, Rosenthal and Fode (1963) demonstrated that telling undergraduate researchers that rats were "bright" as opposed to "dull" resulted in "bright" rats performing better in a discrimination task. Similarly, Rosenthal and Jacobson (1966) demonstrated that, compared to a **control group** (*participants who receive no treatment or any set of participants to which the treatment group is being compared*), randomly assigned first- and second-grade students gained 10 to 15 IQ points after their teachers were falsely informed that the students had scored highly on a "test for intellectual blooming" (p. 115). (The effect was not obtained among students in Grades 3–8.) In neither study was there any evidence of cheating, such as falsely reported scores, among the researchers or teachers; Rosenthal et al. believed the effect was much more subtly produced, probably without the researchers' or teachers' intentional efforts or awareness.

Given that we cannot avoid conflicts of interest entirely and that we cannot avoid bias entirely when we have a conflict of interest, you may be tempted to think that there is nothing to be done, that we simply cannot meet the ethical obligation for scientific integrity, but you would be mistaken in that conclusion. There are tactics we can employ to overcome bias in a conflicted-interest situation. Perhaps the most important of these is becoming aware of the conflict (Cain, Loewenstein, & Moore, 2005). In the Rosenthal et al. studies (Rosenthal & Fode, 1963; Rosenthal & Jacobson, 1966), for example, neither the student researchers nor the teachers were aware that the information presented by the experimenters could bias them; you are in a much better position because you are now aware that even seemingly innocuous information can produce bias. Knowledge of potential bias is necessary but not sufficient; we also need to consider carefully and objectively how the bias could affect our decisions about the research we are reviewing. In other words, we need to think critically about our own thought processes and make our reasoning explicit (Moore & Loewenstein, 2004). Cooper (1998) refers to this as the "*60 Minutes* Test," as in thinking about how we would explain our decision if we were facing one of the interviewers on the CBS television show. Thus, while we make notes about one or another research project, we should keep asking ourselves questions such as these: Is this note correct? Would the project researchers agree with my characterization of their research? or Could someone else argue with my notation? As well, the text in our review must also match the notes we have taken.

In general, then, the ethical obligations with respect to consuming research are somewhat different than those for conducting research, despite the obligations having been based on the same five principles: respect for persons, beneficence, justice, trust, and scientific integrity. With the exception of scientific integrity, as opposed to integrity in general, you probably noticed that the obligations are very much the same as those included in the code of ethics, guidelines for responsible conduct, or other requirements established by the professional organizations of which you are a member. When we are functioning as experts, we must pay considerable attention to the actual and potential impacts we have on nonexperts. Similarly, when we function as research experts, we must pay attention to the actual and potential effects of our actions.

SUMMARY

- Science is not an activity but rather an approach to activities that share the goal of discovering knowledge. One of these activities is research.
- Like any approach, a scientific approach has limitations. These limitations include rational inference, criteria for growth, pragmatic action, and intellectual honesty.
- Rational inference is a limitation on the extent to which we can propose universal truths. Because we must rely on inductive reasoning for such proposals, we cannot prove their accuracy. Thus, we accept theories as temporarily correct, while always assuming that another, better theory is likely to come along.

- Criteria for growth is a limitation on the standards by which to judge the relative merits of explanations. Although such judgments are based on objective observations, we must be aware that the objectivity and relevance of observations are limited to the paradigm on which their relevance and objectivity are based.
- Pragmatic action is a limitation on the practice of research concerning methodological issues. Consensus, based on sound reasoning, is the way we decide how best to practice research.
- Intellectual honesty is a limitation on our willingness to accept a scientific approach. Placing one's faith in the scientific approach, however, does not involve believing in one or another particular theory.
- It is axiomatic that all observations contain some degree of error. Objectivity—the extent to which more than one observer can make the same measurement—decreases measurement error but does not eliminate it.
- Although research reports are written so as to place research results in a theoretical context, it is often the case that the theoretical context was logically derived after the research was conducted. This is a shortcoming when the author suggests that the hypothesis was derived prior to data collection.
- Despite the fact that a scientific approach includes the goal of comprehensively testing theories, there is no such thing as a definitive study. No study produces the final answer to a research question, in part because there is always the possibility that another theoretical context raises additional questions.
- One of the basic assumptions of a scientific approach to research is determinism—every phenomenon has at least one discoverable cause. Although people often confuse determinism with predestination, the two concepts are entirely different. Predestination refers to the belief that events cannot be altered.
- Regardless of the point at which one begins a research project, the project is always related to one or another theory. Variables—logically derived, concrete representations of theoretical concepts—are used to form hypotheses; it is hypotheses that are directly tested in a research project.
- Construct validity refers to the extent to which a variable represents a theoretical concept. Consensus is necessary for validity, but it is possible to misuse a variable on which consensus has been achieved. Avoiding the belief that a variable is the same as the concept it represents prevents such misuse.
- Ethics for scientific research were not formally codified until the second half of the 20th century. The first of these, The Nuremberg Code, was in response to the Nazi atrocities.

- There are five primary principles that guide ethical research: respect for persons, beneficence, justice, trust, and scientific integrity. These same principles can be used to guide ethical consumption of research.
- Respect for persons involves giving researchers proper credit, representing others' work accurately, and providing comprehensive information to those who will benefit from our consumption of research.
- Beneficence involves promoting the public good.
- Justice refers generally to ensuring that risks and benefits associated with research, including the consumption of research, are distributed equally throughout the population.
- Trust involves much more than honesty; those who will read our reports and those who will be otherwise affected by the content of our reports are owed a high degree of competency in the construction of those reports.
- Scientific integrity involves intellectual honesty, accuracy, fairness, and giving careful consideration to and, if relevant, explicit declaration of conflicts of interest and a comprehensive effort to overcome such conflicts.

Exercises

1. Find examples in which people have written or said things that indicate they do or do not understand the rational inference problem in science.
2. Find examples in which people have written or said things consistent with the notion of comparing one theory against another. (Note: This may be somewhat difficult because most popular-press reports usually mention only one theory, if any.)
3. Find examples of reports in which the author(s) claim(s) to be reporting a consensus about scientific conclusions. Can you determine the source of the reported consensus?
4. Find an example of what you think is biased reporting of research. Explain why you think it is biased. Can you think of any conflicts of interest you might have with the report?