2 WHAT IS SCIENCE?

LEARNING OBJECTIVES

1. Identify what distinguishes science from nonscience.
2. Describe the key steps in the scientific method.
3. Identify valid and invalid arguments.
4. List some myths about science.
5. Assess the importance of diversity for science.
6. Construct a scientific model to explain a puzzling observation.

Consider the following five statements. What do they all have in common?

1. Science is a collection of facts that tell us what we know about the world.
2. A scientific theory is one that has been proven.
3. "The sun revolves around the earth" is not a scientific statement.
4. If my theory is correct, then I should observe that rich countries are more likely to be democracies. I do observe that rich countries are more likely to be democracies. Therefore, my theory is correct.
5. Politics cannot be studied in a scientific manner.

The common element in these statements is that they're all wrong. Science isn't a collection of facts that tell us what we know about the world. Scientific theories can't be proven. The statement that the sun revolves around the earth is a scientific statement (even though it's false). The argument outlined in statement 4 is logically invalid and therefore I can't conclude that my theory is correct. And finally, politics can be studied in a scientific manner. We suspect that many of you will have thought that at least some of these statements were correct. To know why all of these statements about science are wrong, you'll need to continue reading this chapter.

Science certainly has its detractors. Some horrendous things have been done in the name of science, been "justified" on scientific grounds, or, at a minimum, been made possible by science.

Although we should never close our eyes to the harm that's sometimes done with science, we believe it's as much a mistake to blame science for what some scientists have done in its name as it is to blame religion for what some believers have done in its name.

But what is science? As we'll see, science is, first and foremost, a method. However, it's also a culture. Some of the negative views of science come from what people perceive the culture of science to be—cold, calculating, self-assured, arrogant, and, perhaps, even offensive. We believe, however, that these perceptions are mistaken when the culture of science is at its best. The scientific method is, at its very core, a critical method, and those reflective individuals who use it are much more likely to be humbled than emboldened. In *The Logic of Scientific Discovery*, Sir Karl Popper ([1959] 2003) reminds us that science isn't a static set of beliefs to be conserved and that all knowledge is tentative. As Socrates points out in Plato's *Apology*, an acute awareness of our own ignorance is always the first step toward knowledge. Part of the culture of science is the willingness to put our ideas to the test. Science isn't about certainty, it isn't merely about the orderly collection of facts, and it isn't about invoking authority to protect our ideas from uncomfortable evidence. Instead, science is about asking tough questions and providing answers that invite criticism. Science is about recognizing the limits of our knowledge without lapsing into irresponsible cynicism. And science is about using the best logic, methods, and evidence available to provide answers today, even though we recognize that they may be overturned tomorrow.

Comparative politics is a subfield of political science. But what exactly is political science? Well, it's the study of politics in a scientific way. It's easy to see that, as it stands, this definition isn't particularly informative. For example, *What is politics?* And *What is science?* In the next chapter, we answer the first of these questions and seek to demarcate politics from other forms of social phenomena. In this chapter, though, we focus on the second question—*What is science?*

WHAT IS SCIENCE?

Is science simply a body of knowledge or a collection of facts, as many of us learn in high school? While there was a time when many scientists may have defined science in this way, this definition is fundamentally unsatisfactory. If this definition of science were accurate, many of the claims about how the universe works, such as those developed through Newtonian physics, would now have to be called unscientific, because they've been replaced by claims based on more recent theories, such as Einstein's theory of relativity. Moreover, if science were simply a collection of statements about how the world works, then we wouldn't be able to appeal to science to justify our knowledge of the world without falling into the following circular reasoning:

"Science is a collection of statements about how the world works."

"How do we know if these statements are accurate?"

"Well, of course they're accurate! They're scientific!"

The body of knowledge we call "scientific" may well be a product of science, but it isn't science itself. Rather, science is a method for provisionally understanding the world. The reason for saying "provisionally" will become clear shortly. Science is one answer to the central question in epistemology (the study of knowledge): "How do we know what we know?" The scientist's answer to that question is, "Because we have subjected our ideas to the scientific method." Science is a quest for knowledge. At this point, you might say that there are many ways to seek knowledge. Does this mean that meditation, reading scripture, and gazing at sunsets are all scientific activities? Although we agree that these are all ways of seeking knowledge, none of them is scientific. Science is a particular quest for knowledge. It's a pursuit of knowledge in which the scientist continually subjects their ideas to the cold light of logic and evidence.

Although science isn't the only route to knowledge, it may be unique in its emphasis on self-criticism. Scientists, like other scholars, can derive their propositions from an infinite number of sources. For example, Gregory Derry (1999) tells the story of how August Kekulé made an extremely important scientific breakthrough while hallucinating—half asleep—in front of the fireplace in his laboratory one night. He'd spent days struggling to understand the spatial arrangement of atoms in a benzene molecule. In a state of mental and physical exhaustion, his answer appeared to him as he "saw" swirls of atoms joined in a particular formation dancing among the embers of his fireplace. In a flash of inspiration, he saw how the pieces of the puzzle with which he had been struggling fit together. This inspired understanding of the physical properties of organic compounds didn't become a part of science that night, though. It did so only after the implications of his vision had withstood the critical and sober onslaught that came with the light of day. Thus, although flashes of insight can come from a variety of sources, science begins only when we ask, "If that's true, what else ought to be true?" And it ends—if ever—when researchers are satisfied they've taken every reasonable pain to show that the implications of the insight are false and have failed to do so. Even then, however, the best answer isn't the final answer—it's just the best "so far."

So, science is the quest for knowledge that relies on criticism. The thing that allows for criticism is the possibility that our claims, theories, hypotheses, ideas, and the like could be wrong. Thus, what distinguishes science from "nonscience" is that scientific statements must be falsifiable—there must be some imaginable observation or set of observations that could falsify or refute them. This doesn't mean that a scientific statement will ever be falsified, just that there must be a possibility that it could be falsified if the "right" observation came along. Only if a statement is potentially testable is it scientific. We deliberately say "potentially testable" because a statement doesn't have to have been tested to be scientific. All that's required is that we can conceive of a way to test it.

What sorts of statements aren't falsifiable? Tautologies aren't falsifiable because they're true by definition. For example, the statement "Triangles have three sides" is a tautology. It's simply not possible to ever observe a triangle that doesn't have three sides because *by definition* if an object doesn't have three sides, it's not a triangle. It's easy to see that this statement isn't testable and hence unscientific. Tautologies, though, aren't always so easy to spot. Consider the following statement: "All good students get high grades." Is this a tautology? This statement may be true, but unless we can think of a way to identify good students without referring to their grades, then it's just a definition and is, therefore, unscientific. In other words, whether this statement is scientific depends on how we define *good students*.

Other statements or hypotheses aren't falsifiable, not because they're tautological, but because they refer to inherently unobservable phenomena. For example, the claims "God exists" and "God created the world" aren't falsifiable because they can't be tested. As a result, they're unscientific. Note that these claims may well be true, but it's important to recognize that science has nothing to do with the truth or falsity of statements. All that's required for a statement to be scientific is that it be falsifiable. It should be clear from this that we're not claiming that "non-science" is nonsense or that it lacks meaning. This would clearly be a mistake. Nonfalsifiable statements like "God exists" may very well be true and have important and meaningful consequences. Our claim is simply that they don't form a part of science. Having defined *science* as a critical method for learning about the world, we can now evaluate the basic elements of the scientific method in more detail.

THE SCIENTIFIC METHOD

Although there's no scientific method clearly written down that's followed by all scientists, it's possible to characterize the basic features of the scientific method in the following manner.

Step 1: Question

The first step in the scientific process is to observe the world and come up with a question or puzzle. The very need for a theory or explanation begins when we observe something that's so unexpected or surprising that we ask, "Why did that occur?" Note that the surprise that greets such an observation, and that makes the observation a puzzle worth exploring, implies that the observation doesn't match some prior expectation or theory we held about how the world works. Thus, we always have a preexisting theory or expectation when we observe the world. If we didn't have one, we could never be surprised, and there would be no puzzles.

Step 2: Theory or Model

Once we've observed something puzzling, the next step is to come up with a theory or model to explain it. In what follows, we'll talk of theories, models, and explanations interchangeably. Scientists use the word *theory* to describe a set of logically consistent statements that tell us why the things we observe occur. It's important that these statements be logically consistent. Otherwise we have no way of determining what their empirical predictions will be and, hence, no way to test them. Put differently, theories that are logically inconsistent should not, indeed cannot, be tested, because we have no way of knowing what observations would truly falsify them.

What should theories or models look like? It's useful to think of our starting puzzle or observation as the end result of some previously unknown process (Lave and March 1975). We can then speculate about what (hidden) processes might have produced such a result. In effect, we try to imagine a prior world that, if it had existed, would have produced the otherwise puzzling observation before us. This prior world then becomes our model explaining the observation.

Notice that this process of imagining prior worlds is one place—but surely not the only one—where imagination and creativity enter the scientific process. What scientists do to stimulate this creative process is itself not part of the scientific method. Essentially, anything goes. Whatever means we use to stimulate speculation about a prior world, if we can show through logical deduction that *if* that prior world existed, it would have produced the puzzling observation we started with, then we have a theory, or model. Note that we have only *a* theory. We don't necessarily have *the* theory. This is why we continually test the implications of our theory.

The model that we end up with will necessarily be a simplified picture of the world. It's impossible to have a descriptively accurate model of the world as an infinite number of details would have to be captured in such a model. Pure description is impossible. Models are always going to leave many things out. As with all arts, much of the skill of modeling is in deciding what to leave out and what to keep in. A good model contains only what's needed to explain the phenomenon that puzzles us and nothing else. If we made our models too complex, we'd have no way of knowing which elements were crucial for explaining our puzzling observation and which were superfluous. The purpose of a model isn't to describe the world but to explain it, so descriptive accuracy isn't a core value in model building. Details are important only to the extent that they're crucial to what we're trying to explain. For example, if we're interested in explaining an aircraft's response to turbulence, it isn't important whether our model of the aircraft includes TV screens on the back of the passengers' seats. In fact, such inconsequential details can easily distract our attention from the question at hand. Another benefit of simple models is that they invite falsification because they make it very clear what we shouldn't observe. The more amendments and conditions placed on an explanation, the easier it is for scholars to dismiss apparently contradictory evidence.

It's important to remember that models are always developed with a specific goal in mind. This means we should evaluate models in terms of how useful they are for achieving that goal. As the Dutch economist Henri Theil (1971) once said, "models should be used, not believed." To emphasize this point, it can be helpful to think of models as being similar to maps. Like models, maps are simplified pictures of the world designed for a specific purpose. Consider the subway map of any city. The subway map is always a simplification of the city and, indeed, an inaccurate simplification in the sense that it provides inaccurate information about the relative distances between, and geographic positions of, particular locations. Despite this, the map is incredibly useful if our goal is to move efficiently around the city using the subway system—the purpose for which the map was designed. Of course, this map would be less useful if our goal was to walk above ground from one location to another. As with a map, we mustn't judge the value of a model in some abstract sense but in terms of how well it helps us understand some particular aspect of the world and explain it to others.

Step 3: Implications (Hypotheses)

Once we have a model, the third step in the scientific process is to deduce implications from the model other than those we initially set out to explain. Why do we say "other than those we initially set out to explain"? Well, presumably the model we construct will provide a logical explanation for the puzzling observation with which we started. After all, that's what it was designed to do! In other words, there's no way that a model can ever be falsified if only the observations that were employed to develop the model in the first place are used to test it. To actually test the model and allow for the possibility that it'll be falsified, we'll have to find other implications that can be deduced from it. We must ask ourselves, "If the prior world we created to explain the phenomena we originally found puzzling really did exist, what else ought to exist? What else should we be able to observe?" As before, there's often room for incredible imagination here, because the complete list of logical implications of a model is seldom self-evident.

Good models are those that produce many different implications. This is so because each prediction represents another opportunity for the model to fail and, therefore, makes the model easier to falsify. This is good because if the model fails to be falsified, we gain more confidence in its usefulness. Fertile models—models with many implications—are also desirable because they encourage the synthesis of knowledge by encouraging us to see connections between ostensibly disparate events. Good models also produce surprising implications. They tell us something we wouldn't know in the absence of the model. Models aren't particularly useful if they tell us only what we already know. Surprise, however, is best appreciated in small doses. If every implication of a model is surprising, either everything we thought about the world is wrong, or the model is.

Step 4: Observe the World (Test Hypotheses)

The fourth step is to examine whether the implications of the model are consistent with observation. Remember that the goal isn't to dogmatically uphold the implications of our model or defend them in order to prove how right they are. On the contrary, we should try our best to falsify them, because it's only after a theory has withstood these attempts to overthrow it that we can reasonably start to have confidence in it. While we should test as many implications as possible, testing those that are most likely to be falsified is particularly important. Always submit a model to the harshest test you can devise.

It's standard practice to stop and ask if other models—models that describe altogether different processes—might also explain the phenomena of interest. When this is the case (and it almost always is), it's incumbent upon scientists to compare the implications of those other models with the implications of their own model. Although it's always the case that competing models have some of the same implications (otherwise they couldn't explain the same observations to begin with), it's typically the case that they'll differ in some of their implications (otherwise they're not different models). The trick for a researcher is to identify these points of conflict between the different models and identify the relevant observations in the real world that would help them decide between them. This is what scientists refer to as a critical test. Ultimately, if a critical test is possible, observation will prove decisive in choosing between the models. This is because we know that there's only one world and the creative scientist has managed to get competing theories to say contradictory things about it. Only one of the models can be consistent with the real world.

Step 5: Evaluation

If we observe the implications deduced from our theory, we say that our theory has been corroborated. We can't say our theory has been verified or proven. This important point is one we'll return to in more detail in the next section.[1] That we can never prove a scientific explanation is why we earlier called science a method for "provisionally" understanding the world. Our theory may or may not be true. All we can conclude if observations are consistent with our theoretical implications is that our theory hasn't yet been falsified. We can't rule out that it won't be falsified the next time it's tested. As you can see, the scientific method is an inherently critical method when it's "successful" (when a theory's predictions seem to be borne out), because it's precisely under these circumstances that it's most cautious in the claims it makes.

BOX 2.1 AN EXAMPLE OF THE SCIENTIFIC PROCESS

THE CASE OF SMART FEMALE ATHLETES

Because student athletes often miss classes to compete, they frequently submit a letter from the athletic director asking for cooperation from their professors. Over the years, a certain professor has noticed through casual observation that women engaged in athletic competition frequently perform better academically than the average student. It's puzzling why female athletes would perform better despite missing classes. Can you think of a model—a process—that might produce such a puzzling observation?

You might start with the following conjecture:

- Female athletes are smart.

This is an explanation, but it's not a particularly good one. For example, it comes very close to simply restating the observation to be explained. One thing that could improve the explanation is to make it more general. This might lead you to a new explanation:

- Athletes are smart.

This model is certainly more general (but not necessarily more correct). Still, there are at least two problems with this model as things stand. First, it has no sense of process. It basically says that athletes share some inherent quality of smartness that leads them to perform better academically. In effect, this only pushes the phenomenon to be explained back one step; that is, we now need to know why athletes are smart. Second, the model comes close to being a tautology. It essentially says that athletes perform better academically because they're defined as being smart. This is problematic, as we saw earlier, because tautologies aren't falsifiable—they can't be tested; hence, they're not part of the scientific endeavor.

[1] Many scientists, however, slip into the language of verification when reporting their results. Instead of simply saying that their test has failed to falsify their hypotheses or is consistent with their theory, they'll claim that the test has shown that their theory is correct. For example, they might claim that their test shows that wealth causes democracies to live longer when, in fact, all they can conclude is that they were unable to falsify or reject the claim that wealth causes democracies to live longer.

This might lead you to look for a new explanation or model that includes some sort of process that makes female athletes appear smart. You might come up with the following model:

- Being a good athlete requires a lot of hard work; performing well academically in college requires a lot of work. Students who develop a strong work ethic in athletics are able to translate this to their studies.

This is a much more satisfying model because it provides a process or mechanism explaining why female athletes might be more academically successful than other students. An appealing feature of the model is that the logic of the argument applies not only to female athletes but to any athlete. Indeed, it applies to any person involved in an activity that rewards hard work. Thus, we might generalize this model by removing the specific reference to athletes:

- **Work Ethic Theory:** Some activities provide a clear, immediate, and tangible reward for hard work—in fact, they may provide an external stimulus to work hard (coaches shouting through bullhorns, manipulating rewards and punishments based on effort, and so on). Individuals who engage in these activities develop a habit of working hard and so will be successful in other areas of life as well.

At this point, you should stop and ask yourself whether there are any alternative explanations for why female athletes are successful. Can you think of any? One alternative explanation is the following:

- **Excellence Theory:** Everyone wants to feel successful, but some people go long periods without success and become discouraged. Those individuals who experience success in one area of their life (perhaps based on talent, rather than hard work) develop a "taste" for it and devise strategies to be successful in other parts of their life. Anyone who achieves success in nonacademic areas, such as athletics, will be more motivated to succeed in class.

Another alternative explanation is the following:

- **Gender Theory:** In many social and academic settings, women are treated differently from men. This differential treatment often leads women to draw inferences that certain activities are "not for them." Because many athletic endeavors are sex specific, they provide an environment for women to develop their potential free from the stultifying effects of gender bias. The resulting sense of efficacy and autonomy encourages success when these women return to gendered environments like the classroom.

We now have three different or competing models, all of which explain the puzzling observation we started with. But how can one evaluate which model is best? One way is to test some of the implications that can be derived from these theories. In particular, we'd like to find some new question(s) to which the three models give different answers. In other words, we'd like to conduct a critical test that would allow us to choose among the alternative reasonable models.

We might start by wondering whether being an athlete helps the academic performance of women more than men. Whereas the Work Ethic Theory and the Excellence Theory both predict that being an athlete will help men and women equally, the Gender Theory predicts that female athletes will perform better than nonathletic women but that male athletes will have no advantage over nonathletic men. Thus, collecting information on how well male and female athletes perform in class relative to male and female nonathletes, respectively, would allow us to distinguish between the Gender Theory and the other theories.

But how can we distinguish between the Excellence Theory and the Work Ethic Theory? One difficulty frequently encountered when trying to devise critical tests is that alternative theories don't always produce clearly differentiated predictions. For example, we just saw that the Excellence Theory and the Work Ethic Theory both predict that athletics will help men and women academically. It turns out that these two theories have other predictions in common as well. The Excellence Theory clearly suggests that success in any nonacademic area of life is likely to encourage academic success. In other words, the Excellence Theory predicts that academic success will be associated with success in other areas of life. The problem is that success in many of these nonacademic areas may require hard work. As a result, if we observe, for instance, accomplished musicians performing well in our political science classes, it will be difficult to discern whether this is because they learned the value of hard work in music and transferred it to political science (Work Ethic Theory) or because they developed a "taste" for success as musicians that then inspired success in political science (Excellence Theory). In effect, the Excellence Theory and the Work Ethic Theory both predict that academic success will be associated with success in other areas of life.

If we want to distinguish between the Work Ethic Theory and the Excellence Theory, we need to imagine observations in which they produce different expectations. Sometimes, this requires further development of a theory. For example, we might expand the Excellence Theory to say that those people who develop a taste for excellence also develop a more competitive spirit. If this is true, then the Excellence Theory would predict that student athletes are likely to be more competitive and will perform better than other students even when playing relatively frivolous board games. Since even the most driven athletes are not likely to devote time to training for board games, the Work Ethic Theory predicts that athletes will perform the same as nonathletes in such trivial pursuits. Thus, we could look at the performance of athletes and nonathletes at board games to distinguish between the Excellence Theory and the Work Ethic Theory.

The three critical tests we've come up with and their predictions are listed in Table 2.1. All that's now required is to collect the appropriate data and decide which model, if any, is best.

TABLE 2.1 ■ Three Critical Tests

	Theory		
Question	**Gender**	**Excellence**	**Work ethic**
Will athletics help women more than men?	Yes	No	No
Is academic success associated with success in other areas of life?	No	Yes	Yes
Are female athletes more successful at board games than women who are not athletes?	Yes	Yes	No

It's worth noting that there's considerable overlap between the predictions of our three theories. This is often the case in political science settings as well. The crucial point isn't that each theory should yield a complete set of unique predictions, but that our theories should have sufficiently many distinct predictions that we can use observation to help us make decisions about which theories to embrace, however tentatively. Table 2.1 lists just some of the predictions that might help us to distinguish between the three theories outlined above. Can you think of others?

Although we can't ever prove our theories, we can claim that some theories are better corroborated than others. As a result, we can have more confidence in their conclusions. One might think that a theory that's been subjected to multiple tests is better corroborated than one that hasn't been subjected to many tests at all. However, this isn't always the case. If we keep testing the same implication over and over again, it's not clear how much an additional test actually adds to the degree to which the theory is corroborated. What really matters isn't so much how many times a theory has been corroborated, but the severity and variety of the tests to which it has been subjected. This, in turn, will depend on the degree to which the theory is falsifiable. Again, this is why we like our models to be simple and have multiple implications. In general, we'll have more confidence in a theory that's survived a few harsh tests than a theory that's survived many easy ones. This is why scientists often talk about the world as if it were black-and-white rather than gray. Bold statements should be interpreted not as scientific hubris but rather as attempts to invite criticism—they're easier to falsify.

What happens if we don't observe the implications deduced from our theory? Can we conclude that our theory is incorrect based on one observation? The answer is "probably not." It's entirely possible that we haven't observed and measured the world without error. Moreover, if we believe that human behavior is inherently probabilistic, we might not want to reject theories on the basis of a single observation. In a world in which our tests are potentially fallible, we shouldn't relegate a theory to the dustbin of intellectual history the minute one of its implications is shown to be false. Instead, we must weigh the number, severity, and quality of the tests that the theory's implications are subjected to and make a judgment. And most important, this judgment should be made with an eye toward what would replace the theory should we decide to discard it. This is why some scientists say that it takes a theory to kill a theory. Further, if we do embrace a new theory and disregard an alternative, it should be because the new theory is more consistent with all of the implications of both theories. Developing a new theory that explains the facts that the old theory found inconvenient without also explaining the many facts that the old theory accurately predicted is called *ad hoc* explanation. Because this practice doesn't expose the new theory to falsification as strenuously as it does the old theory, it's not consistent with sound scientific practice.

AN INTRODUCTION TO LOGIC

In the previous section, we talked in a rather casual way about constructing and testing scientific explanations. To better appreciate the important connection between theory construction and theory testing, it's useful to devote some time to the study of logic. The study of logic is first and foremost about learning to be careful about how we construct and evaluate arguments.

Throughout our lives, we're confronted by people trying to convince us of certain things through arguments. Politicians make arguments as to why we should vote for their party rather than the party of their opponents. National leaders provide arguments for why certain policies should be implemented or abandoned. Lawyers make arguments as to why certain individuals should be found guilty or innocent. Professors make arguments as to why students should spend more time in class rather than at parties. It's important for you to know when these arguments are logically valid and when they're not. If you can't distinguish between a valid and an invalid argument, other people will be able to manipulate and exploit you. You'll be one of life's suckers. In this section, we give you some tools to determine whether an argument is valid or not.

Valid and Invalid Arguments

What's an argument? An argument is a set of logically connected statements, typically in the form of a set of premises and a conclusion. An argument is valid when accepting its premises compels us to accept its conclusions. An argument is invalid if, when we accept the premises of an argument, we're free to accept or reject its conclusions. One way to represent an argument is in the form of a categorical syllogism that consists of a major premise, a minor premise, and a conclusion. The major premise is typically presented as a conditional statement, such as "If *P*, then *Q*." The "if" part of the conditional statement (in this case "If *P*") is called the *antecedent*, whereas the "then" part of it (in this case "then *Q*") is called the *consequent.* An example of a conditional statement is "If a country is wealthy [antecedent], then it will be a democracy [consequent]." The minor premise consists of a claim about either the antecedent or the consequent in the conditional statement (major premise). The conclusion is a claim that's thought to be supported by the premises.

Four types of conditional argument can be represented with a syllogism—arguments that affirm or deny the antecedent and those that affirm or deny the consequent. Which of these four types of argument are valid, and which are invalid? Recall that a valid argument is one such that if you accept that the premises are true, then you're compelled to accept the conclusion as true. Let's start by considering what happens when we affirm the antecedent. An example is shown in Table 2.2.

The major premise states, "If *P* is true, then *Q* must be true." The minor premise says that "*P* is true." Together, these premises compel us to accept that the conclusion is true. As a result, the argument is valid. In other words, the major premise states, "If a country is wealthy [antecedent], then it will be a democracy [consequent]." The minor premise says, "The observed country is wealthy." It logically follows from this that the observed country must be a democracy. You may be able to immediately see why this argument is valid simply by reading it. If you can't, though, don't worry—we find that many students struggle to evaluate the validity of an argument simply by reading it. And besides, trusting our intuition about which arguments seem valid can often lead us astray. An alternative approach that's often useful for evaluating the validity of an argument like this one is to look at it graphically in set-theoretic

TABLE 2.2 ■ Affirming the Antecedent: A Valid Argument

	General Form	Specific Example
Major premise	If *P*, then *Q*	If a country is wealthy, then it will be a democracy.
Minor premise	*P*	The country is wealthy.
Conclusion	Therefore, *Q*	Therefore, the country will be a democracy.

form. In Figure 2.1, we show the general form of a categorical syllogism in set-theoretic form. The major premise "If *P* is true, then *Q* must be true" indicates that the set of cases where *P* occurs is a subset of the cases where *Q* occurs. This is captured graphically in Figure 2.1 by the way that we've drawn a small "P" circle inside the larger "Q" oval. As we'll see, we can use Figure 2.1 to evaluate the logical validity of all four of our categorical syllogisms.

Recall that we're currently interested in evaluating the logical validity of a categorical syllogism that affirms the antecedent like the one shown in Table 2.2. The type of categorical syllogism we're dealing with is given by the minor premise. In our case, the minor premise states that "P is true." Thus, we're affirming the antecedent. How does Figure 2.1 help us evaluate the validity of this type of argument? We're supposed to accept the premises of an argument when determining if the argument is valid or not. Therefore, we must accept the minor premise that "P is true." This is equivalent to accepting that we're located somewhere in the P circle in Figure 2.1. You can now see that the minor premise in a categorical syllogism tells us where we're located in Figure 2.1. As Table 2.2 shows, the conclusion we're being asked to accept in an argument affirming the antecedent is that Q must be the case. Are we forced to accept this conclusion? The answer is yes. We can see from Figure 2.1 that if we're located in the P circle, then we must accept that we're also in Q. This because the entire P circle is inside the Q oval. In other words, if P is the case, then Q must be the case as well. We hope it's now clear why accepting the minor premise affirming P compels us to conclude Q and thus why a categorical syllogism that involves affirming the antecedent like the one in Table 2.2 is a valid form of argument.

Now let's consider what happens when we deny the antecedent. An example is shown in Table 2.3. We know we're dealing with a categorical syllogism that involves denying the antecedent because the minor premise states "Not P." As we've seen, the major premise "If *P*, then *Q*" can be

FIGURE 2.1 ■ Major Premise: If P, Then Q

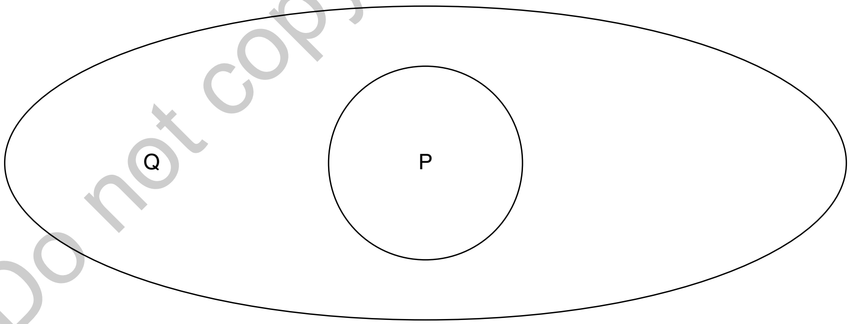

TABLE 2.3 ■ Denying the Antecedent: An Invalid Argument

	General Form	Specific Example
Major premise	If *P*, then *Q*	If a country is wealthy, then it will be a democracy.
Minor premise	Not *P*	The country is not wealthy.
Conclusion	Therefore, not *Q*	Therefore, the country will not be a democracy.

represented in set-theoretic terms by Figure 2.1. The difference from the previous example is that the minor premise now asserts that *P* isn't the case. In other words, we're not located in the P circle in Figure 2.1. If we accept this, as we have to do because this is a premise of the argument, are we compelled to accept the conclusion that *Q* is not the case? The answer is no. Figure 2.1 clearly illustrates that not being in the P circle doesn't preclude us from being in the Q oval. It's certainly possible to be in Q without being in P. As a result, it doesn't logically follow from observing "not *P*" that *Q* isn't the case. Therefore, this is an invalid argument. This is because we can contradict the conclusion (not *Q*) without running into a contradiction with either the major premise or the minor premise. Since a valid argument compels us to accept its conclusion if we accept the premises, this is sufficient to demonstrate that arguments denying the antecedent are invalid.

Let's briefly look at the specific example of denying the antecedent in Table 2.2 that's written in words instead of in terms of P and Q. Does it follow from the fact that the observed country isn't wealthy that it won't be a democracy as the argument claims? Intuitively, we can imagine that there may be reasons why a country is a democracy even though it isn't wealthy. Indeed, one example of a nonwealthy democracy is India. An important point here, though, is that the argument is invalid, not because we can come up with an example of a real democracy that isn't wealthy (India), but rather because we're not compelled to accept the conclusion based on accepting the major and minor premises. It may be confusing for readers that there's no direct connection between the factual accuracy of an argument's conclusion and the validity of the argument itself—a valid argument can have a conclusion that's factually false and an invalid argument can have a conclusion that's factually true. If we restrict our attention only to whether the argument is valid as it applies to our democracy example, we must ask, "Does the major premise claim that wealth is the only reason why a country will be a democracy?" The answer is clearly no. The major premise states only what will happen if a country is wealthy. It makes no claim as to what might happen if a country isn't wealthy. It's for this reason, and this reason alone, that the argument is invalid.

Now let's consider what happens when we affirm the consequent. An example is shown in Table 2.4. We know we're dealing with a categorical syllogism that involves affirming the consequent because the minor premise states "Q." As we've seen, the major premise "If P, then Q" can be represented in set-theoretic terms by Figure 2.1. The difference from the previous examples is that the minor premise now asserts that *Q* is the case; that is, it affirms the consequent. In other words, we're located somewhere in the Q oval in Figure 2.1. If we accept that the premises are true, are we compelled to accept the conclusion that *P* is the case? The answer is no. Figure 2.1 clearly illustrates we can be located in Q without also being located in P. As a result, the argument is invalid—we're not compelled to accept the conclusion based on the premises.

TABLE 2.4 ■ Affirming the Consequent: An Invalid Argument I

	General Form	Specific Example
Major premise	If *P*, then *Q*	If a country is wealthy, then it will be a democracy.
Minor premise	*Q*	The country is a democracy.
Conclusion	Therefore, *P*	Therefore, the country is wealthy.

Let's briefly look at the specific example of affirming the consequent in Table 2.4 that's written in words instead of in terms of P and Q. Does it follow from the fact that the observed country is a democracy that it will be wealthy? An argument that affirms the consequent confuses necessity and sufficiency. Although the major premise states that wealth is sufficient for democracy—wealthy countries will be democracies—it doesn't assert that wealth is necessary for democracy. In other words, the major premise doesn't state that wealth is the only cause of a country's democracy. Consequently, we can't conclude from the fact that a country is a democracy that it's also wealthy—it may be wealthy or it may not be. Recall that to show that an argument is invalid, it's not necessary to show that its conclusion is false. All we have to show is that it doesn't have to be true.

Finally, let's consider what happens when we deny the consequent. An example is shown in Table 2.5. We know we're dealing with a categorical syllogism that involves denying the consequent because the minor premise states "Not Q." As we've seen, the major premise "If P, then Q" can be represented in set-theoretic terms by Figure 2.1. The difference this time is that the minor premise now asserts that Q isn't the case. In other words, we're not located in the Q oval in Figure 2.1. If we accept that the premises are true, are we compelled to accept the conclusion that P isn't the case? The answer is yes. The P circle is entirely inside the Q oval in Figure 2.1. Thus, accepting the premise that we're not in *Q* necessarily means that we can't be in *P*. This is a valid argument—we're compelled to accept the conclusion based on the premises. In the context of the specific example of denying the consequent in Table 2.5 that's written in words, the major premise indicates that all wealthy countries are democracies and the minor premise states that the country isn't a democratic one. If the premises are both true, then it logically follows that our country can't be wealthy.

Our brief foray into the study of logic indicates that if complex arguments can be broken down into categorical syllogisms, then it's possible to classify all arguments into one of four types according to whether they affirm or deny the consequent or antecedent. As we've seen, two of these arguments are valid and two of them are invalid. Specifically, affirming the antecedent and denying the consequent are valid arguments—if you accept the major and minor premises, you're compelled to accept the conclusion. In contrast, denying the antecedent and affirming the consequent are invalid arguments—if you accept the major and minor premises, you're not compelled to accept the conclusion. These results are summarized in Table 2.6.

TABLE 2.5 ■ Denying the Consequent: A Valid Argument I

	General Form	Specific Example
Major premise	If *P*, then *Q*	If a country is wealthy, then it will be a democracy.
Minor premise	Not *Q*	The country is not a democracy.
Conclusion	Therefore, not *P*	Therefore, the country is not wealthy.

TABLE 2.6 ■ What Types of Conditional Arguments Are Valid?

	Antecedent	Consequent
Affirm	Valid	Invalid
Deny	Invalid	Valid

Testing Theories

We think it's important for you to be able to distinguish between valid and invalid arguments so that you're not manipulated or exploited by others. However, this brief introduction to logic is also important because it tells us something about the way that scientists test their theories and explanations. Suppose we want to explain why rich countries are much more likely to be democracies than poor countries. One possible explanation for why this might be the case is given in the following statements[2]:

1. Living in a dictatorship is risky—if you're one of the dictator's friends, you'll do extremely well; but if you're not, you'll do extremely poorly.
2. Living in a democracy is less risky—democratic leaders have to spread the goodies (and the pain) around more evenly. This means you're less likely to do extremely well or extremely poorly in a democracy.
3. Rich people are less likely to take risks than poor people because they have more to lose. This means that countries with lots of rich people are more likely to be democracies than dictatorships.

This short explanation provides reasons why rich countries might be more likely to be democracies than poor countries. How good is this explanation, though? Does this argument have any testable implications? One implication is that rich democracies should live longer than poor democracies. This is because people in rich democracies should be less likely to take the "risk" of becoming a dictatorship. In contrast, people in poor democracies might wonder what they have to lose.

How can we use observations of the real world to evaluate our proposed explanation? It's often the case that the implications of an explanation are more readily observable than the elements of the explanation itself. Consider the example we're using. Although it may be possible to compare the distribution of good and bad outcomes in dictatorships and democracies, the claims that people differ in their propensity to take risks and that this propensity is related to their level of income are difficult to observe. This is because the propensity to take risks is an internal and psychological attribute of individuals. For similar reasons, scholars typically evaluate their explanations by observing the real world to see if the implications of their explanations appear to be true based on the assumption, "If my theory is true, then its implications will be true." If we take this to be our major premise and the truth or falsity of the theory's implications as the minor premise, then we might be able to use observations to draw inferences about our theory or explanation.

Suppose our theory's implications were borne out by our observation that rich democracies live longer than poor democracies. Can we conclude that our theory is true? If we were to do so, we'd be engaging in reasoning that affirmed the consequent. This fact is shown more clearly in

[2] This is a simplified version of an argument presented by Adam Przeworski (2001). It will be discussed more fully in Chapter 5.

Table 2.7. As you know by now, affirming the consequent is an invalid form of argument. The major premise says only that if the theory is correct, then the implications should be observed. It never says that the only way for these implications to be produced is if the theory is correct. In other words, processes other than those described in our theory may produce the observation that rich countries live longer than poor countries. Put differently, the mere fact of observing the predicted implication doesn't allow us to categorically accept or reject our theory.

Suppose now that our observations did not bear out our theory's implications; that is, we didn't observe that rich democracies live longer than poor democracies. Can we conclude that our theory is incorrect? If we were to do so, we'd be engaging in reasoning that denies the consequent. This fact is shown more clearly in Table 2.8. As you know by now, denying the consequent is a valid form of argument. In other words, by accepting the premises, we're compelled to accept the conclusion that our theory isn't correct.

If we compare the two previous examples, we can see an important asymmetry regarding the logical claims that can be made on the basis of "confirming" and "disconfirming" observations. When an implication of our theory is confirmed, the most we can say is that the theory may be correct. This is because neither of the two possible conclusions—our theory is correct or our theory is incorrect—contradicts our major and minor premises. In other words, we can't say that our theory is correct or verified. In contrast, if we find that an implication of our theory is

TABLE 2.7 ■ Affirming the Consequent: An Invalid Argument II

General Form	Example	Specific Example
If *P*, then *Q*	If our theory *T* is correct, then we should observe some implication *I*.	If our theory is correct, then we should observe that rich democracies live longer than poor democracies.
Q	We observe implication *I*.	Rich democracies live longer than poor democracies.
Therefore, *P*	Therefore, our theory *T* is correct.	Therefore, our theory is correct.

TABLE 2.8 ■ Denying the Consequent: A Valid Argument II

General Form	Example	Specific Example
If *P*, then *Q*	If our theory *T* is correct, then we should observe some implication *I*.	If our theory is correct, then we should observe that rich democracies live longer than poor democracies.
Not *Q*	We do not observe implication *I*	Rich democracies do not live longer than poor democracies.
Therefore, not *P*	Therefore, our theory *T* is incorrect.	Therefore, our theory is incorrect.

inconsistent with observation, then we're compelled by logic to accept that the theory's false—this is the only conclusion that's consistent with our observation. Thus, although we can know that a theory must be incorrect in light of a disconfirming case, all that we can say in light of a confirming case is that a theory may be correct (it may also be wrong). What does this mean? It means that we're logically justified in having more confidence when we reject a theory than when we don't. This, in turn, implies that the knowledge encapsulated in theories that haven't been rejected remains tentative and can never be proven for sure—scientific theories can never be proven. Even if we're utterly convinced that our major and minor premises are true, all that we can logically conclude from a confirming instance is that the theory hasn't yet been falsified.

This asymmetry between confirming and disconfirming cases led the philosopher of science Sir Karl Popper ([1959] 2003, 280–81) to conclude the following:

> The old scientific ideal of *episteme*—of absolutely certain, demonstrable knowledge—has proved to be an idol. The demand for scientific objectivity makes it inevitable that every scientific statement must remain *tentative for ever*. . . . With the idol of certainty . . . there falls one of the defenses of obscurantism which bar the way to scientific advance. For the worship of this idol hampers not only the boldness of our questions, but also the rigor and integrity of our tests. The wrong view of science betrays itself in the craving to be right; for it is not his *possession* of knowledge, of irrefutable truth, that makes the man of science, but his persistent and recklessly critical *quest* for truth.

If confirming observations don't prove that our theory is correct, does this mean they're of no use whatsoever? The answer is no. Imagine that we start with a set of implications derived from a theory and then observe some facts. In other words, let's start with the theory and then observe the world. If we do this, it's possible that our observations will contradict our theory. If it turns out that our observations are consistent with our theory, then we can have a greater measure of confidence in our theory because it withstood the very real chance of being falsified. We can't say that our theory's verified or confirmed, just that we have more confidence in it. If our observations are inconsistent with our theory, we can draw valid inferences about the truthfulness of our theory—we can conclude that it's wrong. This approach to doing science, which forms the basis of the scientific method described earlier, is called falsificationism. Falsificationism is an approach to science in which scientists generate or "deduce" testable hypotheses from theories designed to explain phenomena of interest. It emphasizes that scientific theories are constantly called into question and that their merit lies only in how well they stand up to rigorous testing. Falsificationism forms the basis for the view of science employed in this book.

The approach to science we've described here takes a clear stance in the debate between "deductive" and "inductive" approaches to learning. The deductive approach to learning formulates an expectation about what we ought to observe in light of a particular theory about the world and then sets out to see if our observations are consistent with that theory. The inductive approach to learning, on the other hand, starts with a set of observations and then tries to ascertain a pattern in the observations that can be used to generate an explanation for the observations. Induction is problematic because to be successful it must rest at some point on the fallacy of affirming the consequent—the fact that observation precedes theory construction means the theory is never exposed to potential falsification! Popper ([1959] 2003) suggests that, in fact, the biggest problem with induction isn't so much that it's wrong but that it's impossible. Observational facts don't just present themselves to observers. We always decide which facts to pay attention to and which to ignore. As we noted earlier, the hunch that tells us what to observe and what to ignore, that is, what constitutes a puzzle worth explaining, constitutes a theory. In this respect, scholars who claim to be engaged in an inductive inquiry are actually engaged in an implicit deductive endeavor. If it's true that we're "all deductivists" as Popper claims, then the argument for deduction amounts to a claim that it's better to use theory explicitly than to use it implicitly.

Having described the scientific method, we'd like to briefly dispel certain myths that have developed about science. Some of these myths have been promoted by opponents of the scientific project, but others, unfortunately, have been sustained by scientists themselves.

MYTHS ABOUT SCIENCE

The first myth is that science proves things and leads to certain and verifiable truth. This isn't the best way to think about science. It should be clear by now from our discussion that the best science can hope to offer are tentative statements about what seems reasonable considering the best available logic and evidence. It may be frustrating for you to realize this, but science can speak with more confidence about what we don't know than what we do know. In this sense, the process of scientific accumulation can be thought of as the evolution of our ignorance. We use the scientific method because it's the best tool available to interrogate our beliefs about the (political) world. If we hold on to any beliefs about the (political) world, it's because, after we've subjected them to the most stringent tests we can come up with, they remain the most plausible explanations for the phenomena that concern us. Those who discuss the role of scientific expertise in policy debates related to things like climate change and COVID-19 sometimes seem to support the notion of science as "certain" and "verifiable" truth by declaring that we should "follow the science," as if "science" speaks for itself with unquestionable authority. In fact, *science* does nothing. *Scientists* speak with varying degrees of consensus, and they often change their mind. The fact that they do, shouldn't lead us to cynicism. Rather it should be taken as a sign that scientific progress may be taking place. Mature thinkers can, with humility, recognize when others possess expertise they lack, while at the same time reserving the right to remain skeptical. At the end of the day, science is about weighing the preponderance of evidence while keeping a critical eye open for opportunities to improve our understanding.

The second myth is that science can be done only when experimental manipulation is possible. This is clearly false. For theories to be scientific, they need only be falsifiable. There's no claim that the tests of these theories must be carried out in an experimental setting. Many of the natural sciences engage in research that isn't susceptible to manipulation. For example, all research on extinct animals, such as dinosaurs, must be conducted without the aid of experimental manipulation because the subjects are long dead. In fact, there's also no claim a theory must be tested before it can be called scientific. Albert Einstein presented a special theory of relativity in 1905 that stated, among other things, that space had to be curved, or warped. It took fourteen years before his theory was tested with the help of a solar eclipse. No scientist would claim that Einstein's theory was unscientific until it was tested. Put simply, scientific theories must be potentially testable, but this doesn't mean that they stop being scientific if they're yet to be actually tested.

The third myth is that scientists are value neutral. It's necessary here to distinguish between the method of science and the individuals—the scientists—who engage in science. The scientific method itself is value neutral. Science is simply a method that involves generating and evaluating logically consistent sets of falsifiable statements about the world. Scientists, though, may not be value neutral (Haraway 1988; Longino 1987). It's important to remember that the pursuit of knowledge about the world is closely entangled with attempts by people to change the world. As a result, the types of research questions that are asked and the interpretation of scientific results are likely to be infused with the specific values and biases held by individual scientists and those who use their research. The lack of diversity in most scientific disciplines, whether in terms of gender, sex, race, income, class, sexuality, religion, ethnicity, and so on, along with the power structure that exists in many societies, means that some research areas are less studied than others and that certain viewpoints are excluded or less privileged than others when it comes to interpreting scientific evidence (Carroll and Zerilli 1993; Collins 1986, 1989; D. E. Smith 1974). In effect, the knowledge that's produced by science is socially constructed. This is one of the many reasons for trying to promote the diversity of those involved in the scientific endeavor. The fact that scientists may not be value neutral means that we should be very clear about the limits of our knowledge and not encourage others to act upon knowledge that's not highly corroborated. Moreover, we should try to conduct our studies in such a way that someone who doesn't share our biases can determine if our arguments and evidence are reasonable.

It's been argued that science is predicated on two rules (Rauch 1993). First, no one gets the final say on any issue—all knowledge claims are, for the reasons outlined in this chapter, open to criticism. Second, no individual has a personal claim of authority about whether scientific statements are true or not. Taken together, these two rules create a social system that makes it possible that even though individual scientists will have biased perspectives, others, who hold different biases, will have incentives to check their work. As a community, scientists with many different biases will use the scientific method to check the claims that are being made in an attempt to reach a consensus that's independent of the biases held by individual scholars. This is yet another reason why having a diverse group of scientists is valuable.

The fourth myth, that politics can't be studied in a scientific manner, can easily be dispelled by now. Our description of the scientific method clearly shows that this myth is false. The study of politics generates falsifiable claims and hence generates scientific statements. The implications of these theories of politics can be tested just like the implications of any other scientific theory. We'll further demonstrate that politics can be studied in a scientific manner in the remaining chapters of this book. The fact, though, that our subjects can read our work and change their behavior makes our job quite a bit harder than if we were working in one of the natural sciences.

BOX 2.2 DIVERSITY AND SCIENCE

Debates about diversity often focus on issues of representation, fairness, equity, or social justice. Are different groups appropriately represented? Do some groups face discrimination? While these issues are important, we might also wonder whether increased diversity brings substantive benefits. In other words, does it lead to better performance or produce better outcomes? In terms of science, does it help us to better understand and predict things? The answer, in short, is that it can do, at least under certain conditions. In what follows, we look at the benefits (and costs) of diversity (Page 2007a, 2017).

One way to think about diversity is in terms of cognitive diversity or what we might call toolbox diversity. Loosely speaking, cognitive diversity has to do with the way that people think about, interpret, and interact with the world. Scott Page (2007a) argues that there are at least four different components of cognitive diversity. The first component has to do with someone's *perspective*. People who have different perspectives look at the world differently and therefore approach and represent problems differently. For example, one person might view criminal activity as primarily a security problem, whereas others might view it as more of an economic, social, or psychological problem.

The second component has to do with someone's *heuristics*. Heuristics refer to the methods, techniques, and routines that people use to solve problems and understand the world around them. Some people might approach new problems by using simple rules of thumb or by making analogies to problems they've solved in the past. Others might use more complicated algorithms or methods. People with distinct perspectives see problems differently (or see different problems), while people with distinct heuristics approach problem solving in different ways.

The third component has to do with someone's *interpretations*. Interpretations refer to how people take real world things and place them into categories. For example, someone who sees Barack Obama might place him in the category of African American men, whereas others might place him in the category of United States presidents or the category of Democrats. These categories are likely to lead people to think about Barack Obama in certain ways and hold particular expectations about him. Interpretations shape the predictions we make about how things work because they indicate the dimensions, attributes, or categories we think are important.

The fourth component has to do with someone's *predictive model*. Predictive models refer to how people think things fit together and therefore use their interpretations to make predictions. Diverse predictions can result from diverse interpretations and diverse predictive models. People with the same interpretations will make different predictions if they have different predictive models about how the world works.

Each of us "possesses all of these: perspectives, heuristics, interpretations, and predictive models. And each of us differs in the particular collection of these tools that we hold inside our heads" (Page 2007b, 12). We can think of the collection of these things as the cognitive repertoire or toolbox that each of us can bring to bear when we attempt to understand the world. Our own personal toolbox grows and changes over the course of our lives and is influenced by our preferences, our education and training, our life experiences and practice, and our culture. A diverse group in this sense is thus a group that contains people with different cognitive toolboxes.

How is cognitive diversity related to identity diversity? Identity diversity has to do with the different components of one's identity such as gender, sex, race, ethnicity, sexual orientation, class, religion, age, nationality, and physical qualities. It should be clear that cognitive diversity and identity diversity are distinct concepts and don't have to go together. Cognitive diversity doesn't necessarily imply identity diversity and vice versa. For example, a group may have considerable identity diversity in terms of gender, race, class, sexuality, age, and so on, but if the individuals all have similar life experiences, are trained in a similar way, or have come through the same educational or career "pipeline," then it's likely to have low levels of cognitive diversity. That said, empirical evidence suggests that identity diversity is often associated with cognitive diversity in practice.

Cognitive diversity is important for problem solving. In one study, Lu Hong and Scott Page (2004) present a mathematical theorem showing that groups of cognitively diverse problem solvers are likely to outperform groups of high-ability problem solvers in many settings. Groups comprised of the highest performing individuals often don't exhibit much cognitive diversity. They tend to be trained in a similar way or see the world in a similar way. As a result, they get stuck in the same place when trying to solve a problem. In more cognitively diverse groups, there's often an individual or two who, although less able when it comes to problem solving on their own, can provide a different perspective or heuristic that can get the group moving again or avoid getting stuck in the first place.

This won't always be the case. Certain conditions must be met for the "diversity trumps ability theorem" to hold. First, we must be faced with a difficult problem or task. Cognitive diversity isn't required if the problem or task is easy. Second, the group members must all be individually capable. In other words, ability still matters. Third, the "right type" of cognitive diversity is important. Including individuals trained as, say, anthropologists or medical doctors on a team working to solve a physics or engineering problem may increase cognitive diversity in some general sense but not in a way that's relevant for solving the team's problem. Fourth, the group must be drawn from a large population of people and be sufficiently large for cognitive diversity to be possible. Small groups can't be diverse. These conditions help to indicate when we can expect cognitive diversity to be beneficial.

Cognitive diversity is also important for making good predictions. Groups of people often make better predictions than individuals. This is why some people speak about the "wisdom of crowds" (Surowiecki 2004). While crowds aren't always better, it's the case that predictions from diverse groups are always better than predictions from similarly capable groups that are less diverse (Page 2007a, 2017). This is because the error in a crowd's prediction is equal to the average error of the people in the crowd minus the diversity in their individual predictions,

$$\text{Crowd Error} = \text{Average Error} - \text{Prediction Diversity}.$$

This is a mathematical identity and hence something that's always true; it's not a matter of belief. We might call this the "diversity prediction theorem." It tells us that adding more able individuals (to reduce average error) and more diverse individuals (to increase prediction diversity) can both improve a group's predictive accuracy.

While diversity can produce better performance, it can sometimes have negative consequences as well. This can be the case, for example, if it creates communication or coordination problems. People with drastically different perspectives, heuristics, interpretations, and predictive models may find it difficult to trust or communicate effectively with one another, thereby making problem solving more difficult. We also need to recognize that diverse groups may include individuals who hold conflicting preferences or values. We might refer to this as preference or value diversity. The studies we've discussed previously assume there's some true or correct answer out there in the real world just waiting to be found. In many situations, though, things are more complicated than this. What if there's no objectively correct truth? This can happen when our values or preferences come into play, such as when we're making policy about, say, abortion rights or choosing how to trade-off individual freedom and public health concerns during a pandemic. And even if we all agree on the goal at hand, such as, say, reducing child mortality, individuals may hold diverse views on the best way to achieve it. When people hold diverse preferences and values, conflict can occur and group decisions, along with the process by which those decisions are made, can become contested. This is something we'll look at in more detail in Chapter 9 when we investigate problems with group decision making. The costs associated with preference or value diversity are likely to be higher in some settings than others. For example, they're likely to be higher when we're trying to solve problems related to poverty and policing where our values and preferences come into play than when we're trying to find a medical cure for a disease or an explanation for dark matter. The bottom line is that cognitive diversity is almost always helpful, but value or preference diversity often isn't. As Scott Page (2007a, 300) puts it, "To think different is good. To want differently isn't, at least not necessarily."

What does this all mean for identity-based diversity and performance? Identity diversity helps performance to the extent that it produces (relevant) cognitive diversity. And it hinders performance to the extent that it produces (relevant) value or preference diversity. This suggests that identity diverse groups are likely to exhibit high variability in their performance depending on whether cognitive diversity and value diversity are high or low. In contrast, groups that exhibit less identity diversity are likely to see less variability in their performance.

These predictions, which are graphically presented in Figure 2.2, are largely born out in empirical research looking at the impact of identity diversity on group performance. Much of the debate in the literature is about whether the mean performance of diverse groups is better than that of non-diverse groups. Overall, the empirical evidence suggests that there may not be large differences in mean performance across these two types of groups (Page 2007a). Reflecting this, we've drawn Figure 2.2 assuming that the mean level of performance, μ, is the same across diverse and non-diverse groups. Importantly, the argument presented here indicates that high performance is much more likely in diverse groups. This is indicated by the fact that the dashed distribution for diverse groups in Figure 2.2 has "fatter" tails. If this is true and we're interested in solving big, complicated problems, as we are in science and politics, then this suggests that diverse groups may be our best bet for achieving this. Indeed, increased diversity may be our only way to solve some of these problems.

Another implication of the argument presented here is that if we can reduce value and preference diversity or, perhaps more plausibly, create systems, institutions, rules, or environments that allow us to work through our differences more peacefully and productively, then we can more effectively leverage the very real benefits of diversity while minimizing the potential costs.

FIGURE 2.2 ■ Probability Distribution of Performance for Non-Diverse and Diverse Groups

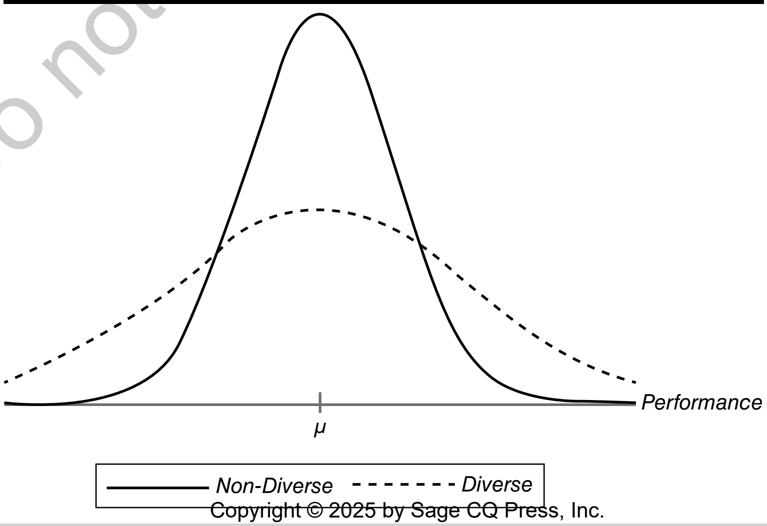

CONCLUSION

In this chapter, we've argued that it's useful to think about politics in a scientific manner. We've also tried to offer a clear view of what most practicing scientists have in mind when they use the word *science.* It's a fairly minimalist view. What unites all scientists is the idea that one ought to present one's ideas in a way that invites criticism and refutation (Popper 1962). It's incumbent upon the scientist to answer the question "What ought I to observe if what I claim to be true about the world is false?" This view of science recognizes that scientific knowledge is tentative and should be objective. Although it's certainly likely that our prejudices and biases motivate our work and will creep into our conclusions, the goal of science is to present our conclusions in a way that will make it easy for others to determine whether it's reasonable for people who don't share those prejudices and biases to view our conclusions as reasonable.