To understand episodes in science, we have to learn more about how scientists decide whether a model really does fit the world.

CONVINCING OTHERS

It is not enough for an individual scientist to decide that a model fits. Other scientists must be persuaded to make the same decision. This requires data and arguments that will appeal to scientists approaching the subject with a wide variety of different interests, backgrounds, and skills. Franklin, for example, came to agree that DNA has a helical structure long after Watson, Crick, and Wilkins were quite convinced it must. Pauling fairly quickly agreed that the double-helix model was correct after seeing both the model and the x-ray data in London. Other scientists in the field quickly agreed after hearing about it or reading the papers by Watson, Crick, Wilkins, and Franklin.

SPREADING THE WORD

Once most of the scientists directly involved in a scientific area decide that a model fits, a much slower process begins by which that conclusion spreads to the general nonscientific public. Scientists may be involved in this process, as Watson has been, but so are many others, particularly teachers, journalists, and even filmmakers. It is at this stage of the game that the rest of us learn most of what we know about science. To understand and evaluate scientific information as presented for the general public requires learning how to use what is presented to reconstruct features of the models and the decision-making processes that went into producing the information in the first place. We turn now to this task.

2.3 MODELS AND THEORIES

Scientists often describe what they do as constructing **models.** Understanding scientific reasoning requires understanding something about models and how they are used in science. In fact, there are at least three different types, or uses, of models to keep in mind.

SCALE MODELS

Watson and Crick were helped greatly by actually trying to construct a physical model of DNA. This was a model in the ordinary sense in which model airplanes and dollhouses are models. They are all **scale models**. The big difference between Watson and Crick's model and more familiar scale models is the extreme nature of the scale, which, in the case of the DNA model, was roughly a billion to one. That is, an inch in the model represented roughly one one-billionth of an inch in an actual DNA molecule.

Scale models are widely used in science and even more widely used in engineering. We can learn a lot about the wind resistance of various automotive designs, for example, by testing scale models of automobiles in small wind tunnels. This is much easier, and cheaper, than building wind tunnels large enough to hold full-sized cars. Nevertheless, when you find scientists talking about models, they most likely are not talking about scale models.

ANALOG MODELS

In *The Double Helix,* Watson talks about noticing spiral staircases and thinking that the structure of DNA might be like a spiral staircase. He also had the example of Pauling's α-helix. Here we would say that Watson was using a spiral staircase and the α-helix as **analog models** for the DNA molecule. He was suggesting that the DNA molecule is analogous to the α-helix or to a spiral staircase. One might also say that he was *modeling* the structure of DNA on that of the α-helix or a spiral staircase.

The most famous analog model in modern science is that of the solar system as an analog model for an atom. The nucleus of an atom, containing protons and neutrons, is said to be analogous to the sun. The electrons are said to be analogous to planets circling the sun. There is no doubt that this analogy between the solar system and atoms was extraordinarily fruitful during the first half of the twentieth century. It suggested all sorts of questions that formed the basis of much research (e.g., How fast are the electrons moving around in their orbits? Are the orbits circular or elliptical?). In investigating such questions, scientists learned much about atoms. In particular, they learned about many respects in which atoms are not like the solar system. In the end, a good analogy often leads to its own demise.

Analog models are typically most useful in the early stages of research when scientists are first trying to get a handle on the subject. At this point, almost any suggestion as to how they might construct a new model may be helpful. At later stages, when the question turns to *evaluating* how well the new model fits the real world, the original analog model is less useful. In trying to convince Wilkins and Franklin that they had the right structure for DNA, Watkins and Crick did not appeal to features of spiral staircases. Nor did they simply appeal to Pauling's success with the α-helix. Similarly, facts about the orbits of the planets were not used as evidence that the solar system model of the atom is correct. For these evaluations, other evidence was needed.

In sum, thinking about analog models may be very useful in attempting to *understand* a proposed new model. Analog models are much less useful when attempting to *evaluate* a proposed new model.

MODELS AND MAPS

The models most commonly referred to in scientific contexts are *theoretical models*. In attempting to understand what theoretical models are, it is helpful to invoke an analogy between theoretical models and maps (i.e., to use maps as analog models for theoretical models). Maps are more abstract than scale models but still less abstract than typical theoretical models.

Before proceeding, you are encouraged to produce a map of your own. This could be a map showing a trip between home and school, between a dormitory and a classroom building, or between home and work. Figure 2.5, for example, is a map that depicts part of my own university campus, including the main library and the building housing my department.

On my map, as shown in Figure 2.5, there is a solid arrow. What is the object to which the arrow is pointing? Stop and answer this question before reading any further.

FIGURE 2.5

A partial sketch of the Twin Cities campus of the University of Minnesota.

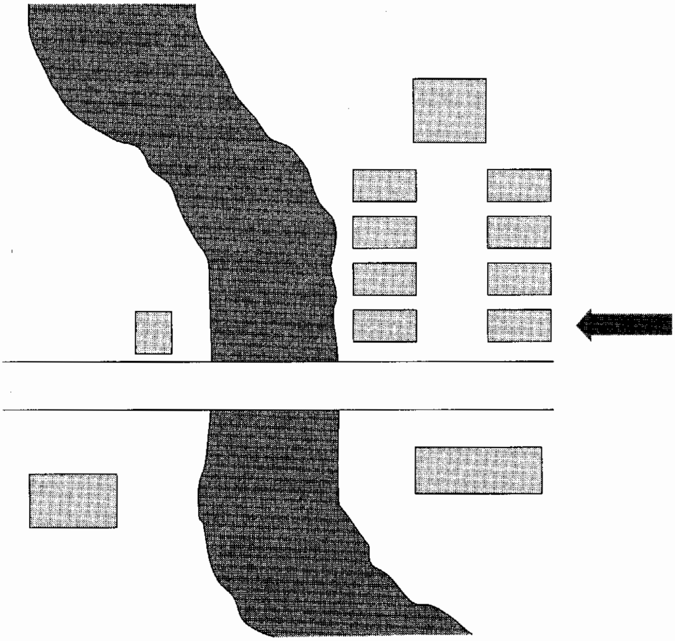

The usual answer is that the arrow is pointing to a building, presumably the university building in which my office is located. What would you think if I were to tell you that your answer is mistaken? Not only is it mistaken, it is not even close to being right. It is totally off base. No doubt, you would begin to suspect that there is some trick being played on you. You would be right. But the trick has an important point.

The correct answer is that the arrow is pointing to a rectangle drawn on the page. That is, quite literally, what the arrow is pointing to. The reason one is inclined to say the arrow points to a building is that it is pretty clear from the map and accompanying text that the rectangle *represents* a building. One therefore *interprets* the arrow as pointing to what the rectangle represents rather than to the rectangle itself.

The point of this exercise is that *a map is not the same thing as what it represents*. In the case of maps, no one is likely to make this mistake. After all, you can fold up a street map and put it in your pocket; you cannot fold up the city and put it in your pocket. Nor are you likely to mistake a scale model for the thing modeled. Surely no one was in danger of confusing Watson and Crick's wire and tin scale model for a real molecule of DNA. Theoretical models, as we shall see, are another matter.

Granting that a map is distinct from what it maps, what is the relationship between the two? It is true, but not very informative, to say that the map "maps" the area mapped. It is somewhat more informative to say that the map *represents* the area mapped. The next question is, How does a map manage to represent a particular space?

The first part of the answer is that a map exhibits a particular *similarity of structure* with the space mapped. In the case of maps, the particular similarity of structure is spatial. The spatial relationships among marks on a street map, for example, correspond to the spatial relationships among the streets in the city represented by the map.

The second part of the answer as to how a map manages to represent the area mapped is that we have a whole set of fairly well understood *social conventions* for constructing and reading maps. Without these conventions, a map would be just a piece of paper with lines drawn on it. The conventions for street maps are so well known that most people are not even aware of them. But this is a special case. Few people knew the conventions for interpreting Watson and Crick's scale model of DNA. One had to know a lot of physical chemistry to be able to interpret a particular tin plate as representing a purine base. Franklin could do it easily and quickly recognized that the three-chain model did not have enough places to attach water molecules.

The analogy between maps and models suggests further interesting questions. One is, Could there be a perfect map (e.g., a perfect map of Chicago)? The answer depends on what one means by a "perfect map." Suppose it means a map that contains a perfectly accurate representation of every feature of the city. Is that possible? Hardly. To represent every feature would mean representing every alley, house, garage, tree, bush, broken sidewalk, and abandoned car. It would mean representing not just the locations of buildings but their height as well. That is an impossible task. So, one way in which maps are not perfect is that they are *incomplete* in the sense that they represent only *selected features* of their subject, such as streets, and ignore others such as heights of buildings.

Restricting our attention to those features that are represented, there remains the question of how *accurately* those features are represented. For example, does our map of Chicago accurately portray the distance between the Water Tower and the Chicago River or between Michigan Avenue and Halsted Street? And is it accurate to the nearest 10 yards? yard? foot? inch? Clearly, no map is going to be perfectly accurate down to a fraction of an inch.

In summary, a map can be used to represent a place because there exists a set of social conventions that allow us to interpret features of the map as representing features of the place. All maps are incomplete in that they do not represent all features of the place represented. And no map gives a perfectly accurate account of the features that are represented. Nevertheless, there remains a similarity of structure between the map and the place represented. They are similar in some specifiable respects and to some specifiable degree of accuracy. All these things hold as well for the relationship between theoretical models and the parts of the world they represent.

THEORETICAL MODELS

A theoretical model is part of an imagined world. It does not exist anywhere except in scientists' minds or as the abstract subject of verbal descriptions that scientists may write down. When Watson was building the three-chain model, for example, he could have written down a description of what a DNA molecule would be like, if it were like this model. His description could have begun: "The model has three sugar phosphate backbones that twist in a helical structure with bases arranged. . . ." This description obviously could not describe a real DNA molecule, because we now know that DNA has only two chains, not three. What it describes, rather, is a *possible* molecule that turned out not to exist at all.

Watson and Crick did build a scale model with three chains. What is the relationship between that scale model and the corresponding theoretical model? The scale model can be used in place of words to characterize the theoretical model. One simply says, "The theoretical model has three sugar phosphate chains with bases arranged like this" and, then, points to the scale model. This strategy works because there is a similarity of structure between the scale model and the theoretical model. We can understand that similarity if we know the conventions used in building the scale model (e.g.,that red wires stand for hydrogen bonds).

Why can we not just stick to scale models and dispense with the notion of theoretical models? Because not all theoretical models have corresponding scale models. Watson, for example, never completed a scale model of the two-chain molecule with like bases bonded together. But this model existed as a theoretical model. Watson even described it in a letter to Max Delbrück. More fundamentally, scientists construct theoretical models of a whole variety of complex processes for which it would be difficult, if not impossible, to build working scale models.

To keep the idea of a theoretical model from becoming too mysterious, it helps to realize that we all frequently create things like theoretical models. For example, we can imagine giving a party, including imagining who comes with whom and who says what to whom. Here, we are constructing a theoretical model of a complex social event. The party may never occur, or if it does, it may be nothing like originally imagined. In the process of doing science, scientists imagine all sorts of complicated things and processes, including large molecules such as DNA. Some of these imagined possibilities turn out to have counterparts in the real world; others remain mere possibilities.

THEORETICAL HYPOTHESES

The most important question we can ask about a theoretical (or scale) model is whether it is, indeed, similar to the world in the intended respects and to the intended degree of accuracy. By way of shorthand, we often simply ask, "Does the model fit the world as intended?" Even more simply, "Does the model fit?"

There is another way of talking about the fit between models and the world that is often used by scientists, journalists, and other commentators on science. To accommodate this way of talking, we must introduce some additional terminology.

When scientists make the claim that their model is, in fact, similar to the world in the desired respects, we can say that they have formulated a *theoretical hypothesis* and that they are claiming that this theoretical hypothesis is *true*. A **theoretical hypothesis,** then, is a statement (claim, assertion, conjecture) about a relationship between a theoretical model and some aspect of the world. It asserts that the model is indeed similar to the world in indicated respects and to an implied degree of accuracy. If the model is similar to the world, as claimed, then the theoretical hypothesis is **true.** If the model is not similar to the world, as claimed, then the theoretical hypothesis is **false.**

FIGURE 2.6

A picture of the relationship between a model and the real world. The hypothesis that the model fits the real world may be either true or false depending on whether or not the model actually fits.

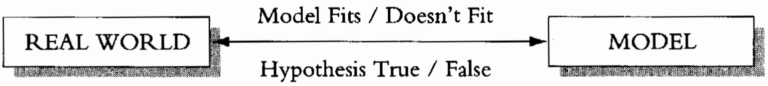

For example, early in the story, Watson and Crick formulated the theoretical hypothesis that DNA has a helical structure with three polynucleotide chains. That hypothesis was shown to be false. They later formulated the theoretical hypothesis that DNA has a two-chain structure. That hypothesis turned out to be true. In general, asking whether a specified theoretical hypothesis is true or false is just another way of asking whether the corresponding theoretical model fits the real world. This relationship is pictured in Figure 2.6.

"Truth" is a heavy-duty concept whose meaning has been debated by philosophers and others for 2,000 years. Theories of the nature of truth are normally discussed in courses in logic and philosophy. The above discussion of the truth or falsity of theoretical hypotheses falls into the category of what is usually called *the correspondence theory of truth*. But there is no need to enter these troubled waters here. For the practical purposes of understanding and evaluating reports of scientific findings, it is sufficient to think in terms of the fit between a model and the world. And here the analogy of the fit between a street map and the streets of the corresponding city provides a useful guide. If that is not enough, we can always fall back on the more restricted relationship between a scale model and the real thing. If that does not work, an abstract inquiry into the nature of truth is unlikely to be of much help.

Finally, in everyday speech, the word *hypothesis* often carries the connotation of a claim that is highly speculative—a conjecture without any real support. Thus we may reply to a claim we dispute by saying, "Well, that is one hypothesis." Here, the implication is that there are other equally plausible hypotheses that we might propose.

For the purpose of developing a systematic framework for understanding and evaluating scientific findings, it is best to ignore this way of talking. For us, the claims made on behalf of both the three-helix model and the two-helix model are both hypotheses. Indeed, in our preferred terminology, all general scientific claims are "hypotheses." The difference is that some hypotheses are well supported by the evidence and others are not. The important thing is to learn to distinguish between those hypotheses that are well supported and those that are not. That tells you which hypotheses it is reasonable to regard as true and which not.

THEORIES

Everyone knows that scientists produce *theories*. Yet up to now, we have not explicitly talked about theories as such. The reason is that *theory* is a quite vague and often ideologically loaded term. The main reason for calling something a theory may be to give it honorific status or, alternatively, to call it into question. Which function is served by using the word *theory* depends largely on the intentions of the speaker and the nature of the audience. Here we use a more neutral analysis. We already have in hand all the elements of such an analysis.

For our purposes, a **scientific theory** has two components: a **family of models,** which may include both scale models and theoretical models, and a **set of theoretical hypotheses** that pick out things in the real world that may fit one or another of the models in the family.

In 1953, there was basically just one model, the double helix, and one chemical substance, DNA, to which it was applied. But several years later, when people began talking about "the theory of molecular biology," there was a whole family of similar models and a range of other substances to which they were applied, including ribonucleic acid(RNA). We will encounter other theories that clearly include many distinct, but similar, models.

Like the word *hypothesis,* the word *theory* often carries the common meaning of something speculative. Those who question the theory of evolution, for example, commonly claim that evolution is "merely a theory" and not a "fact." Here again, we shall reject this common usage. For our purposes, the theories of molecular biology and the theory of evolution are all theories. Whether they are *also* facts depends on whether the corresponding models fit the world or, alternatively, whether the corresponding theoretical hypotheses are true. If so, they are facts; if not, they are not facts.

The important question is, How well is each theory supported by the evidence? Here, the relevant distinction is not between theory and fact but between theory and *data,* a distinction that is crucial for evaluating theoretical hypotheses and, thus, theories.

2.4 DATA FROM THE REAL WORLD

Everybody knows that to determine whether a proposed model fits the world some information is needed about the part of the world in question. But not all information is relevant. We use the term ***data*** (singular, *datum*) to refer to all the special information that may be directly relevant to deciding whether the model in question does fit. There are several general characteristics that such data must have.

The first feature that information must have to function as data is that it be obtained through a process of **physical interaction** with the part of the real world under investigation. The interaction may be *active,* as when one does experiments on the materials in question. Or, the interaction may be *passive,* as when radio astronomers measure radio frequency signals from distant galaxies. In either case, the data result from a physical interaction with the relevant part of the real world.

A second general feature of data, as opposed to mere information, is that relevant differences can be **reliably detected.** Detection may be a simple matter of looking, as when we observe a chemical solution change from blue to green. More often, detection requires elaborate instruments that produce outputs among which a scientist can discriminate just by looking. Often, these outputs are computer printouts of tables of numbers or of graphs. Figure 2.7 provides a schematic picture of the relationship between the real world and some data.

FIGURE 2.7

A picture of the relationship between the real world and data generated through a physical process of observation and experimentation.

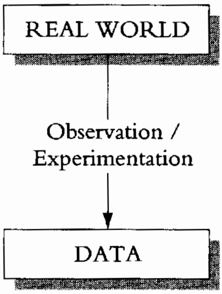

Among important bits of data in the story of the double helix were Chargaff's results on the one-to-one ratio of purines to pyrimidines in DNA, Franklin's results on the amount of water in DNA samples, and of course, Franklin's x-ray pictures of DNA. In each case, these data were obtained by actually working with samples of DNA.

Among information used by Watson and Crick that did not count as data in favor of their model of DNA was Pauling's discovery of the helical structure of α-kerotin. This information was influential in the decision by Watson and Crick to investigate helical models of DNA. Acquiring this information, however, required no physical interaction with samples of DNA. It could not, therefore, play a role as data for their hypothesis that DNA fits a helical model.

Here is one respect in which the analogy between maps and theoretical models breaks down. We can discover that a street map is deficient simply by finding a street that is not on the map. That requires no special skills or instruments. We can just look and see. Most models in modern science are not like that. Modern science typically investigates things that are very small (DNA), very large (our galaxy), very far away (distant stars), or otherwise inaccessible (the center of the earth). In all these cases, we cannot just look to see whether a proposed model fits. This fact has profound implications for how we might *evaluate* whether a model fits the real world.

2.5 PREDICTIONS FROM MODELS

The fact that we can only interact indirectly with the objects of a scientific investigation means that we can only investigate the fit between parts of the objects and some limited aspects of a proposed model. It is important, therefore, to be able to figure out what kind of data an object that did fit the proposed model would produce in the circumstances of a particular experiment. Scientists often speak of using a model to make **predictions** about what kind of data would be produced. This requires some explanation.

Sometimes, scientists use a model to make predictions in the literal sense of trying to say ahead of time what the data will be like. Often, however, predicting the data simply means being able to use the model to determine what the data should be like

FIGURE 2.8

A picture of the relationship between a model and a prediction obtained by reasoning about the model in the given experimental context.

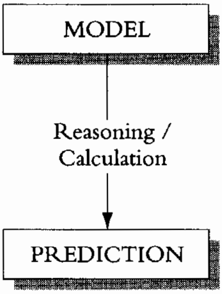

if the proposed model does fit the real world, even though the experiment has already been done. For example, the double-helix model was said to "predict" the Chargaff ratios even though Chargaff's experiments were done several years earlier. Similarly, the model allowed Crick to calculate the kind of x-ray pattern a double helix would produce and thus, in this respect, "predict" the kind of picture that Franklin, at that point unbeknownst to Watson and Crick, already had in her possession.

The example of Crick's prediction of the x-ray pattern exhibits another important feature of making predictions from models. It requires more than just the model in question. It also requires that one have a well-attested model of the experimental setup. Thus, Crick had also to have a good model of how x-rays are diffracted by atoms. Otherwise, he could not calculate what the pictures should look like if the x-rays were being diffracted by a helically shaped molecule.

Figure 2.8 provides a schematic picture of the relationship between a model and a prediction derived from the model. Note that here the arrow represents not a physical interaction with the world but a process of reasoning or calculation based on our understanding of the theoretical model under consideration.

2.6 The Components of a Scientific Episode

We are now in a position to construct our own model of a scientific episode in which data are used to evaluate whether a particular theoretical model fits the real world. Our model has four components: (1) a **real-world** object or process under investigation; (2) a **model** of the real-world object or process; (3) some **predictions,** derived from the model, describing what the data should be like if the model does, in fact, fit the real world; (4) some **data** generated through the sorts of interactions with the real world assumed in the predictions derived from the model. It is helpful to arrange these components as shown in Figure 2.9. This arrangement reveals four important relationships among the four components.

FIGURE 2.9

The four elements of an ideally complete report of a scientific episode involving a theoretical hypothesis.

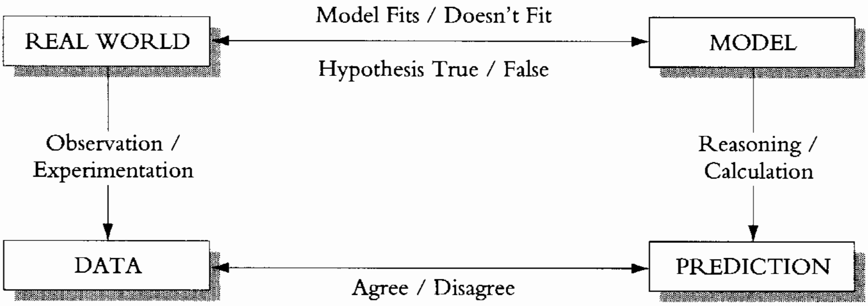

First, the relationship between the real world and the model is expressed by a *theoretical hypothesis* asserting that the model fits the real world. It is understood that the model fits only in some respects and then only to some specified degree of accuracy. If the model does not fit accurately in the intended respects, then the theoretical hypothesis is false.

The model and the prediction are related by *reasoning* or calculation. The real world and the data are related by a *physical interaction* that involves observation or experimentation.

Figure 2.9 also contains a relationship not previously noted, namely, one between the data and the prediction. If what is going on in the real world, including the experimental setup, is similar in structure to the model of the world, including a model of the experimental setup, then the data and the prediction should *agree*. That is, the actual data should be as described by the prediction. On the other hand, if the real world and the model are not similar in the relevant respects, then the data and the prediction may *disagree*.

To understand the relationships pictured in Figure 2.9, it is helpful to contrast the top part of the figure with the bottom and also the left side with the right. The top part of Figure 2.9 pictures the relationship between the real world and the model in question. Are the model and the real world similar in the respects under study and to an appropriate degree of accuracy? This relationship is typically not open to direct inspection. We cannot look at DNA in a test tube and see the helical structure. The bottom part of the figure, by contrast, pictures a relationship that can be evaluated by relatively direct inspection. Scientists can examine the data and see whether they agree with the predictions derived from the model.

The left side of Figure 2.9 pictures relationships existing in the physical world. The data are generated through physical interactions with bits of the real world. The right side of the figure, by contrast, pictures relationships that are mainly symbolic. The model exists mainly as a description of a possible type of object. This is so even if in giving the description we refer to a physical model such as Watson's scale model of DNA. Predictions derived from the model are likewise just descriptions of results that would be obtained in specified circumstances, circumstances that might not ever be realized.

It is important to keep these relationships in mind because *evaluating* hypotheses regarding the fit between models and the world turns out to depend crucially on whether the data and the prediction agree. That is, agreement at the bottom of the diagram is used to evaluate fit at the top. As we shall soon see, however, there is more to evaluating the fit between a model and the world than just agreement between data and predictions.

Finally, it should be realized that Figure 2.9 provides a model of fully developed scientific episodes that contain all four components arranged to make possible an evaluation of how well the model fits the real world. Many episodes, and thus many reports of scientific findings, do not include all four components. It is common, for example, to find reports that describe only the part of the real world under investigation together with some new data. There may be no mention of models or predictions. Similarly, we often find discussions of new models of real-world entities or processes with no mention of data or predictions. Occasionally, we find accounts of models of real-world things that include predictions but no discussion of data. We can learn a lot from such reports. Unless all four components are present, however, there may be nothing we can subject to an independent evaluation.

2.7 EVALUATING THEORETICAL HYPOTHESES

We are now ready to begin developing a general scheme that can be used by nonspecialists to evaluate scientific hypotheses as reported in various popular and semitechnical sources. We will continue with examples from the story of the double helix. Later in this chapter, we work through several completely different examples.

The basic idea behind the evaluation of hypotheses is to use the agreement or disagreement between data and predictions, information that is relatively accessible, to evaluate the fit between a model and the real world, something that is not directly accessible. Ideally, there are only two possible cases: either the prediction and the data agree or they do not agree. We will treat these cases separately, beginning with the case in which the prediction and data disagree. That turns out to be the simplest of the two cases.

In less than ideal situations, it may not be clear whether the prediction and data agree. Agreement may be a matter of degree. In such cases, it is difficult for a nonspecialist to make any independent evaluation of whether the model in question fits the world. Here, the best we can do is rely on the informed judgment of specialists. Unfortunately, when agreement between data and predictions is unclear, specialists often disagree among themselves about how well the model might fit the real world. In such cases, the only safe course for the nonspecialist is to regard the data as inconclusive and to suspend judgment about the model until more decisive data become available. If use of the model in question is relevant to some practical decision that needs to be made, the problem, then, is to make that decision in a manner that takes proper account of the uncertainty as to whether the corresponding theoretical hypothesis is true. This latter sort of situation is treated in Part Three of this text.

EVIDENCE THAT A MODEL DOES NOT FIT THE REAL WORLD

The story of the double helix provides a clear example of a model that was judged not to fit the real world—the three-chain model of DNA. Here, the decisive data were provided by Franklin's experimental measurements of the amount of water contained in samples of DNA. The three-chain model yielded a prediction as to how much water such a DNA molecule would accommodate. This prediction could be made simply by examining the scale model as long as one could interpret the model and knew enough physical chemistry to judge where water molecules might fit into the structure. The trouble was that the prediction from the model gave a value for the amount of water that was only one-tenth the amount Franklin had measured. So there was a clear *disagreement* between the experimental data from real samples of DNA and the prediction based on the three-chain model of DNA. This situation is pictured in Figure 2.10.

In this case, we are tempted to conclude without further ado that the hypothesis is false (i.e., real DNA molecules do not closely resemble the proposed three-chain model). Franklin immediately drew that conclusion, although Watson and Crick took a little longer to come around. As nonspecialists reading about this episode, we could just follow their lead. But if we are attempting to reach an *independent* evaluation, we cannot be quite so decisive. There are two possibilities that militate against so hasty a conclusion, neither of which tend to be accessible to a nonspecialist.

One possibility is that *the data were mistaken*. That is, Franklin's experiment yielded a mistaken value for the amount of water. There are all kinds of things that, unbeknownst to anyone, might have gone wrong with the experiments to yield a value for the amount of water ten times greater than the actual amount. Only people experienced

FIGURE 2.10

The elements of the episode involving Watson and Crick's three-chain model of DNA.

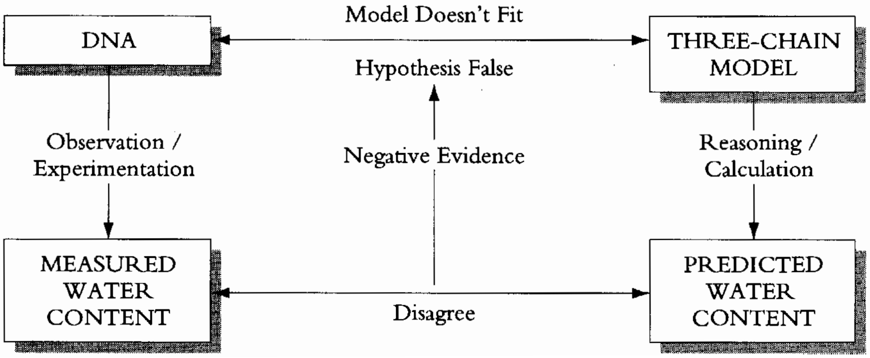

with the actual apparatus and experimental techniques can reliably judge how likely it is that something was seriously wrong with the experiment.

A second possibility is that, through misunderstanding of the model itself or because of a mistaken model of the experimental apparatus, *the prediction was mistaken.* A proper understanding of the model or the experiment might have yielded a predicted value in agreement with the actual data. Again, this is something for which a nonspecialist must rely on the judgments of the experts.

For these reasons, we will take reports of clear disagreement between data and predictions only as a basis for concluding that there is *good evidence* that the hypothesis is false. That is, there is a good reason, although not necessarily a conclusive reason, for believing that the model does not adequately represent the real world.

EVIDENCE THAT A MODEL DOES FIT THE REAL WORLD

One nice feature of the double-helical model for DNA is that having the sugar phosphate backbones on the outside left a lot of room for water molecules to attach themselves. So the double-helix model yielded a prediction for the amount of water in agreement with Franklin's data. Should we take that as evidence for thinking that the double-helix model adequately represents the physical structure of DNA?

As a matter of fact, Watson and Crick did not treat the agreement between the amount of water predicted by the double-helix model and the measured amount of water as a basis for arguing in favor of the double-helix hypothesis. Why not? Because they knew of many possible ways to build models with the required places for water. It could be done with a variety of three-helix models, for example, so long as one put the backbones on the outside. Thus, predicting the measured amount of water provided no basis for distinguishing the two-helix model from a variety of three-helix models. There was, therefore, no basis for regarding this agreement between prediction and data as evidence that the two-helix hypothesis, rather than some three-helix hypothesis, was true.

This explains why the x-ray data were regarded as being so important. According to Crick's calculations, a double helix should produce a quite distinctive pattern (i.e., a pattern unlikely to result from molecules with a significantly different structure). Thus, agreement between the predicted x-ray pattern and actual x-ray pictures provided a reliable basis for distinguishing between a double-helical structure and a variety of other structures. In this case, therefore, agreement between the prediction and the data did provide evidence in favor of the double-helix hypothesis. The components of this case are pictured in Figure 2.11.

The moral of this story is that mere agreement between a prediction and relevant data is not enough to provide a basis for thinking that a theoretical hypothesis is true. *Agreement counts only when such agreement would have been very unlikely if the hypothesis were clearly false,* which is to say, if some significantly different model provided a better fit to the real world. Ignoring this moral puts us in great danger of thinking that we have evidence in favor of what is, in fact, a false hypothesis.

At this point, we might have the following worry. No matter what the data happen to be and no matter what model is being considered, is it not always possible at least to imagine there being some completely different model that, nevertheless, just happens

FIGURE 2.11

The elements of the episode involving Watson and Crick's two-chain model of DNA.

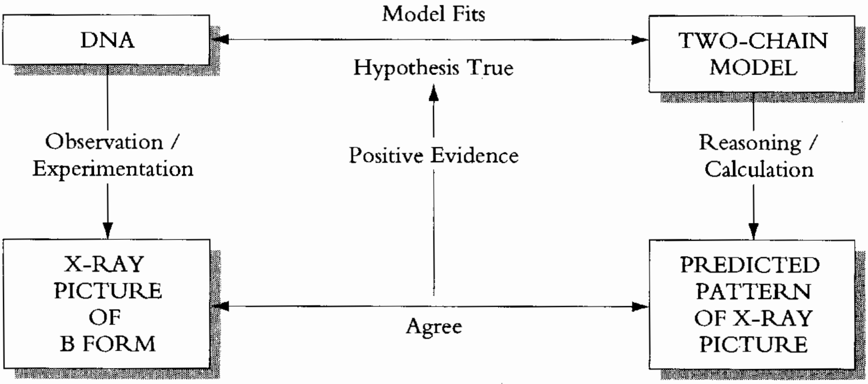

to yield the same prediction as the model under consideration? Does this not mean, therefore, that agreement between a prediction from a model and observed data can never provide any basis for thinking that the model fits the world?

This sort of difficulty has been voiced by scientific skeptics since the development of Greek astronomy nearly 2,000 years ago. The reply is that it is never enough for an alternative hypothesis to be imaginable in the abstract. It must be plausible relative the general scientific understanding of the kinds of models being used at the time by scientists working in the same general area. In the typical case, there will not be more than a few such plausible alternatives. We can, therefore, have good evidence that the model in question is the best fitting among all plausible alternatives. We cannot expect more from any scientific investigation.

The above theoretical worry has a more practical consequence. It is often difficult for a nonspecialist to judge independently whether there are very many other plausible models that would also yield predictions in agreement with existing data. In many cases, therefore, all the nonspecialist can do is rely on the reported judgments of specialists as to whether there are any such alternative models. These judgments, unfortunately, are often more implicit than explicit. We must read very carefully to determine whether there is any consensus on the availability of other plausible models yielding the same predictions. With experience, however, we can learn to recognize the kinds of hints from which we can infer the existence of the relevant consensus.

Finally, as in the case of disagreement between predictions and data, we cannot take even initially unlikely agreement between a prediction and data as a definitive basis for deciding whether a proposed model fits the real world. The most we can conclude is that there is *good evidence* for thinking that the model fits. The possibility of serious experimental errors or of mistakes in determining what prediction the model yields precludes a definitive conclusion on the basis of any single experiment. Even for cases in which it seems to the nonspecialist that scientists themselves reach a definitive conclusion on the basis of a single experiment, they may be relying in part on knowledge of other experiments, as well as on their considerable experience with the experimental and theoretical techniques involved.

2.8 A PROGRAM FOR EVALUATING THEORETICAL HYPOTHESES

In this section, we reduce the process of evaluating reports involving theoretical models to six easy steps. This does not require learning anything new. It is just a matter of organizing what we already know into a kind of "program" for doing an analysis. The advantage of developing such a program is that we can have in our heads a simple, uniform scheme for the evaluation of all sorts of scientific reports, a scheme that is easy to remember and to apply.

The program has two parts. The first four steps instruct us to identify the four basic components in a complete episode. These steps provide a basis for *understanding* the episode. If all four components are reported, we can go on to the final two steps, which constitute an *evaluation* of the associated theoretical hypothesis. If not all four components are identifiable, it may not be possible to perform an evaluation.

The Program

Step 1. Real World. Identify the aspect of the real world that is the focus of study in the episode at hand. These are things or processes in the world that can be described mostly in everyday terms together with a few widely used scientific terms. Do not use terms introduced to characterize particular models to be evaluated.

Step 2. Model. Identify a theoretical model whose fit with the real world is at issue. Describe the model, using appropriate scientific terminology as needed. A diagram may be helpful in presenting a model.

Step 3. Prediction. Identify a prediction, based on the model identified, that says what data should be obtained if the model actually provides a good fit to the real world.

Step 4. Data. Identify the data that have actually been obtained by observation or experimentation involving the real-world objects of study.

Step 5. Negative Evidence? Do the data agree with the prediction? If not, conclude that the data provide good evidence that the model does not fit the real world. If the data do agree with the prediction, go on to Step 6.

Step 6. Positive Evidence? Was the prediction likely to agree with the data even if the model under consideration does not provide a good fit to the real world? This requires considering whether there are other clearly different, but also plausible, models that would yield the same prediction about the data. If there are no such alternative models, the answer to the question is "No." In this case, conclude that the data do provide good evidence that the model does fit the real