scientific explanation? and How are scientific knowledge claims to be justified? As such, philosophy of science draws upon other areas of philosophy, especially logic, metaphysics, and epistemology. In subsequent chapters I examine a number of influential views of science advanced by philosophers and note some of their strengths as well as raising a variety of objections. As was noted at the outset, this inquiry has consequences because scientists often use perspectives from philosophy of science either to support their own work or criticize that of their opponents. As we proceed to look at a variety of answers in subsequent chapters, it is important to bear in mind that these issues are controversial and that there are a variety of different views on most important questions in philosophy of science. The answers philosophers have offered to these questions are fallible and philosophers should not be used as final authorities to resolve important questions in science. Rather, their views ought to be carefully examined and evaluated by all scientists who invoke them in making decisions about their scientific endeavors.

Before proceeding, however, a final qualification is needed. Twentieth century philosophy in Western countries has been broadly split into two traditions. What is commonly called *analytical philosophy* has been the dominant tradition in most of the English speaking world as well as much of Scandinavia and parts of Germany. A prominent focus of the analytic movement has been the logical analysis of language and statements, including scientific statements, made in language. *Existentialism* and *phenomenology*, on the other hand, have dominated philosophical work on much of the European continent. These traditions have been much more concerned than the analytic tradition with the subjective elements of human existence and have attempted to describe these systematically and explore their implications for a variety of human endeavors, including scientific inquiry. This is not the proper place for a detailed discussion of the relative merits of the two approaches. Both have shown considerable interest in science, but from adopted quite different perspectives toward science. The views discussed in this text are primarily drawn from the analytic tradition. However, the approach to philosophy of science pursued by this tradition has been so transformed in recent decades that it no longer bears clear affinities to its philosophical heritage of logical analysis but is more concerned with giving a faithful account of actual scientific practice.

2

Logical Positivism: The Received View in Philosophy of Science

INTRODUCTION: THE ORIGINS OF LOGICAL POSITIVISM

Logical Positivism emerged and became the dominant philosophical perspective on science in the first half of the century. Although its popularity has declined in recent decades, it continues both to set the agenda for many ongoing philosophical discussions and to provide the criteria that many scientists, including those in the cognitive sciences, use to judge what is good science. Logical Positivism arose in Austria (the Vienna Circle), Germany (the Berlin School), and Poland in the 1920s, but many of its principal theorists, including Rudolf Carnap, Herbert Feigl, Hans Reichenbach, and Carl Hempel moved to the United States with the rise of Nazism. Their views have subsequently been treated as part of the mainstream of "analytic philosophy" as practiced in the English-speaking world. Many of these founders were physicists and mathematicians who recognized that developments in physics, particularly the emergence of quantum mechanics and relativity theory, seemed incompatible with accepted wisdom about the nature of scientific investigations. As admirers of science generally, and especially of the new physics, the Logical Positivists set out to explicate the nature of science with a view to showing what made it a reliable source of knowledge.

The two terms that comprise the name "Logical Positivism" provide a good introduction to it. The term *positivism* comes from the philosophy of August Comte, an early 19th century philosopher who was skeptical of philosophical systems and of metaphysics generally and emphasized knowledge based on experience. He took science to be the paradigm of knowledge, citing as its strength the fact that it was empirically grounded in experience. More influential than Comte in providing this part of the foundation for Logical Positivism, however, were the classical 17th century Empiricists, especially David Hume, and their more contemporary descendants such as Ernst Mach. In accord with both Comte and the empiricists, the Logical Positivists maintained that all knowledge must be grounded on experience,[1] although the specific nature of this grounding was a matter of some dispute.

The term *logical* reflects the role that modern symbolic logic (see chapter 1) played in the views of the Logical Positivists. The Positivists used the resources of logic in attempting to provide a formal rendering of the structure of science. Because ordinary discourse often fails to adhere to the standards of symbolic logic, the Positivists found it necessary to propose the use of formal languages designed to adhere to the canons of symbolic logic in order to present their analyses. Their hope was that such a clear, formal presentation of science would ground the claim that science is a source of knowledge as well as help to resolve issues in science that had resulted from lack of precision.

An obvious difficulty that any practicing scientist will recognize in this program is that actual scientific thinking often fails to adhere to the strict canons of logical thinking. However, the Logical Positivists were not attempting to account for all scientific activity. First, they distinguished between *the context of discovery*, in which scientific hypotheses were developed, and *the context of justification*, in which they were rationally assessed (see Reichenbach, 1966). They held that the context of discovery might well be nonlogical. To cite a famous example, Kekule is supposed to have developed his proposal for the structure of the benzene radical while gazing upon the pattern of a flame from a burning log. He interpreted the flames as atoms dancing in snakelike arrays, and when one of the snakes seemed to grasp its tail, forming a ring structure, that suggested to him the ring structure for benzene. On the other hand, the process by which the correctness of this idea is tested is thought to depend on the logical relation between the hypothesis and the evidence that supports it. Hence, the Positivists proposed to leave to psychologists the task of explaining how scientists discovered new ideas and focused their attention on articulating the procedures of justification, whereby scientific theories could be shown to be true on the basis of evidence.

Even in the context of justification, however, the Positivists recognized that practicing scientists often do not adhere to the canons of formal logic. What the Logical Positivists maintained was that justificatory reasoning of scientists, if it was good, could be "reconstructed" to accord with the construal of scientific reasoning based on modern logic. The Positivists were, in effect, proposing normative standards for science. They were claiming that science that adhered to these standards, or that could be reconstructed so as to conform to them, constituted good science that provided knowledge about the world.

In order to clarify the basic conception of scientific justification advanced by the Logical Positivists, I explore three critical features of their account. The first is the theory they advanced about how terms in scientific laws had meanings (which captures the Positivists' allegiance to empiricism). The second is the deductive-nomological method of explanation and the related hypothetico-deductive model of justification, and the last is the axiomatic view of theories. (These last two together reflect the commitment to logical analysis.) Because the Positivists devised their conception of science with the physical sciences principally in mind, I rely heavily on examples from those disciplines in initially presenting the Positivists' program, and consider how Positivism might apply to examples from cognitive science in the concluding section of this chapter.

THE VERIFIABILITY THEORY OF MEANING

The Logical Positivists attributed many of the confusions and uncertainties of science, particularly those found in the social and behavioral sciences, to unclarity in the language. Even more strongly, they claimed that the quandaries that beset other areas of human inquiry, including politics, religion, and areas of philosophy like metaphysics, resulted from unclear use of language. When language is not governed by strict rules of meaning, the Positivists contended, confusion sets in and people may end up producing utterly meaningless statements. In calling a statement meaningless, the Positivists were not merely asserting that the statement was false but something worse—the statement was not really understandable. The kind of statement the Positivists had in mind is a statement like "God is love." Consequently, they viewed theological debates, for example, not as substantive debates for which there were objective answers, but simply as confused discourse. The remedy for such confusion was to attend carefully to the principles governing meaningful discourse and to restrict oneself to those domains where language could be used meaningfully. The Positivists did admit that language could serve other functions than making true or false statements. For example, they thought that literature and poetry could be used to arouse emotional responses or inspire action. But science, they maintained, was concerned with truth and therefore had to restrict itself to discourse for which clear principles of meaningfulness were available.

---

[1] To capture this aspect of their position, some of the logical positivists subsequently came to prefer the label "Logical Empiricism" for their endeavor.

In their discussions of meaning the Positivists followed the classical Empiricists in linking knowledge to experience, but they advocated one important change. The classical Empiricists treated ideas as the units of thinking and viewed these ideas as causal products of sensory experience. The Logical Positivists rejected ideas as fuzzy entities. Rather, they took linguistic entities—sentences and words—to be the basic vehicles of meaning. They proposed the criterion of verification to explain how these linguistic entities could be appropriately related to experience. According to this criterion, *the meaning of a sentence was the set of conditions that would show that the sentence was true.* Although these conditions would not actually occur if the sentence was false, we could still state what would be the case if it was true. Because only sentences and not individual words could be true or false, the meaning of words had to be analyzed in terms of their roles in sentences. This account of meaning became known as the *verifiabilty theory of meaning.*

Some sentences, the Logical Positivists maintained, could be directly verified through experience. Sensory exposure could tell us directly that these sentences were true or false. The Positivists referred to these sentences variously as *protocol sentences* or *observation sentences.* There was considerable disagreement amongst Positivists as to which sentences counted as such. Some, like the early Carnap (1928/1967), restricted observation sentences to those characterizing our phenomenal experience (e.g., "I am sensing a blue color patch now"). Others, like Nuerath (1932), maintained that sentences about observable parts of the world (e.g., "The sun is shining") could be directly verified. For the most part, Positivists took observation sentences to refer to physical states of the world, producing a bias toward that which is physically observable.

Other sentences in a language could not be verified directly through experience. This is particularly true of sentences that contain theoretical terms (e.g., *force*) that do not directly refer to observable features or objects. To explicate the meaning of these terms the Positivists focused on ways in which the truth or falsity of sentences using these terms could be determined indirectly via other sentences that were observational. Here logical analysis became important, for the Positivists had to explain the logical relationship between two sentences whereby one could serve to explicate the meaning of the other. Initially, a number of Positivists proposed to "translate" all sentences referring to theoretical entities into observation sentences (Carnap, 1923). Because they limited themselves to the tools of symbolic logic, the kind of translation with which the Positivists were concerned was not aimed at preserving the connotation of the theoretical sentence, but at identifying sentences that were true under the same empirical conditions. Thus, translations consist of biconditional sentences that assert that one statement (the theoretical statement) is true if and only if another, possibly complex statement (the observation statement) is true. These statements have an unusual characteristic. Because they only articulate the meaning of one sentence in terms of another sentence, they do not depend on experience in any way and so cannot be refuted by experience. Such statements are often referred to as *analytic statements* to distinguish them from ordinary sentences whose truth depends upon how the world is.[2]

This attempt to explicate the meaning of all scientific discourse in terms of observation conditions is closely related to the very influential doctrine, associated with the American physicist and mathematician Percy Bridgman (1927), of operational definitions. According to this doctrine, in introducing a theoretical concept, it is necessary to specify operations through which one can confirm or disconfirm statements using that term. Bridgman's notion of an operational definition extends the Positivists' conception of an observation term by supplying procedures for producing the requisite observations.

One of the issues in cognitive science to which the verifiability theory of meaning has been applied is the question of whether machines can think. In order to render this into a meaningful question, the Positivists required that it be translated into a sentence that can be confirmed or disconfirmed observationally. Turing's (1950) famous test for machine thinking provides the kind of operational definition of thinking that would be required. Turing proposed that we should accept a machine as thinking when we could not distinguish its behavior (e.g., in answering questions and carrying on a dialogue) from that of a thinking human being. Of course, we also confront problems in deciding whether another human being is thinking, or is simply an automaton. The verificationist theory of meaning, however, advocates the same treatment of this case—explicate what thinking is in terms of the kinds of behavior a thinking being would perform. This treatment construes the concept of thought as referring not to some unobservable activity but as something detectable in the behavior of organisms or computers. (For a philosophical analysis of thought in terms of behavior, see Ryle, 1949.)

The criterion that theoretical terms have to be translatable into observational terms was quickly recognized to be too strong. First of all, it is common for theoretical terms to be linked with experience in more than one way. This is particularly true for measurement terms for which there may be several different observational criteria. Generally, scientists will not accept just one of these as the definition, but view them as giving alternative criteria. Some of these may be discounted if several of the others all support a common measurement. This practice cannot be understood if one insists that there be a single definition translating theoretical terms into observational terms. Secondly, a number of theoretical terms, for example, dispositional terms like *soluble,* may not be translatable into observational terms. An object's property of being soluble cannot be correlated directly with observable features

---

[2] The latter kind of sentence is then referred to as a *synthetic statement.*

of the object except when the object is placed in water. Many soluble objects will never be placed in water. Even worse, the dispositional term cannot be translated into a conditional sentence (e.g., if it is placed in water, then it will dissolve). The reason is that in symbolic logic a sentence of the form "if———, then . . ." is defined as true if the antecedent is false (see previous chapter). This would make any object that was never placed in water soluble.

To account for the meaning of such terms, which contemporary science seems clearly to require, positivists attempted to weaken their verifiability conditions. Carnap (1936, 1937) proposed that a dispositional term like *soluble* could be translated by the following sentence (which he termed *reduction sentence*):

> "If x is placed in water, then x will dissolve if and only if x is soluble."

Such a reduction sentence overcomes the previous objection because it does not imply that something never placed in water is soluble. It also has the consequence that under conditions where the test conditions are never investigated (e.g., where the object is destroyed before it can be placed in water) we will not be able to determine the truth of the theoretical sentence. Unfortunately, this means that the initial aspirations of the verifiability criterion are not achieved because there will be reduction sentences for terms even though we may be powerless to verify the actual applicability of the term in specific instances. But at least, according to the Positivists, we know what conditions we claim hold when we make a statement using the term.

THE DEDUCTIVE-NOMOLOGICAL MODEL OF EXPLANATION AND THE HYPOTHETICO-DEDUCTIVE MODEL OF THEORY DEVELOPMENT

So far we have focused on the criterion for assessing the meaningfulness of scientific statements, but the goal of science is not only to make meaningful statements, or even true meaningful statements. For the Logical Positivist, the basic tasks of science were to explain phenomena in nature and to predict their occurrence. These tasks, as we see, were closely intertwined. Following a tradition that goes back at least to Aristotle, the Positivists maintained that explaining an event consisted of deriving a statement describing that event from statements of scientific laws and statements describing antecedently known empirical facts (initial conditions). Thus, deduction plays a central role in their account of explanation and the Positivists adopted what has been termed the "covering law" or "deductive-nomological" (D-N) model of scientific explanation. This basic model is represented by the following schema, in which L_1 through L_n represent general laws, C_1 through C_n represent initial conditions, and E represents the event to be explained:

$$\begin{array}{c} L_1, L_2, \ldots L_n \\ C_1, C_2, \ldots C_n \\ \hline \text{Therefore, } E \end{array}$$

The laws that are required for this schema are conditional statements of the form "If X happens, Y will happen." The initial conditions tell us that X has happened. An example of a law statement would be a statement of the form "If a human being is deprived of vitamin C for a certain number of days, the individual will suffer scurvy." The initial condition might then be that a particular individual had been deprived of vitamin C for that number of days. This would provide an explanation of why the individual experienced scurvy. (For an extended discussion, see Hempel, 1965, and Nagel, 1961.)

A couple of features of this general schema of explanation should be noted. First, in order to explain an event, according to the Logical Positivists, it was not sufficient simply to point to a factor that might have caused the event. For example, noting that someone threw a rock at it would not be sufficient to explain why the window broke. An explanation requires a complete derivation of the event from general laws and known facts. Hence, laws play a central role in any explanation. Second, according to this view there is a symmetry between explanation and prediction. They both have the same logical structure such that a derivation of the sort required for explanation would, if carried out before the event, serve to predict the event. The difference is simply a temporal matter that predictions are made before the event has occurred, while explanations are offered for events that have already taken place. Some critics have found this symmetry to be counterintuitive, because sometimes, at least, it seems that we may be in position to explain events that we could not have predicted. For example, after the fact, we may explain traffic accidents even though we never acquire sufficient information about the initial conditions or develop a precisely articulated law that would predict accurately when accidents would occur. We may be able to assign responsibility for an accident to a faulty traffic light, but still not have enough information about the movement of vehicles to predict when an accident would occur. (This type of objection is developed further by Scriven, 1962, whom I discuss further in the next chapter.) The Positivists, however, defend this symmetry, and would reject the claim that we have explained the event if we do not have sufficient information to determine just what accident would occur. If we do develop such information, then we can both explain why the particular accident occurred, and, if we had had the information in advance, we would have been able to predict it.

So far we have focused on deterministic explanations, where every time the initial conditions satisfy the law statements, the consequent will follow. Several positivists attempted to generalize this model to include probabilistic laws that hold that given a specified set of conditions, there is a certain probability that an effect of the specified kind will follow. In such a case we do not have a strict deduction of a statement specifying the event that occurred, but something weaker—a demonstration that a certain kind of event was likely. For example, taking a particular drug may cure a disease most of the time, but not in all cases. In such a case we can explain and predict the cure by appealing to the statistical relationship and the fact that the person had the disease and took the drug. Hempel (1962) thus proposed a modification of the D-N model to allow for "inductive-statistical" explanations where one could infer that the event was highly probably. This strategy only worked, however, for events whose probability was raised above .5 by the statistical regularity and the initial conditions. Events that follow upon one another with relatively low frequency, such as acquiring lung cancer after smoking, could neither be explained nor predicted on this account. Many critics found this consequence of the attempt to broaden the D-N approach to cover statistical explanation counterintuitive (see chapter 3).

Returning to the context of deterministic laws, we should note that a deduction of the kind called for by the D-N model of explanation would only count as an explanation if the laws adduced in the explanation were true. The law statements required in D-N explanations are generalizations that cover a potentially unlimited number of events. That is, they are statements of the form "If any object contains exposed iron, then it is subject to rusting" $[(x)(Fx \rightarrow Gx)]$. The fact that such statements apply to a potentially infinite number of states of affairs may, at first, seem to render them unverifiable and thus, given the verification theory of meaning, meaningless. But the Positivists took the very deductive relationships used in D-N explanations to give meaning to the law statements. Moreover, they viewed the events that were explained by the laws as themselves providing the evidence for the truth of the law. Thus, the previous law statement would be confirmed by particular events of iron rusting.

The Positivists called the procedure for developing scientific laws the ***hypothetico-deductive*** (H-D) method. The basic idea of the H-D method is that scientists begin with an event that requires explanation. Hempel (1966), cites the example of Semmelweis' work during the 1840s on childbed fever to illustrate the method. Semmelweis noted that a large proportion of the women who delivered children in his hospital contracted an often fatal illness known as "puerperal fever" or "childbed fever." The rate, moreover, was much higher in the wards where physicians handled the deliveries than in wards where midwives were in charge. To explain that it was necessary to propose a hypothesis from which the difference between the two wards could be derived. Prior to Semmelweis, a number of hypotheses had been proposed, but none seemed adequate to explain the differences between the two wards. Semmelweis, however, found a clue when a fellow physician came down with a fatal disease much like childbed fever after receiving a puncture wound while performing an autopsy. Semmelweis offered the hypothesis that "cadaveric matter" picked up during the autopsy might be the agent responsible both for his colleague's disease and the cases of childbed fever.

Having developed a hypothesis (recall that for the Positivists how hypotheses were arrived at was not a matter for logical inquiry), the task was to discover whether the hypothesis was true. If it was, it could provide the law needed to explain the event. Semmelweis would be reasoning circularly if the only evidence he offered for the hypothesis was the event he started out with. But the hypothesis is a general statement and so could be tested by considering other initial conditions, and deriving predictions about what would happen under those conditions. If these predictions turn out to be true, the initial hypothesis would be confirmed; if the predictions turn out false, the hypothesis would be disconfirmed. In Semmelweis' case, he proposed a test in which physicians would begin to wash their hands in chlorinated lime before examining patients. He predicted that the rate of childbed fever in the physicians' ward would decease (this was a consequence derived from the new initial conditions and the hypothesized law). This prediction proved true, providing evidence for the truth of the hypothesis.

Both the deductive-nomological model of explanation and the hypothetico-deductive model for developing explanations seem eminently plausible when one considers cases like that of childbed fever. Both, however, encounter some basic difficulties that were recognized by the Positivists themselves. The D-N model requires that one of the premises in the deductive explanation of an event be a law. Explicating what makes a statement into a law, however, is a difficult problem given the tools of symbolic logic upon which the Positivists relied. It was clear that a law statement had to be a true general statement of the form: "For all x, if x is F, then x is G" $[(x)(Fx \rightarrow Gx)]$. (For example, for any person, if the person is infected with cadaveric matter, then the person contracts childbed fever.) However, it was also clear that this is insufficient because we would not want to count all true general statements as laws. For example, if it were true that I only carried \$1 bills in my wallet, then the following would be a true general statement: for all x, if x is a bill in my wallet, then it is a \$1 bill. But this intuitively is not a law (Goodman, 1947). The reason is that there does not seem to be any reason except chance or my perversity for me to carry only \$1 bills in my wallet. It is commonly thought that laws are more than general statements that happen to be true. We think they tell us something about the limits of how things ***must be***.

Sometimes people try to characterize what is distinctive about laws by saying that they must be able to support counterfactual claims, that is, claims about what would be the case if the facts were different than they are. Counterfactual claims commonly take the form "If something were F (infected with cadaveric matter), then it *would be* G (ill with childbed fever)." This rules out the previous example concerning currency in my pocket because few of us think that if someone put a $10 bill in my wallet, it would become a $1 bill. The problem, however, is that counterfactual claims cannot be represented within basic symbolic logic. Some philosophers have proposed a variety of logics to handle counterfactual claims that are commonly referred to as *modal logics*. Such logics contain operators that specify what is possible or what is necessary. (See Stalnaker, 1968, for an attempt to apply modal logics to scientific laws.) Some Positivists (e.g., Carnap, 1956; Reichenbach, 1956) explored this avenue. However, for strict Positivists, the commitment to the verifiability account of meaning prevented use of modal logics. Nothing in experience could ground a distinction between common generalizations and modal statements because the only evidence we can acquire is that which supports the generalization. Counterfactual circumstances, by definition, do not arise and hence cannot be called upon to mark the difference between a true generalization and a true counterfactual. The only route open to the Positivists, therefore, was to attempt to differentiate law statements from merely true generalizations in terms of how the statements are accounted for by theories. Generalizations that are supported by theories have greater empirical support and hence are more likely to be true in new circumstances (Hempel, 1966). The role of such theories in the Positivists' view of science is the focus of the next section. What is important to note here is that being explained by a theory is the only factor to which the Positivist can appeal to distinguish laws from universal generalizations.

The use of symbolic logic also poses problems for the hypothetico-deductive analysis of hypothesis development. It was recognized by David Hume (1740/1888) that inductive evidence could never establish definitively the truth of any general claim. It is always possible that there might a counterevidence to a general claim that simply had not been discovered as yet. Yet the Positivists wanted to maintain that collecting evidence confirming a hypothesis should increase our confidence in its truth. The reason is clear: Confirmed predictions were the only vehicle recognized by the H-D model for gathering evidence for the truth of hypotheses or laws. But given the commitment to standard symbolic logic, even this is in jeopardy. A number of paradoxes were brought forward to challenge the assumption that confirming evidence should strengthen our belief in particular hypotheses.

One of these paradoxes (commonly known as the *Raven Paradox*) depends on the fact that a law statement of the form

For all x, if x is F, then x is G

is logically equivalent to the statement

For all x, if x is not G, then x is not F.

If F stands for *raven* and G for *black*, then the law "All ravens are black" (e.g., for all x, if x is a raven, then it is black) is logically equivalent to "All things that are not black are not ravens" (e.g., For all x, if x is not black, then it is not a raven). To test the first statement, the H-D model would lead us to examine ravens to see if they are black. The more black ravens we encounter, the greater support for the law (as long as we do not encounter ravens that are not black). But the form to which it is logically equivalent only requires us to examine things that are not black and test the prediction that these things will not be ravens. Every nonblack object that you see that is not a raven will confirm the putative law. So you can sit in the room you are now in and test the law that all ravens are black by making sure all the nonblack objects in the room are not ravens. Something clearly seems to have gone wrong!

Confronted by this and other logical peculiarities,[3] the Positivists sought to refine their account of how evidence could confirm hypotheses (see Swinburne, 1971, for a review). The Positivists' commitment to symbolic logic and, in particular, their commitment to laws being fundamentally universal generalizations, however, lay at the heart of these problems. Hence, they were not easily resolved. Moreover, the moves made to rescue Positivism tend to cloud the clear and intuitive picture of the nature of explanation and confirmation that the Positivists' account seemed to offer.

THE AXIOMATIC ACCOUNT OF THEORIES

I noted previously that the Positivists proposed to differentiate laws from accidental generalizations by appeal to the fact that laws can be grounded in scientific theories. When they spoke of theories, the Positivists generally had in mind such large-scale frameworks as Ptolemy's or Copernicus' astronomy, which offered basic accounts of how various celestial bodies moved with respect to each other; Newton's mechanical theory, which offered a basic set of principles relating the motion and attraction of objects; and the germ theory of disease, which offered an account of what caused diseases and how they spread. The idea underlying the Positivist account of theories is that just as they claimed that an event is explained by showing how a statement

---

[3] Another such paradox was Goodman's (1955) "grue" paradox. By definition, an object is grue if it is green before time t and blue after that time. If time t is the year 2050 and we are looking at an object before that year and it appears green, we cannot conclude that the object is green. It might actually be grue. Give the possibility of predicates like "grue" it is impossible to determine what hypothesis is actually confirmed by current evidence.

about the event could be derived from a law, so a law (e.g., a law about the free fall of bodies on the surface of the earth) is explained by deriving it from a theory (e.g., Newton's mechanical theory that specified the force of attraction between any two objects). A theory was thus a structured network of statements from which one could derive specific laws (see Hempel, 1966, and Nagel, 1961). A model of the kind of structure they had in mind is found in Euclidean geometry. At the core of Euclidean geometry are a set of primitive terms and postulates. From these postulates, various axioms can be derived. In a like manner, the Positivists proposed that scientific theories could themselves be rendered as deductive structures in which we could identify a set of primitive terms and postulates. The particular laws would be the axioms that we could derive from these assumptions and postulates.

Thus, for the Positivists, a theory is best viewed an axiomatic structure.[4] Although the Positivists recognized that most theories are not presented in such an axiomatic fashion, they claimed that theories could be axiomatized and offered thermodynamics as an example of a theory which had already been axiomatized. They argued, moreover, that such axiomatization could be helpful to scientists. First, it would introduce rigor into scientific discourse, forcing scientists to be precise in characterizing notions that might otherwise be left intuitive and hence unclarified. Second, it would allow scientists to discover some of the consequences of a theory that they had not anticipated. This would allow them to make additional predictions so as to carry out additional tests of the theory and to appreciate the full explanatory power of the theory.

The Logical Positivists also envisioned that the process of axiomatizing theories could bring unity to science. Imagine, for example, that the Copernican theory of astronomy has been successfully axiomatized. Someone might question the basic postulates of astronomical theory and demand an explanation of them. How should a Copernican respond? The Positivists proposed that the Copernican should proceed just as in other cases where an explanation is sought—by seeking more general statements from which the Copernican laws could be derived. In this case, the more general statements would not be statements about astronomical phenomena, because the assumption is that the astronomical theory was already complete. Rather, the Copernican would try to generalize beyond astronomy, developing general physical principles that apply not just to astronomy but to all other physical objects. This is the enterprise Newton carried out successfully, showing that the basic postulates of astronomical theory are themselves axioms derived from a more basic physical theory. Astronomy was thus subsumed within physics. Eventually, the Positivists proposed, all sciences could be subsumed into one theoretical edifice, that of unified science.

The process of unifying science by deriving the principles of one science from those of another is commonly spoken of as *theory reduction*. I return to this topic in chapter 5. For now, however, we should note just a few aspects of this view. First of all, it assumes that science is basically a cumulative enterprise. Scientists continually incorporate the results of previous inquiries into ever larger theoretical networks. Second, it views the laws of specialized disciplines, such as physiology or psychology, as derivative laws which, in principle, can be derived from the most basic laws of physics. Hence physiology and psychology, according to this view, will eventually be subsumed within physics as a special application of physical laws. The compartmentalization of science into separate disciplines with their own theories and laws is, for the Positivists, simply a result of the incompleteness of current inquiry. Once we have axiomatized the theories in these disciplines, we will be able to integrate them into one broad account of nature.

SUMMARY OF LOGICAL POSITIVISM

The Logical Positivists offered a systematic and highly attractive view of the project of science. They proposed a theory of meaning that showed how scientific discourse was grounded in sensory experience and thus certain to be meaningful. They provided an account of explanation that used deduction to show how particular events could be explained by laws and an account of confirmation that showed how particular events provided evidence for the laws that were developed. Finally, they showed how the laws of each science could be unified in to axiomatic structures and ultimately grounded in a unified account of nature. Many have found this view of science very attractive. (For further details regarding the Positivists' program, see Suppe, 1977, and Brown, 1979. For a useful collection of readings from many of the Positivists, see Ayer, 1963.)

In concluding this chapter it is worth briefly indicating how these doctrines of Logical Positivism might apply to cognitive science. A number of psychologists of an earlier generation adopted the Positivists' conception as a guide for developing their own science. A variety of the doctrines discussed here had significant impact on a number of behaviorists such as Spence and Skinner. This is particularly true of the verificationist theory of meaning, which was taken by many behaviorists to show the illegitimacy of positing or recognizing mental events except insofar as they could be explicitly linked to observable behaviors. The deductive-nomological model of explanation and the axiomatic view of theories also had profound influence on such

---

[4] The Positivists and their descendants have not been in total agreement on the virtues of axiomatization. Suppes (1968), Kyburg (1968), and Feigl (1970) have been strong advocates of axiomatizing theories, while Hempel (1970) points to the limits of such an approach. Generally, axiomatization has been most favored by those focusing on examples from physics, but some philosophers of biology (notably, Williams, 1970, and Rosenberg, 1985) have been strong advocates of solving problems by first axiomatizing the theories in question.

behaviorists as Clark Hull, whose learning theory was a highly developed axiomatic structure. Moreover the hypothetico-deductive method of theory development has been emphasized in the teaching of scientific methodology in psychology throughout the reign of both behaviorism and cognitivism.

Although Positivism has been less popular during the recent reign of cognitivism, one can illustrate the basic claims of the Positivists equally by showing how they could be applied to theories of recent cognitive science. A central field of research over the past decade has been the structure of human concepts and the processes of categorization. Philosophers and others have commonly thought of categories as mathematical sets, that is, structures with well-defined conditions of membership (generally referred to as *necessary* and *sufficient conditions* of membership). This idea was already challenged by the philosopher Wittgenstein (1953), who used the example of the concept "game," and questioned what could be the necessary and sufficient conditions for being a game. He argued that there were no necessary and sufficient conditions that defined the category games, but that games were only related by "family resemblance." To develop a scientific theory of concepts, the Positivists would insist, it is necessary first to provide criteria for the meanings of terms used in the theorizing, especially terms like *concept*. Eleanor Rosch (1975, 1978) developed a tool that might be viewed from a Positivist's perspective (although probably not from Rosch's) as providing an operational definition of concepts. She asked subjects to evaluate the typicality of particular instances of category. From this she showed that with a variety of categories, both natural and artificial, subjects would not only quite willingly evaluate how typical a particular example of the category was but also that there would generally be a high rate of intersubject agreement (but see Barsalou & Sewall, 1984). Thus, most Americans would judge a robin to be a very typical bird, and a chicken to be quite atypical. From the Positivists' perspective, the typicality measure can be construed as providing a basis in observation in terms of which we can understand the meaning of the mentalistic notion of a concept.

As we saw, Positivists construed the task of science as explaining and predicting phenomena in nature via laws. From a Positivist's perspective, then, we need to see if cognitive scientists have been able to come up with hypotheses about the behavior of concepts. One hypothesis that has been offered by a number of cognitive scientists is that concepts are stored in the mind as prototypical instances and a metric in terms of which new instances are compared to the prototype. Thinking involves bringing this concept into working memory and formally manipulating it in some rule governed way. The hypothetico-deductive method requires that once such a hypothesis is advanced, a variety of consequences must be derived from it that can be tested. One prediction is that subjects should be slower in making judgments about less typical instances of a category than more typical instances since they would be further from the prototype, a result that is confirmed by the data of Rosch (Rosch, 1975; Rosch & Mervis, 1975)[5] and others (see Smith & Medin, 1981, for a review). A further prediction is that there will be higher agreement between subjects when asked to judge whether objects closer to the prototype are members of a category than those further away. This result is also borne out (e.g., McCloskey & Glucksberg, 1978). From a Positivist's perspective, at this point the hypothesis has been supported by tests, but it has not yet been taken up into a general, axiomatizable theory. But one can anticipate the kind of theory that the Positivists' account of science would envision. Researchers would need to embed the idea of prototype and metric into a general theory of the structure of the cognitive system so that the idea of concepts being encoded as prototypes with a metric would be a derivable consequence.[6]

It does seem reasonably easy to take instances of work in cognitive science, such as the work on concepts and categories, and explicate it using the Positivists' account of what a science is. On the surface, the Positivist's account seems to offer a compelling account of the character of scientific explanations. Despite its initial force, I noted in the course of this chapter some problems that confronted the full articulation of this view. During Positivism's heyday many felt that these could be resolved and we would then have a clear account of the nature of scientific inquiry that could justify the claim of science to provide knowledge of nature. In recent decades, however, this optimism has waned. Although some philosophers remain convinced that the basic picture of science offered by the Positivists is correct, many have found the objections to be fatal and have begun to look for alternatives. In the next chapter I begin to examine the criticisms that have been raised against Logical Positivism, whereas in chapter 4 I discuss alternatives to the Positivist conception.

---

[5] Rosch (1978) explicitly disavows the idea that her data should support the claim that concepts are stored in terms of prototypes and metrics, although many psychologists have so interpreted her.

[6] For a quite different perspective on the significance of this work on concepts and categorization, see chapter 6.