3

Explanation

EXPLANATION, THEORY, AND HYPOTHESIS

When we ask for an explanation, we could be asking for a number of things. If I arrive late for an appointment with you, for example, you might ask me to explain why I'm late. "Sorry, I got caught in traffic," I might reply. Here what you have asked for and what I have provided is an excuse. Or to take another example, you might ask your math teacher to explain how to solve a particularly nasty problem. Here you are asking to be shown how to do something. But suppose you were to notice something curious. On the front of my shirt, just below the pocket, there is a bright red stain. "What happened?" you might ask, pointing at my shirt pocket. In effect, you are asking neither for an excuse nor to be taught how to do something. Instead, you are asking for the reason why something is the case, the reason for the mess on my shirt. "Oh that," I reply. "It's my red pen. It must have leaked again." In speaking of a "scientific explanation" we are speaking of an explanation in this latter sense: an account of how or why something is the case.

Explanations in science are often identified with causes, as in the example just above. To give a causal explanation of something is to set forth those events that led up to the thing in question. My pen leaked, I didn't notice it and the ink saturated my shirt. As we shall see, causal explanations play a central role in many scientific investigations. But scientists often rely on other ways of explaining, ways that utilize notions other than effects and their antecedent causes. In this chapter we will discuss these basic explanatory strategies and then consider how scientists respond when confronted with rival explanations for a single event or set of facts. But first, we need to do a bit of terminological ground clearing. Two notions are closely associated with explanations in science—*theory* and *hypothesis*. What they involve and how they differ will be our first topic.

At the most basic level, both hypotheses and theories are types of scientific explanation. A hypothesis can be anything from a vague hunch to a finely detailed, though speculative, account of how or why something has come to be the case. In general, however, the point of characterizing an explanation as a hypothesis is to note that it is tentative and unproven. You may have heard of the Tunguska blast. In June of 1908 a mysterious explosion occurred over the skies of Siberia. The impact of the blast decimated 830 square miles, destroying somewhere in the neighborhood of eight million trees. The precise details of what caused the Tunguska blast are still not fully understood. Several hypotheses have been advanced, some quite plausible, some highly questionable. Among the latter is the claim that the blast was caused by the destruction of an alien spacecraft hovering in Earth's atmosphere. Perhaps the most widely accepted is that the blast was caused by the explosion of a stony meteoroid as it reacted to the friction of Earth's atmosphere, finally coming apart somewhere between four and six miles above Siberia. Yet both claims must be regarded as hypotheses, since the exact details of what happened remain in question.

By contrast, "theory" does not always imply the kind of tentativeness associated with hypotheses. A theory may be a well-developed, well-confirmed body of explanatory material, as in the big bang theory, the theory of evolution, or the germ theory of disease. But often, people say things like, "That's only a theory," meaning roughly, "That's only your opinion of why so and so happened." To make matters worse, many of the things referred to in science as theories are subject to serious question. In astronomy, for example, one highly questionable alternative to the big bang theory is nonetheless referred to as a theory, the steady state theory. Moreover, broad explanations that were once embraced but finally discarded are still referred to as theories. The notion that the earth is at the center of the universe has long been rejected but it is nonetheless still called the Ptolemaic theory.

What typifies theories in science is the breadth and depth of their explanatory power. A hypothesis typically will offer an explanation for a limited range of phenomena, a single event, or a fact. Theories tend to be more general structures capable of explaining a much wider variety of phenomena. Moreover, theories will often contain well-confirmed rules and principles that reveal underlying explanatory similarities between apparently quite diverse phenomena. With four basic principles of motion and a handful of definitions, Isaac Newton was able to explain the behavior of just about anything with mass, from the tiniest of particles to the stars and planets. Similarly, the theory of evolution by natural selection offers a coherent picture of how the vast array of organisms on our planet has developed, and does so by reference to successive applications of a profoundly simple procedure.

As you can see, "theory" and "hypothesis" are used to cover a lot of ground and there is no simple and straightforward line of demarcation between the two. The net effect is that when someone speaks of a theory or a hypothesis, we may not be entirely clear what they mean. With a few exceptions in what follows, we can avoid any potential confusion by speaking simply of explanations. Explanations which share with hypotheses a kind of tentativeness, we can call *novel* or *proposed* explanations or something similar. Explanations which are well established, like some theories, we may simply characterize as *received*, *established*, *generally accepted*, etc.

CAUSATION

One way to explain how or why something has occurred is to give an account of the events leading up to it. Why, for example, when we were small children, did teeth, carefully tucked under our pillows, vanish only to be replaced by money? Because while we were sleeping our parents removed the teeth and replaced them with money. Why is there a circular crater several miles in the diameter in the Arizona desert? Because a large meteor survived its trip through Earth's atmosphere intact; its crash produced the crater. Why is smoking on the increase among young adults? In part, because the tobacco industry targets this segment of the population in much of its advertising. In each of these cases, a cause for a particular effect is identified, and with each we understand something of why the phenomenon in question is the case.

Causal explanations are common in our daily lives. Imagine I've arrived late for a lunch engagement. "Sorry I'm late. The traffic was horrendous," I say. My excuse is to the effect that something out of my control caused me to be late. Or suppose the street out front of the restaurant where we are meeting is flooded. You venture the guess that all of the drains are clogged with leaves. Your guess involves a causal explanation. The leaves covering the drains have caused the street to flood. Causal relationships are not always simple or straightforward. Consider some of the complexities we must face in thinking about causes and their effects.

First, effects can be the result of a combination of causes. It may be, for example, that my lateness was in part caused by a traffic jam. But suppose that while hung up in traffic I ran low on gas and so had to stop and fill up. Suppose also that neither event, alone, would have made me late. My being late has been caused by a combination of events.

Second, both causes and effects can be about groups rather than individual facts or events. To claim, for example, that cigarette smoking causes lung cancer is to claim that lung cancer will occur with greater frequency among the group of people who smoke than among those who do not.

Third, effects may result from several distinct causes. We know that cigarette smoking causes lung cancer. But other things—exposure to asbestos and genetic abnormalities, for example—can cause lung cancer as well. In some cases, a series of discrete causes will be responsible for an overall effect, though each factor will be responsible for only part of the effect. In the early 1990s, violent crime rates in the United States fell precipitously after 15 years of constant increase. As it turns out, several factors were responsible for the decline, though each was responsible for only part of the overall decline. Among the more pronounced causal factors were: increased use of capital punishment, increased rates of imprisonment, increased number of police, changes in the illegal drug market, and the legalization of abortion approximately 20 years earlier.[1] A rough estimate is that the hiring of additional police accounted for roughly 10% of the 1990s crime drop and that decreased demand for crack cocaine was responsible for another 15%.

Fourth, effects need not invariably be associated with a given causal factor. Though cigarette smoking is indeed a cause of lung cancer, many cigarette smokers will neither contract nor die of lung cancer. There is today a good deal of evidence that children who drink fluoridated water will have fewer problems with tooth decay than children who do not drink fluoridated water. Though there is clearly a causal link between fluoride and tooth decay, it does not follow that children who drink fluoridated water will be completely free of decay problems.

Fifth, causal explanations can be negative, as in our last example. Fluoridation of the water supply prevents tooth decay. Similarly, many people believe there is a causal link between vitamin C and the common cold. Regular doses of vitamin C, it is claimed, will decrease your chances of contracting a cold.

Finally, causal explanations can involve a sequence of linked events. If, say, A causes B which in turn causes C, A is often referred to as a *proximate* cause of B and a *remote* cause of C. B in turn is a proximate cause of C. So for example, if I trip and bump into the table where you are seated, causing your water glass to spill into your lap, my tripping is the proximate cause of the movement of the table and the remote cause of the mess in your lap.

CORRELATION

Closely related to the notion of a causal explanation is that of a correlation. Indeed, people often assume that if two things are correlated they are causally linked. But this assumption is often wrong. A correlation is nothing more than a comparison between a pair of characteristics within a population. Those characteristics are correlated if they display some regular, measurable variance. The simplest sort of correlation involves the comparison of two groups, one having a given characteristic and the other lacking it. If a second characteristic occurs at different frequencies in the two groups, it is correlated with one of the two. Suppose, for example, that we compared two groups of people, all between ages 30 and 49. The first group all have completed at least four years of college, while the second have less than four years. Suppose also that we were able to look at the average annual income of the two groups and were to find that the income of the first group is, on average, 20% higher than that of the second group. This means there is a correlation between education and income in the groups of people we have considered.

Correlations can be positive or negative. If a characteristic occurs at a greater frequency in one group than in the other, it is positively correlated with the first group; if it occurs at a lesser frequency, the correlation is negative. By contrast, if the characteristic occurs at roughly the same frequency in both groups, there is no correlation between the characteristic and either group. In our example, we have uncovered a positive correlation between education and income. Suppose instead we had found that the income of those having four or more years of college was actually lower than that of people with less education. This finding would suggest a negative correlation between the two factors. Had we found no real difference in levels of income, we would have had to conclude that, insofar as we can tell, there is no correlation between level of education and income. (This does not mean there is no such correlation. All we can conclude is that our quick check of the data available shows no correlation!)

Correlations can also hold between pairs of characteristics within a single group. Within a group, if two measurable characteristics vary in a somewhat regular and predictable fashion, they are correlated. Suppose, for example, we had at our disposal a large amount of information about the freshman class at a small local college. Examining the data we find what appears to be an interesting relationship between first-semester grade point averages (GPA) and Scholastic Aptitude Test (SAT) scores. About 100 students completed the first semester. In most cases, say 75 or so, we find that GPA varies directly with SAT score. That is, if we arrange these 75 students in order of ascending SAT score, we find a corresponding increase in GPAs; the higher the SAT score, the higher the GPA. For the other 25 or so students, we find no regular variance. Some students with relatively high SAT scores have relatively low GPAs and vice versa. Some with average SAT scores have relatively high, some relatively low GPAs. Despite these exceptions, our findings suggest a positive correlation between SAT score and GPA, at least in the group we have examined. Had we found just the reverse—had we found that for most students, GPA diminished when SAT scores increased—we would have uncovered a negative correlation between SAT score and GPA. Suppose instead we were to discover no regular variance between SAT scores and GPAs; many students with relatively high SAT scores had average or low GPAs, while many with relatively low SAT scores had average or high GPAs. This would suggest that no correlation exists between SAT score and GPA in the freshman class of the college.

As our last examples suggest, correlation is seldom an all-or-nothing matter. A perfect correlation between two characteristics would require a one-to-one correspondence between changes in the two. (In our example, increases in SAT score would need to be accompanied by increases in GPA in all 100 cases to establish a perfect correlation.) But particularly when groups of subjects are large, the fact that a correlation is somewhat less than perfect does not undercut its potential significance, perhaps as a predictor of one characteristic in cases where we know something about the breakdown of the other. Presuming, in our example, that we have uncovered a fairly consistent positive correlation between SAT score and first semester GPA, we may be able to predict something about a new college student's chances of success, based on his or her SAT score. But here we need to introduce a crucial note of caution. Any inference we draw about an individual, based on the evidence of a correlation, assumes a causal connection between the correlated characteristics. And this assumption is not always warranted. The fact that two things are correlated does not, by itself, indicate that the two are causally linked.

Why this is so can be seen in the following examples. If we were to examine a group of similar people, say, members of a single trade or profession, we could probably unearth a number of correlations. We might find, for example, a correlation between age and income, established either by showing a regular variance between age and income for the whole group, or by showing that people above and below a certain age have, on average, different income levels. We would probably also find a correlation between age and the use of reading glasses. Given these correlations, it is likely we will also find a correlation between income and the use of reading glasses! Now, none of these correlations seems to be a coincidence. There seems to be a clear link between age and the need for reading glasses. But the link in the other two cases is much more tenuous. Advancing age does not cause one's income level to rise, nor does income have any bearing on the need for reading glasses. The link in these two cases is undoubtedly explained by some other factor or series of factors. For example, in most trades or professions, the longer one works at a job, the more one generally makes. This, then, accounts for the correlation between age and income.

To make matters worse, a correlation may be evidence of nothing more than coincidence, a "mere correlation." This is because unrelated things can vary in regular, measurable ways. For a number of years now, two things have regularly increased: the sale of Burger King Whoppers and the number of minutes per day that children watch television. Come to think of it, recent increases in Whopper sales are correlated, negatively, with a gradual but regular decrease, in the same period, in the number of people who go bowling! And since we are an aging population, I suspect we could also dredge up a correlation between Whopper sales and the purchase of reading glasses. These new correlations, of course, suggest nothing more than the fact that lots of things, many of them not causally related, vary over time in somewhat regular ways.

All of this is not to say that the search for correlations is not an important component of causal research. Indeed, if two things are causally linked, they will be correlated, and so evidence of a correlation may provide some initial evidence for a causal link. But the simple fact that two things are correlated is, by itself, not evidence of a causal link. In Chapter 5 we will look closely at the ways in which claims about causal links and their attendant correlations are tested. For now, it is enough to keep in mind that correlations do not necessarily indicate causal links and, for this reason, are of less explanatory value than are facts about causal links.

Though causal explanation plays a central role in science, there are several other ways of explaining, all of which have a distinct role to play in scientific investigation. Consider another pair of examples. Why do our eyelids blink open and shut several times every minute? To keep the surface of the eye moist. Why does a gun "kick" as it is discharged? Because of a well established physical law: for every action there is an equal and opposite reaction. Neither of these explanations involves a cause, at least in any straightforward sense. We normally think of causes as events that precede the things they bring about. Physical laws do not cause things to happen in this sense. Nor is the blinking of an eyelid caused by the fact that it thereby keeps the eye moist.

An explanation may focus on the function or role a thing plays in some larger enterprise or on a law that accounts for the behavior in question. These

QUICK REVIEW 3.1 Causation and Correlation

Causation
Two things are causally linked if one proceeds and is responsible for the other. Suppose your car won't start because its battery is dead. The dead battery is the cause and your car's failure to start, the effect. Effects can have more than a single cause, and there may be many causes for similar effects. Several causal factors are responsible for the behavior of the stock market and a market decline can be caused by a variety of factors. Causal relationships can hold between individual events or between large classes of events as in the claim that megadoses of vitamin C can reduce occurrences of the common cold. If events are causally linked they will be correlated, but correlations do not necessarily indicate causal links.

Positive Correlation
In two populations, P and Q are positively correlated if a greater percentage of Ps than non-Ps have Q. Suppose that nationwide, people with cell phones have, on average, a higher income than people without cell phones. Cell phone ownership and income are positively correlated.

In a single population, if a regular increase in one trait, P, is accompanied by a regular increase in another, Q, then the two are positively correlated. Suppose worker productivity at a plant increases as pay increases, though with some exceptions. Worker productivity is positively correlated with income.

Negative Correlation
In the two populations, P and Q are negatively correlated if a smaller percentage of Ps than non-Ps have Q. Suppose regular users of the local public library (once or more a month) watch, on average, much less TV than sporadic or nonusers of the library. TV-watching and library use are negatively correlated for the group in question. In a single population, P and Q are negatively correlated if a regular increase in P is accompanied by a regular decrease in Q. Suppose that the number of visits to the library per month increases as the average number of hours watching TV decreases. Library use and TV watching are negatively correlated.

No Correlation
In two populations, P and Q are not correlated if there is no difference in levels of Q in P and not-P. If equal percentages of males and females are left-handed, there is no correlation between left-handedness and gender. In a single population, two traits are not correlated if there is no regular variance between the occurrence of the two. Suppose we were to record both the number of checks written and the number of soft drinks consumed per month by a randomly chosen group of people. We would probably find no evidence that variation in one trait is a predictor of a variation in the other. This suggests there is no correlation between the two.

Perfect Correlation
An invariant relation between two traits; for every change in one trait there is a consistent change in the other. In most species of trees, age in years is perfectly correlated with the number of rings in the tree's trunk; the older the tree the greater the number of rings, without exception.

strategies were used in the examples just above. Explanations can also make reference to underlying processes and causal mechanisms, techniques often used to enrich causal explanations. Since each of these ways of explaining plays an important role in scientific inquiry, let's take a closer look at what each involves.

CAUSAL MECHANISMS

Earlier we noted that a causal explanation can involve a series of linked events. One cause can be more remotely connected to its effect than another, more proximate cause. A causal mechanism is nothing more than a series of proximate causes that intervene between a remote cause and its effect. Researchers often comment that though they have evidence for a causal link, they lack a clear understanding of the mechanism involved. This tells us that the sequence of causes which have led to the effect has not been fully fleshed in. Cigarette smoking, for example, is known to be linked to lung cancer. Yet despite the fact that we are quite confident there is a link between smoking and lung cancer, little is know about the mechanism—the physiological process—by which the carcinogens ("carcinogen" just means "cancer-causing agent") in cigarette smoke lead to uncontrolled cell growth in the lungs of the smoker.

A recent study revealed an apparent causal connection between aspirin consumption and the risk of heart attack. According to the study, men who take a single buffered aspirin every other day have a 50% lower chance of having a heart attack than do men who do not take aspirin. Here the connection between aspirin consumption and risk of heart attack seems to be fairly well documented. As it turns out, the causal mechanism by which aspirin reduces the risk of heart attack is also well understood. Aspirin interferes with the first stage of the blood's clotting process. Now, many heart attacks are caused by blood clots in damaged arteries. It seems that when the thin inner wall of an artery is damaged, aspirin inhibits the tendency of minute blood platelets to clot over the damaged area. Thus, aspirin reduces the clotting effect that can lead to serious heart attack.

To take a very different example, one more closely related to every day life, imagine the following.[2] A friend applied for a job she really wanted to get. Yet now she tells us she finds the job utterly uninteresting and probably wouldn't accept it even if it were offered to her. Why the change in attitude? We discover subsequently that our friend learned she had no chance of getting the job. But how, if at all, did this bring about her change in attitude about the job? The answer may well lie in a causal mechanism, often called cognitive dissonance reduction, that makes people cease desiring that which they cannot get; you may be familiar with this mechanism under its more common name, sour grapes. Having learned she wouldn't get the job, our friend adjusted her desires thereby reducing the dissonance caused by wanting something she could not have. No doubt, the notion of cognitive dissonance reduction is a bit less precise than is the mechanism invoked to explain the connection between aspirin and heart disease, and for that reason it would be more difficult to test. But such psychological mechanisms nonetheless play an important role in our attempts at explaining why people behave as they do.

UNDERLYING PROCESSES

In 1828 the Scottish botanist, Robert Brown, discovered that when tiny particles are suspended in a liquid they undergo a constant quivering motion. This phenomenon, called Brownian motion, remained a mystery until it was explained in a 1905 paper by Albert Einstein. Brownian motion is due to the constant buffeting of the suspended particles by the ever-moving molecules of the liquid. In this explanation the movement of the particles in the liquid is redescribed in terms of the properties of the liquid's component parts. Underlying processes, unlike causal mechanisms, do not attempt to "fill in the gap" between cause and effect by positing intervening causes. Rather the point is to redescribe the phenomena only now at a more basic level. Molecular bombardment is thus not the cause of Brownian motion. Molecular bombardment *is* Brownian motion described from the point of view of molecular chemistry, a point of view that sheds considerable explanatory insight into the nature of the phenomena.

Explanation by underlying processes is sometimes said to be reductionistic, in that descriptions of phenomena at one level are reduced to descriptions at another, more basic level. Reductive descriptions can be technical and usually will make use of explanatory notions that do not occur in the original description. You may, for example, be aware that fluorescent lamps are much more efficient than traditional incandescent bulbs. The explanation lies in the way each produces light. When light is produced by incandescent bulbs, the following process takes place. Electrical energy passes through a wire and heats it until it incandesces (glows). The wire, called a filament, typically is made of a metal called tungsten; the enclosing bulb around the filament directs or diffuses the light. The problem is that 90% of the energy put into such a bulb is released in the form of heat, while only 10% results in light. Fluorescent lamps produce light in a different way, by energizing gas. Electrical energy flows into electrodes at the ends of a tube. The electrodes emit electrons, which energize a small amount of mercury vapor held at very low temperatures inside the tube. The energized mercury molecules radiate ultraviolet light, which is in turn absorbed by a phosphorescent coating on the inside of the surface of the tube, thus producing visible light. This process produces very little heat; fluorescent lamps are able to convert almost 90% of the energy they consume into light. So, the amount of electrical energy required by a fluorescent lamp to produce a given amount of light is substantially less than that required by incandescent bulbs. In redescribing incandescence and fluorescence in terms of the behavior of their underlying constituents, we have introduced a host of new explanatory notions: electrons, electrodes, filaments, and gases and the way in which electrons behave under various conditions. In effect, we have explained the greater efficiency of fluorescence over incandescence by looking carefully at what is going on at a more fundamental level in each process.

LAWS

Laws are of two fundamentally different types. The first and perhaps most familiar are typified by the rules and regulations that govern our daily lives. "Buckle up. It's the law." "Income taxes must be paid by April 15." "Please turn off cell phones during the movie." Laws of this type are conventions: they are created and often enforced to regulate human behavior. They can be amended and replaced, obeyed or disobeyed. They range from the formal prohibitions of civil and criminal law to the subtle suggestions of our codes of proper etiquette. By contrast, the laws that play an important role in scientific explanation are of a considerably different type. Unlike conventions, scientific laws are not created to regulate anything. They cannot be legislated into existence, and they cannot be followed or disobeyed, though, as we shall see, they can admit of exceptions. Rather, scientific laws are generalized descriptions of regularities that have been found to occur in some area of nature.

What happens if heat is applied to a closed container of a gas? Pressure increases. Why? An important law governing the behavior of gases, discovered by Joseph Gay-Lussac, provides the answer. Gay-Lussac's law states that if volume is held constant, the pressure exerted by a gas will vary directly with the temperature. So as we increase the temperature of a gas by applying heat, we increase the pressure in the closed container. Such laws derive their explanatory power from their ability to reveal how particular events are instances of generally understood regularities in nature.

We tend to think of scientific laws as being universal, claiming that a particular kind of behavior will occur in all (or no) cases. Thus, Gay-Lussac's Law is universal in that it makes a claim about the behavior of all gases. Similarly, the law of gravity holds for all objects with a mass. Physicists now tell us that nothing can travel faster than the speed of light. But scientific laws need not be universal; some laws claim only that a particular kind of behavior will occur in a certain proportion of cases.

Suppose we were to learn that a good friend, a nurse, has contracted hepatitis B. We are aware that he works in a clinical setting where patients with hepatitis B are regularly treated. We are also aware that recent studies have shown that an alarmingly high number of health care workers contract the hepatitis virus from their clients—one out of four health care workers who are accidentally exposed to the virus will actually contract hepatitis B.[3] It seems a real possibility that our friend's condition is explained, in part, by the statistic we have just cited. The explanation we might give would go something like this:

> Exposed health care workers have a 25% chance of contracting hepatitis B. Friend F is a nurse who works in a setting where the risk of exposure to hepatitis B is high. F has hepatitis B. Thus, it is likely that F has contracted hepatitis B from a client.

Though this explanation involves a law, it is not universal; it does not claim that everyone who is exposed to hepatitis B will contract it. Here we have an example of explanation by law where the law on which we rely claims only that a certain proportion of those exposed will contract hepatitis B.

No doubt it seems odd to call this claim a "law," yet it is certainly law-like, in just the way Gay-Lussac's law is law-like. Both describe regular correspondences. In the case of Gay-Lussac's law, the correspondence is between the pressure, volume, and temperature of a gas; in the case of our latter law, the correspondence is between workers who are exposed to the virus and workers who subsequently contract hepatitis B. The difference is that laws of the latter sort, often called *statistical laws*, enable us to give explanation that must be carefully qualified. It may be that our friend has contracted the hepatitis virus from someone or something other than a client and, as our statistical law tells us, chances are quite good that exposure to clients with the virus will not lead to infection. Thus, we had to qualify our explanation by adding the phrase, "it is likely," to acknowledge the possibility that our explanation may be wrong for this particular case.

The social and behavioral sciences often make use of statistical laws in their explanations. For example, psychologists tell us that people tend to feel that they are under an obligation when something is given to them. (This is sometimes referred to as the law of reciprocity.) Thus, if an inexpensive gift—say, a small calendar—is included with a letter asking for a donation, more people will respond than would if the gift had not been included. This explains why so many trinkets come to us via commercial solicitations. Economists tell us that increases in the rate of inflation are likely to be followed by increased employment and that the scarcity of a product generally increases its desirability. Though these laws are subject to a variety of exceptions they are nonetheless valuable, well established explanatory tools.

Statistical laws often stand behind simple causal explanations of the sort discussed earlier. Consider again one of our examples. I claimed that I was late for a luncheon date because of a traffic jam. I suspect you would accept this excuse in part because you are aware that generally, when people are stuck in traffic, their travel time increases. If you did not believe this statistical law to be true, you would probably not buy my excuse.

FUNCTION

We often explain the things we and others do (and don't do) by reference to our hopes, wants, aspirations, beliefs, etc. "Why," I might ask, "are you only having a salad for lunch?" "Because," you might reply, "I want to lose a few pounds." To explain one thing by reference to the purpose it fulfills is to give a functional explanation. And so, explaining our behavior by reference to what we hope to achieve, as in the example above, gives one sort of functional explanation. Human behavior is not the only thing susceptible to explanation by reference to function or purpose. If you asked me about the rock sitting on my desk, I would offer the following explanation. A heating duct is located just over my desk and whenever the heat comes on, unsecured papers blow about. So I use the rock as a paper weight. Following a similar strategy, we might explain that a carburetor

QUICK REVIEW 3.2 Ways of Explaining

Causes
To explain one thing or event by reference to another that precedes it. *Examples*: "Debris from last night's windstorm caused the power outage." "Excessive alcohol consumption can damage the liver."

Causal Mechanisms
To explain by citing intervening causal factors, factors that explain the effects of a more distant cause. *Example*: "Debris from the storm severed power lines, thus causing last night's power outage."

Laws
To explain an event by referring to a general law or principle of which the event is an instance.

Example: "The fuel efficiency of a vehicle is determined in part by size and weight. This is because acceleration is directly proportional to force but inversely proportional to mass. Thus, the larger the object you want to move, the greater the force you need to apply, and so the more energy you need to expend."

Underlying Processes
To explain something by reference to the workings of its component parts. *Example*: "The chest pain and breathing difficulty symptomatic of pneumonia result from an infection of the lung tissue. The tiny air sacs of which the lungs are composed—called alveoli—fill with inflammatory fluid caused by the infection. As a result, the flow of oxygen through the alveolar walls is greatly impaired."

Function
To explain something by reference to the role it fulfills in some larger enterprise. *Examples*: "Many species of birds build their nests in high places—trees, cliffs, etc.—to protect their young from predators." "The lungs serve both as a means of introducing oxygen into and removing carbon dioxide from the blood stream."

is the component of an internal combustion engine that mixes fuel and oxygen. In both these examples, we explain by specifying the purpose the thing in question serves. The purpose of the rock on my desk is to hold down papers; the purpose of a carburetor is to mix fuel and oxygen.

In the social sciences, functional explanations are indispensable. A historian or economist, for, example, might explain the emergence of a social practice—say, slavery, or liberalized abortion laws—by reference to the role such practices play in some larger social or economic enterprise. Slavery, it seems, was instrumental in the development of economies of scale in the United States in the 18th century. Liberalized abortion laws adopted in the 1970s reflected changing attitudes about the role of women in society and thus provided women greater latitude in making decisions about their future.

As the examples above suggest, functional explanations often make reference to the purpose or purposes of that which is being explained. Because of this, it may seem that functional explanation will be useful in dealing only with human contrivances and behavior. But functional explanations can provide insight into other sorts of cases as well, cases in which "purpose" implies nothing about human intervention, planning or forethought. For example, functional explanations are often used in the biological sciences. One of the most influential figures in the scientific revolution was the British physician, William Harvey (1578–1657). Perhaps Harvey's greatest accomplishment was his discovery that the purpose of the heart is to act as a sort of pump, facilitating the circulation of the blood. Similarly, evolutionary biologists often explain the dominance of a trait within a species on the basis of the advantage it confers on those who have that trait—in other words, on the purpose it serves.

But in such cases, "purpose" need not be understood on the model of human purposes. Rather, "purpose"—as it is used in biological explanations—means something more like "role in some larger enterprise." To give the purpose is to specify that role. So, for example, to wonder about the purpose served by the bright colors of many species of flowers is merely to consider how this trait is beneficial in the propagation of those species. Bees, it seems, are attracted to brightly colored flowers and thus bright coloration tends to enhance the chances of pollination. In a perfectly harmless sense then, the "purpose" of a bright coloration in some flower species is to attract potential pollinators. But to explain by reference to such a trait is not to suggest that anything like conscious planning and deliberation are involved.

THE INTERDEPENDENCE OF EXPLANATORY METHODS

In science, as we have noted, an explanation tells us something about how or why something is the case. Yet rarely will an explanation be so complete as to leave no further whys or hows about the thing in question. It seems that in science the need for explanation rarely comes to an end. This fact is reflected in the interdependence of the various types of explanation we have just considered. Put simply, more than one type of explanatory claim may be involved in a chain of explanations. Knowing, for example, that the function of a carburetor is to mix fuel and oxygen, we might then go on to consider how a carburetor accomplishes this goal. And here we will probably need something like a causal mechanism. We will, in other words, need to consider how a carburetor's parts operate in conjunction with one another to accomplish the proper mixture of fuel and oxygen. To go even deeper, we may want to consider underlying processes by thinking about the chemical reactions that contribute to combustion. A sense of the function something performs can often guide our understanding of how and why it works as it does. (This strategy is sometimes called "reverse engineering": first figure out what a thing is intended to do, then consider how it is designed and built to accomplish that end.)

To take a different kind of case, if we want to understand more about a particular causal connection, we will need to speculate about causal mechanisms that may be involved. A lake is polluted and some of its indigenous species of wildlife begin to diminish. There seems to be a connection. But what is the process by which greater pollution leads to fewer and fewer species?

Similarly, if we want to understand more about why a law-like regularity obtains, we may need to consider underlying processes. Recall our discussion earlier of Gay-Lussac's Law: if volume is held constant, the pressure exerted by a gas will vary directly with the temperature. Why, we might wonder, should this particular relationship between temperature, volume, and pressure hold for gases? The answer to this question requires that we examine the processes underlying the phenomena described by our gas law. In fact, gases are composed of molecules rushing hither and thither at enormous speeds. Pressure on the container holding the gas is a result of gas molecules colliding with the walls of the container. When heat is applied to the container, it is translated into increased activity on the part of the molecules of gas. The result is that the number of collisions with the container increases, thereby increasing the pressure exerted on the container by the gas. (This is a very rough sketch of a basic notion in what is called the kinetic theory of gases.)

Or if we want to understand more about a process underlying something we may need to look more deeply for causal mechanisms and law-like regularities that are considerably more fine-grained in character. To return for a moment to our story about the process involved in fluorescent lighting, why would mercury molecules, bombarded by electrons, radiate ultraviolet light? To answer this question we would need to consider processes that intervene and perhaps even underlie the interaction of electrons and the various component parts of the mercury molecules.

At this point, you may be wondering whether the process of explaining can ever come to an end and if so, what method of explanation is at the root of things. These are deep and profoundly difficult philosophical issues. Some philosophers believe that as a given science matures, claims about causal connections and mechanisms will be replaced gradually by broader and broader laws describing more and more causal phenomena. In this view, the most fundamental kind of scientific understanding is that provided when laws are discovered that reveal something about the interconnectedness of a wide variety of phenomena; the wider the variety, the greater the understanding. Other philosophers would maintain that at least in certain sorts of cases, perhaps all, to explain a thing is to identify its immediate cause or causes, and that when we can find no further intervening mechanism, the process of explanation must come to an end. On this view, law-like statements, no matter how broad and unifying, merely help us to classify and describe the rather more basic causal process at work in nature. For our purposes, however, we need not wrestle with these deep philosophical issues. Suffice it to say, the kind of explanatory claim one will give—whether it be about causes, causal mechanisms, laws, underlying process or something else—will depend on how much one knows and, of course, what it is one wants to explain.

RIVAL EXPLANATIONS AND OCCAM'S RAZOR

Often there will be more than one possible explanation for something that is not well understood. Many of the examples discussed in Chapter 2 involved rival explanations. Are crop circles messages from alien beings? Are they hoaxes? Is cold fusion the result of a chemical or a nuclear reaction? Do people actually leave their bodies during near-death experiences? Or are they suffering from something like a hallucination brought on by the stressful conditions they are under? The first step in sorting through rival explanations is to apply a simple principle, *Occam's Razor* or, as it is sometimes called, the principle of parsimony.

Occam's Razor[4] is named for its author, a Medieval philosopher and monk, William of Ockham (1285-1349). Occam's own version of the Razor is somewhat obscure: "What can be done with fewer is done in vain with more." A more revealing version of this principle for our purposes is the following: given competing explanations, any of which would, if true, explain a given puzzle, we should initially opt for the explanation which itself contains the least number of puzzling notions. The rational behind this admonition should be clear. If a puzzle can be explained without introducing any additional puzzling notions, there is no good reason to entertain any explanation that involves additional puzzles.

To see how this principle applies, consider the following case. Imagine that you are unable to find your keys. You have searched all morning to no avail and you know they should be around the house somewhere because you remember using them to open the door when you came home late last night. One possible explanation is that you've simply put them somewhere that you haven't yet looked. But other explanations are available as well. Perhaps someone who shares the house with you has inadvertently taken your keys instead of theirs. These two explanations rival one another in that both, if true, would serve to explain the phenomena in question. Presumably, at least one of the two is wrong, though in just the right circumstances I suppose they might both be correct.

What makes one explanation more plausible than its rivals is a bit more difficult to say. Let's begin by considering a couple of explanations for your missing keys that are a bit more bizarre than the two we have considered so far. Perhaps someone broke into your house while you were asleep and stole them. Or perhaps they just disappeared into thin air. Compare these two new explanations with the first explanation we proposed above, that you have simply misplaced you keys. Our first explanation is at least fairly plausible in that it makes no reference to other things which themselves stand in need of explanation. Surely, you've misplaced objects before, only to have them turn up even after you were convinced they were lost forever.

Next consider the first of our rather more bizarre explanations: somebody stole your keys. Keep in mind here that the point of an explanation is to make sense of how or why something has happened. If in giving an explanation we invoke events which are themselves quite puzzling, we have really only avoided the question of why the event in question happened. Why would someone break into your house and only take your keys? And why is there no evidence of forced entry? Though I suppose these occurrences could be explained—maybe we are dealing with a clever burglar who intends to return when you are not home—I think you can see that each additional explanation makes the original explanation seem less and less likely. Now a whole string of events would have to occur in order for our second explanation to retain some sense of plausibility. Our final explanation does no good at all. The keys have just "disappeared into thin air"? How does this work? Were they consumed by a tiny black hole? Did they spontaneously melt? In the name of resolving a simple puzzle our final explanation has embraced ideas that are radically anomalous and, judging by what we know of nature, false.

By comparison with our two bizarre explanations, our first explanation—that you have put your keys somewhere you haven't yet looked—fills the bill here. So, to say that one of a series of rival explanations is the most plausible is to say it is the one most in keeping with Occam's Razor. Keep in mind that Occam's Razor does not rule out explanations which themselves involve notions not fully understood. Rather it only suggests that given competing explanations, we should favor the one which involves the least number of problematic notions. Forced to choose between clever burglars and black holes to account for the missing keys, Occam's Razor would suggest the former.

EXPLANATION AND DESCRIPTION

In this and in Chapter 2 we have discussed two key elements of scientific method: observation and explanation. Unfortunately, many reports of extraordinary happenings of the sort discussed in Chapter 2 blur the distinction between these two key notions. Ideally, observations should be couched in purely descriptive language that tells us what occurred: no more, no less. But often, reports of extraordinary events include much more than pure description. Imagine, for example, if someone—let's call them X—were to report awakening in the middle of the night to discover what appeared to be their long-departed grandmother standing at the foot of the bed. X might subsequently claim:

(1) I saw the ghost of my dead grandmother.

But what, precisely, is factual in (1)? What, that is, can we be confident actually happened? That the person had an extraordinary experience is clear. Beyond this it is hard to know just what to say. Consider two rival accounts of what may have happened:

(2) X had a vivid life-like dream in which X's grandmother appeared.

(3) Somebody played an elaborate but vicious prank on X in the middle of the night.

(1) through (3) implicitly contain explanations of the event in question. That is to say, each presupposes the truth of a very different explanation: (1) that

what X actually saw was a ghost; (2) that what X "saw" was part of a dream; and (3) that what X saw was real, but a hoax, not a ghost.

Similarly, many anecdotal reports of the extraordinary contain much more than a simple, objective description of the experience. Such reports often blend fact with untested explanation and are what we might call *explanation laden.* For example, the statement, "The flying saucer hovered over the horizon and then accelerated away at a fantastic rate," tells us a couple of things about a person who might claim to have witnessed such an event. First, the person had an undeniably extraordinary experience. Second, the person believes the proper explanation for the experience is that he or she actually saw an intelligently controlled spacecraft.

In evaluating such a report, we must do our best to separate the descriptive wheat from the explanatory chaff. If we are able to ignore the explanation-laden portions of a report of the extraordinary, we may be able to arrive at a clear sense of what actually was experienced and, thus, what needs to be explained. Think once again of our flying saucer report. Suppose we could establish, for example, that the person making the report actually saw a bright light near the horizon, looked away to call to a friend, looked again and saw only a dim, twinkling light at some distance from the original light. Having gotten clear on this much, we would at least be in a position to think about more plausible rival explanations.

I once spoke with a person who claimed to have lived in a haunted house. He recalled that every few nights he would hear a knocking at the front door despite the fact that there was never anyone there when he opened it. We agreed that a more accurate description of the experience would contain only the salient facts: on several occasion he heard a series of sounds, very much like knocking at the door and the sounds seemed to come from the area of the house near the front door. He also added that he was never near the door when he heard the noise. Once we focused on this new, more objective description, several plausible explanations immediately came to mind; a tree or bush knocking against the house, or perhaps some activity outside or even inside that from a distance sounded like knocking. Now, we may never discover what really happened on those nights when the person in this episode heard a "knocking" at the door. At the very least, however, we know what parts of the story are fact and what parts speculation. And this is the real value of carefully distinguishing between the descriptive and explanatory elements of an extraordinary claim.

ULTIMATE EXPLANATIONS

We have covered a lot of ground in this chapter. We have found that there is no single model for what an explanation should involve. We have also found that the answer to just about any explanatory question gives rise to further, deeper explanatory questions. Now one final issue remains to be tackled. Can the process of explanation be brought to a conclusion? Is it possible that at some point science will provide us with an understanding of nature so deep and broad that nothing further remains to be explained? These are fair questions about which we can only speculate. It is possible that science—at the level of theoretical explanation—will one day come to an end. Perhaps a single final theory will explain everything. In this view, the behavior of a simple mix of particles, fields, and forces will explain everything in nature, from the processes that underlie the structure of the universe to the development of plant and animal life and even the complex interactions of human social institutions. (This view is called "greedy reductionism" by philosopher Daniel Dennett.) Or perhaps a series of discrete theories, taken together, will do the job. Physics will provide a final theory of the fundamental stuff of the world, psychology a theory of human and animal behavior, neurobiology a theory of consciousness, and so on. But it seems just as likely that science will never come to an end and that each scientific discipline will continue to provide closer and closer approximations of what nature is really like, while never completely finishing the job. We might add that history seems to be on the side of this latter possibility. On more than one occasion a scientist has proclaimed that the end of science is near, only to be proven wrong by some new theoretical breakthrough.

No matter what the final outcome, at least this much can be said: science is moving us in the right direction. In every area of science today, progress is being made. Often this involves nothing more than adding new detail to received explanations. In broad outline, for example, the theory of evolution by natural selection closely resembles the theory propounded a hundred and fifty years ago by Darwin and Wallace. But the details of the theory have been fleshed in by more recent developments in fields such as genetics and microbiology. However, on occasion, progress will come at the expense of a widely accepted explanation. New investigative techniques will uncover anomalies which the accepted view cannot explain, and gradually a superior explanation will emerge, one that can accommodate the new anomalies as well as the phenomena explained by the outmoded view.

There is perhaps no better example of the kind of progress science is capable of making than the shift from the Ptolemaic conception of the universe to the Copernican. In the Ptolemaic view, systematized about 140 A.D. by Ptolemy Claudius of Alexandria, the stationary earth stands at the center of the universe and all heavenly objects revolve around it. The Ptolemaic view had considerable explanatory power in that the movement of all celestial objects known at the time—the sun, the moon, the five innermost planets, and the stars—could be explained by a series of complicated calculations, though in ways very different than we would use to explain them today. For example, careful observation revealed that Mars generally moves eastward across the night sky but occasionally appears to move backward for a bit before resuming its eastward course. In the Ptolemaic view, all celestial objects trace out circular orbits around the earth. Ptolemy explained the backward, or retrograde, motion of Mars by introducing the notion of an epicycle—a small circular loop in the orbit of Mars such that, from an earthly perspective, Mars would actually appear to stop and then move backward during its epicycle. A tribute to its explanatory value is the fact that the Ptolemaic view dominated Western thought for more than a thousand years.

In the 16th century, however, Nicholas Copernicus, a Polish scientist and astronomer, proposed a new and radically different view of the cosmos. In Copernicus's view, many of the basic assumptions of Ptolemy were wrong. The sun, not the earth, is at the center of things; two of the planets, Mercury and Venus, occupy orbits nearer the sun than does the earth; and what is more, many celestial movements are to be explained by the fact that the earth rotates on its axis. One advantage of the Copernican view is that it suggests a very different explanation for retrograde motion than does the Ptolemaic view. If, as Copernicus suggested, the orbit of Mars is outside that of the earth, then the double motion of Mars with respect to the earth explains the apparent backward motion of Mars. For in Copernicus's view, we observe the motion of Mars from a location that is itself moving through space, with the net effect that Mars will, on occasion, appear to be moving backward.

There are a number of interesting facts about this particular episode in the history of science. The first, of course, is the enormous shift in thinking about the nature of celestial motion occasioned by the work of Copernicus. One might think the Copernican "revolution," as it is sometimes called, would have ushered in a new level of accuracy and simplicity in the calculation of planetary motions. But as it turned out, Copernicus's explanation was neither more accurate nor even much simpler than that of Ptolemy. Both views explained roughly the same collection of data about planetary motion. Moreover, like Ptolemy, Copernicus had to introduce a number of epicycles into his work to make his explanation fit the facts. The real value, then, of Copernicus's achievement resides in the simple but profoundly new way of thinking about celestial motion that it introduced.

But our story does not end here. Though in rough outline the Copernican view of the universe finally replaced that of Ptolemy, many of the details of the Copernican view were themselves eventually rejected. Copernicus, like Ptolemy, believed that the planets trace out circular orbits around the sun. (In fact, it was this conviction that necessitated the introduction of the occasional epicycle in his calculations.) It remained for Johannes Kepler, nearly a century later, to discover that the planets trace out elliptical orbits around the sun. Kepler thereby reduced the kinds of motion required to explain the observed positions of the planets and did away, finally, with the infamous epicycle. In defense of Copernicus, it must be noted that Kepler had available much more accurate measurements of the movement of the planets than anything available to either Copernicus or Ptolemy. Yet despite the enormous import of Kepler's contributions to our understanding of celestial motion, it remained for astronomers long after the time of Kepler to refine the Copernican world view even further by removing the sun from its exalted position at the center of the universe.

CONCEPT QUIZ

The following questions will test your understanding of the basic ideas introduced in this chapter. Your answers can serve as a brief summary of the chapter.