3

Challenges to Logical Positivism

INTRODUCTION: CHALLENGES TO SPECIFIC THESES OF LOGICAL POSITIVISM

One of the virtues of Logical Positivism is that it is a view of scientific explanation that is articulated in clearly stated, specific doctrines. This enabled careful analysis and criticism. The result has been a number of objections that have focused on specific doctrines. These criticisms served to chisel away at the foundations of Positivism, and subsequently there have been broader criticisms that attack the program as a whole and recommend replacing it with a quite different approach to philosophy of science. In this chapter I look at a number of the more specific objections, reserving for the next chapter an examination of the major alternative approaches to philosophy of science that have been advocated in recent decades.

THE ATTACK ON CONFIRMATION

One of the first major challenges emerged from Karl Popper, a philosopher who was often in dialogue with members of the Vienna Circle and who was viewed by them as a friendly critic. Like the paradoxes of confirmation discussed in the previous chapter, Popper's criticisms focused on the hypothetico-deductive model of a hypothesis development and testing. The Positivists, as I noted, were aware that they could not solve the problem of induction, but thought that at least positive tests of a hypothesis could give support to that hypothesis. Popper (1935/1959 and 1965) contended that this

assumption was false. Just as Hume had shown that it was never possible to prove that a general statement was true, Popper maintained that one could not even show that it was likely to be true. The degree of support we could offer to a general statement was always dependent on the few cases we could examine. If we treat our degree of confirmation of a general claim as a probability measured by the proportion of all cases that we have actually examined and that have turned out positive, we could never establish that a general hypothesis had greater than a very small probability of being universally true.

Popper proposed a radical remedy to this problem. He recommended abandoning the whole endeavor of seeking well-confirmed theories, and proposed instead that scientists focus on demonstrating that some hypotheses are false. His reason for making this shift was that although there is no valid deductive or even good inductive form of reasoning from which we could derive a general proposition from specific instances, there is a valid argument form according to which we can show that evidence disproves an hypothesis. This form is *modus tollens* or denying the consequent:

If *H*, then *P*
Not P
———
Therefore, not *H*.[1]

If H is the hypothesis in question and P is a prediction that follows from it, then when P turns out to be false we can infer that H is false.

Popper went on from this simple logical point to propose a general schema for scientific inquiry. Empirical investigation, he proposed, should attempt to disprove hypotheses. Although this sounds defeatist, Popper has tried to show how it could be the basis of a constructive enterprise, which could result in continual improvement in the accuracy of our science. He characterizes his program as one of conjectures and refutations. A scientist should begin by making conjectures about how the world is and then seek to disprove them. If the hypothesis is disproved, then it should be discarded. If, on the other hand, a scientist tries diligently to disprove a hypothesis, and fails, the hypothesis gains in stature. Although failure to disprove does not amount to confirmation of the hypothesis and does not show that it is true or even likely to be true, Popper speaks of such an hypothesis as *corroborated*. The virtue of a corroborated hypothesis is that it is at least a candidate for being

[1] If we tried to create a similar argument to confirm a hypothesis, we end up with the invalid form of affirming the antecedent:

If *H*, then *P*
P
———
Therefore, *H*.

See the discussion of logic in chapter 1.

a true theory, whereas hypotheses that have been disproved are not even candidates.

The substitution of the term *corroboration* for *confirmation* is not a mere linguistic change for Popper. It represents a broader change and rejection of other aspects of the Positivists' schema. It signals a rejection of the Positivists' attempt to distinguish meaningful from meaningless discourse through the verificationist theory of meaning. Popper substitutes a quite different program of demarcation—one that demarcates scientific from nonscientific discourse in terms of the risk that true scientific theories face of being wrong. In taking the risk of being wrong, the theory is, according to Popper (1965, chapter 1), "forbidding" certain things to happen. That is to say, if the theory is true, then certain things cannot happen. If they do happen, then the theory was not true. This ability to forbid certain things is what gives scientific theories their power, for they distinguish things that are possible if the theory is true from those that are not. The more things the theory rules out, the more powerful the theory. We can speak of it as more informative—it is telling us that, of the circumstances that we otherwise think might be possible, only this limited subset is actually possible. Thus, the more the theory rules out, for Popper, the stronger the theory actually is, for it tells us more about how the world actually is.

Although Popper intends to apply this thesis primarily to theoretical statements, we can illustrate his point with a straightforward, observational claim. There are a variety of conditions that could make the sentence "It is sunny today" false. The truth of the sentence thus informs us that those circumstances that would make the sentence false do not occur. Thus, an informative sentence rules out possible configurations of the world and a highly informative sentence rules out large numbers of ways the world might be. In contrast, a sentence that is certain or almost certain to be true ("It is either sunny or not sunny today") rules out little and so tells us little about how the world really is.

To apply this view to theoretical statements, Popper appeals to the testability of scientific theories. True scientific theories are ones that can be put to critical tests where we can specify in advance what would count against the theory. An example Popper used was the test of Einstein's general theory of relativity, which predicted that light from stars would be distorted in a particular way when the light passed near to the sun. This can only be tested during an eclipse, and Popper greatly admired Eddington's test of Einstein's theory during such an eclipse. Because the prediction was risky it was quite conceivable that it would come out false. On the other hand, Popper found Freudian and Adlerian psychology to be incapable of such tests because both theories could be applied to any possible circumstance that arose. Hence, for Popper, they were not informative, and hence not true scientific theories even if they are meaningful by the canons of the Positivists.

In the process of conjecture and refutation, Popper recommends starting with our current theories. Some of these have passed all tests to date, whereas others have failed one or more tests. These latter theories present problems to the theorist, and the theorist's task is to propose new theories that solve these problems.[2] Popper then imposes three kinds of constraints on such new theories, two of which focus on the theories themselves, while the third focuses on the results of testing the theories. Popper (1965) requires first that the "new theory should proceed from some *simple, new, and powerful, unifying idea* about some connection or relation (such as gravitational attraction) between hitherto unconnected things (such as planets and apples) or facts (such as inertial and gravitational mass) or new 'theoretical entities' (such as field and particles)" (p. 241). The reason for this requirement is that the goal of theorizing is to get at basic features of nature and not merely to posit ad hoc relations so as to solve problems with previous theories. The second requirement is that the new theory should be independently testable, by which Popper means that it should have new testable implications that have not been previously investigated. Here is where the theory must be bold and risky. Moreover, it is as a result of possessing such implications that the theory promises to tell us something new.

A theory will tell us something new even if its new implications are proven false, because we will have learned new empirical findings that future theories must accommodate. But Popper places a third requirement on a good theory: It must pass some of these new, severe tests. The reason is that progress would be stifled if we continued to offer new theories without them passing at least some of these new tests. For one thing, if the new tests constantly turned out negative, we would no longer have any clear idea at what points the theory was in accord with nature and where it was making new, risky proposals. Thus, we would lose the ability to devise crucial tests that could distinguish true from false theories. A crucial test involves finding a situation where a theory that generally fits the natural world can be precisely compared to it in a well-defined circumstance so that a specific piece of the theory could be falsified or further corroborated. Such tests, however they turn out, give us precise information, but they are not possible unless our theories are, for the most part, accommodating our experience.

What the process of conjectures and refutations offers, for Popper, is a means of continually homing in on the truth. We are able to reject ideas that turn out false and use theories that continue to be corroborated. Although

[2] Popper acknowledges that the problems to be solved might not stem simply from empirical falsifications of previous theories, but may be theoretical problems involving how to unify two theories or dispense with seemingly ad hoc principles in current theories. This interest in theoretical problems is developed much more fully by Laudan, whom I discuss in chapter 4.

such corroboration does not give us a logical basis for increasing our confidence that the hypothesis will not be falsified on the next test, it does serve to limit us to an ever narrowing set of hypotheses that might be true. Popper has compared this process to natural selection, for both nonadapted organisms and false theories are weeded out, leaving the stronger to continue in the competition. But Popper points to an important difference. If we are maladapted to our environment, it is we who perish, but if our theories are maladapted, we can let them perish while we pursue better theories. Popper thus speaks of letting our theories die in our stead. (Other theorists have pursued the idea that theory development may be parallel to the process of evolution by natural selection. See, for example, Campbell, 1974a; Toulmin, 1972.)

The implications of Popper's falsificationism for disciplines in cognitive science are rather serious. There are cases in cognitive psychology, for instance, where competing theories do make contrasting empirical claims. For example, Sternberg (1966) compared three different models of how humans might access items stored in memory. Each model predicted a different pattern of reaction times, so Sternberg had subjects memorize a list of numbers and then asked them to identify whether particular numbers were on the list. The experiment thus pitted these models against each other so that those models that did not fit the reaction time data could be rejected.[3] However, it is widely recognized that many theories in cognitive psychology and linguistics, and most simulations in artificial intelligence, fail to accommodate the data already available. It is not necessary for them to make radical new predictions in order to be tested and possibly falsified, for they are clearly already falsified. On Popper's grounds, these theories should be rejected.

The problem, for cognitive science, is that it is very difficult to create theories to accommodate existing evidence. Perhaps, however, what is at fault is Popper's claim that the only way to advance science is by conjectures and refutations, not the practice of cognitive science. This is suggested by some of the psychology of reasoning studies. Taking a lead from Popper, some early studies[4] focused on the fact that humans do not rigorously try to falsify

[3] The three models were that (a) the mind viewed all the numbers at once to determine if the number was on the list; (b) it searched the numbers sequentially and stopped when it reached the queried number to report a positive result and reported a negative result when there were no more numbers on the list; and (c) it searched the list sequentially, but reported neither a positive nor a negative result until it reached the end of the list. The first two models were rejected on the basis of reaction times, whereas only the third, rather counterintuitive, model was corroborated.

[4] One study performed by Wason and Johnson-Laird (1972) asked subjects to try to figure out the rule that lies behind the sequence 2, 4, 6. The rule might be something general like "any three positive whole numbers." The subjects are encouraged to offer additional sequences to test their possible rules before announcing them. Typically subjects begin by assuming the rule is "three sequential even numbers." Rather than trying to falsify this hypothesis by offering

hypothesis, but rather mistakenly go about seeking confirming data. Other studies (see Mynatt, Doherty, & Tweney, 1978) suggest, however, that seeking confirmatory data for early hypotheses may be quite reasonable in a variety of problem environments where the character of the environment and the nature of the problem are not well understood. It may be necessary to acquire some concrete data and structure plausible models before it makes sense to engage in rigorous falsification. The research endeavors in many parts of cognitive science may be at just such a state where exploratory, nonfalsificationist research is required.

Regardless of these difficulties in applying Popper's falsificationism to current cognitive science, it is important to note how his approach to philosophy of science marks a significant break with the picture of science offered by Logical Positivism. In giving up the quest for highly confirmed theories, Popper also foregoes the program of building an ever larger theoretical structure of well-confirmed propositions. Rather, he directs science to make new, bold conjectures that correct for previous failings rather than amplifying and rendering more adequate a theory already supported by a number of confirmations. Thus, he begins to undercut the cumulative conception of science that emerges from Logical Positivism and opens up a concern with how science changes. On the other hand, there are still affinities between Popper's conception of science and that of the Positivists. Although he draws different consequences from the use of logic, and does not present axiomatization as the goal of theorizing, he still holds that modern logic can provide a framework for analyzing scientific investigation and that scientific investigation is to be grounded in empirical inquiry. For example, the focus is still on the context of justification, not discovery, and explanation still involves deriving a fact from a law and a set of initial conditions.

a sequence like "3, 2, 1," however, they try to confirm it by offering sequences like "8, 10, 12." Incidently, this problem can be used to create a useful simulation of the scientific process. This can be done by letting a group represent members of a scientific community competing for the Nobel Prize, with one person (e.g., a teacher) playing Mother Nature. Any member of the group can perform an "experiment" by offering up a sequence, to which Mother Nature gives the "result" by saying "Yes" (e.g., it fits the rule) or "No." Any member of the group can also publish a hypothesized law by asserting it. As in real science, Mother Nature does not say whether the law is right or wrong, but leaves that to be determined by other members of the group through further experiments. The simulation ends when the group chooses someone to receive the prize. During the simulation, both the laws and experimental results may be written on a blackboard so that after the simulation the participants can go back through the events sequentially and discuss the strategies that seem to be at work at each stage in the problem solving activity. One thing that can be pointed out at the end of the simulation is that the rule that the group ends up with may not be the "right" one (i .e., the one Mother Nature was using). In science, even after the Nobel Prize is awarded, it is still possible that we will learn that the hypothesis is wrong.

REPUDIATION OF THE DEDUCTIVE–NOMOLOGICAL MODEL OF EXPLANATION

Although both the Positivists and Popper accepted the correctness of the deductive–nomological model of explanation, it has been criticized by numerous philosophers. These critics contend that something other than deduction from laws and statements of initial conditions is required of explanation. Two alternative views have emerged; one of them requires that an explanation identify a cause for the event being explained, whereas the other treats explanation as a matter of answering certain kinds of questions. On both of these views, it is claimed that an event may be explained even if a description of it has not been derived from more basic laws.

The appeal to causes as explanations is partially motivated by dissatisfaction with the treatment of statistical explanation within the Positivists' framework. I noted previously that the Positivists attempted to generalize the D–N model to handle statistical explanation by holding that in a statistical explanation it is possible to derive a statement that the event in question is likely to occur. One objection to this approach is that it only works with events with greater than .5 probability. It does not allow us to explain relatively low probability events, such as getting cancer after a lifetime of smoking. But often the goal is to explain low probability events. Salmon (1970, 1984) contends that if we view explanation as a matter of identifying causes, we can handle such situations very naturally. He proposes that commonly we identify such causes by a process of *screening off*. To explain the idea of screening off he offers an analogy to a situation where we might be trying to find the source of light. If we can find a place where we can put a screen to cut off the light, then we have identified where the light is coming from. Generalizing, his idea is that if we can identify a way of interrupting an effect, we have found the causal chain. That which we have interrupted can be taken to be part of the causal chain and we can continue up the causal chain until we find the source that generates the causal chain.

The idea of screening off works equally well in probabilistic situations as in deterministic situations. If we have a situation where the effect (for example, a disease) is found in 10% of the population, and we take action that interrupts the causal pathway, then we will reduce the percentage of the population in which the effect occurs. We will then be able to identify the cause which led to 10% of the population initially suffering the disease. Moreover, this idea also works with partial causes. If we do not reduce the percentage to 0, but only to 5%, then we can infer that we have interrupted one of the causal pathways but not all. To capture what is happening in situations of these kinds, Salmon introduces the idea of statistical relevance. With unrelated events, the probability of both events occurring together is simply the product of the probability of each occurring. When the probability of both events occurring together is greater than the product of each occurring, then Salmon speaks of the two events being statistically relevant to each other. An example of such statistical relevance is given in the case of smoking. The probability of being a smoker and dying prior to age 70 is greater than the probability of death prior to age 70 times the probability of being a smoker. In such cases, Salmon maintains, either one event is the cause of the other, or both are effects of a common cause. Once we identify the existence of a common cause situation, we can use techniques like screening off to find this common cause.

One factor that is particularly noteworthy about Salmon's approach to explanation is that explanation does not involve a deduction of a statement of the event to be explained, as the Positivists had maintained. Explanation is not a relation between laws and statements of events at all; rather, it is an answer to a "why" question that cites the event or entity that is causally responsible for the phenomenon. (In offering this view, Salmon rejects the Positivists' approach of examining only the language of science. Instead, Salmon directs his attention to the events in the world and their causal interactions.) Salmon's approach to explanation also differs from the Positivists in giving a central role to causation. In the Positivists' account, the concept of cause had no special status and one might well explain an event without knowing its cause. Some laws invoked in explanations might state causal relationships, but that was not required. In fact, the Positivists had no resources for distinguishing causal laws from other generalizations embedded in the axiomatic structure of a theory. By focusing on procedures whereby one can interrupt causal chains, Salmon proposes to identify causal relations and allow them to assume a central place in an account of explanation.

The second alternative to the D–N model begins by focusing on the context in which people seek explanations. Both Bromberger (1968) and Scriven (1962) contend that explanation begins when someone asks a question because he or she is missing some information. What counts as an explanation will depend on what information the person is lacking. For example, a person may ask why a certain window broke. What the person may not know is that it was hit by a baseball. In this case, telling the person what hit the window will suffice for an explanation. In other contexts, the person may know that the window was hit by a baseball, but will want to know why being hit by a baseball led to the window's breaking. Then the person will be seeking information about the structure of windows that explains why they break when hit with certain kinds of objects whereas other objects do not cause breakage.

One result of placing the issue of explanation in this context of answering questions is that giving an explanation may not always be the same kind of activity. Asking for an explanation may involve asking for quite different things in different contexts. Sometimes it may be a request for a scientific theory and a theoretical account about how things happen, but not always. Sometimes it may be a request for an identification of an unknown cause. In general, however, the critics maintain that when a person asks for an explanation, the person does not require a derivation of a description of the event from a general law and a statement of initial conditions. It is here that the break with the Positivists appears, for they held that one did not have an explanation if one did not have the deductive structure. The Positivists did acknowledge that when, in practice, a scientist is asked for an explanation, he or she might not provide the whole deductive account, but only part of it. This, however, they viewed as a matter of shorthand, believing that to explicate the character of the explanation, one had to offer the whole deductive structure. The critics contend, however, that such a deductive structure need not be present even implicitly in an adequate explanation.

Bromberger and Scriven make a similar point about theoretical explanation. For the Positivists, a theoretical explanation involved development of an axiomatic structure in which the particular law was derived from more basic axioms. Both Bromberger and Scriven contend that even when the questioner asks for a theoretical account, the request may not be for an axiomatic structure. In fact, often producing such an axiomatic structure will leave the questioner asking for further explanation. The person may be able to handle the derivation within the axiomatic structure and still not feel that he or she understands the event. To explain the event to this person may require providing a model of the phenomenon in question so that the person can see how different factors, as described by the equations in the axiomatized theory, interact with one another.

In some disciplines of cognitive science (e.g., linguistics) there are large theoretical structures that can be viewed as giving explanations comparable to D–N explanations. The computer programs that underlie cognitive simulations may also have sufficient logical structure that it may be plausible to view them as subsuming the behavior to be explained under laws in a D–N style. But in other disciplines relevant to cognitive science, such as cognitive psychology, the explanatory endeavors are often far more modest. Researchers do not have elaborated theories from which they can derive specific behaviors once given initial conditions. Rather, a particular aspect of behavior is noted, and researchers try to identify a factor that might explain it. For example, Tulving's (1983) distinction between an episodic memory system (devoted to memory for particular events) and a semantic memory system (devoted to memory for general propositions such as those about the meanings of language or general facts about the world) can be viewed as part of an attempt to explain different sorts of results that can be produced on various memory tasks. The theory is not presented axiomatically in a full set of laws from which particular events are presented, but much more informally by describing the differences in the two proposed memory systems and then

showing that these differences would account for the different results on memory tasks for the two kinds of memory. What Tulving did, then, can be seen as an attempt to isolate a potential causal factor that affects memory behavior without creating a D–N structure for deriving statements of memory behavior from general laws and statements of initial conditions. Insofar as Tulving's work is fairly typical of explanatory endeavors of cognitive psychologists, these endeavors might be better understood from the perspective of the critics of the D–N model, rather than trying to force them into the structure of D–N explanations.

These challenges to the deductive–nomological model have reduced the desire of some philosophers to construe all explanations as deductive and all scientific theories as objects to be axiomatized in the manner of Euclidean geometry. However, they have also led to a situation where there is no clear and widely accepted model of what a law or theory is or how they are to figure in explanation. The core of the Positivists' view of science has been attacked, but there is no general agreement on a replacement. But, as I show next, even more basic assumptions upon which the Positivists' view was built have been brought under attack, suggesting that the kind of philosophical enterprise in which they were engaged is misguided.

CRITIQUE OF THE ANALYTIC–SYNTHETIC DISTINCTION

In the previous chapter we saw that in applying the verificationist theory of meaning the Logical Positivists invoked *analytic statements*. These statements often are characterized as statements that are true in virtue of the meanings of the words contained in them. Hence they are not dependent upon evidence. In this respect they are distinguished from *synthetic statements* that make substantive empirical claims for which evidence is appropriate. The distinction between analytic and synthetic statements was fundamental to the Positivists' enterprise and indeed to much of analytic philosophy, which has tried to resolve philosophical problems by analyzing the meanings of important concepts. In addition to analytic sentences that spelled out the meaning of theoretical sentences, the Logical Positivists were also concerned with two other classes of analytic propositions—mathematical propositions and logical propositions. Because the truth of these statements could be secured independently of experience, they could be used in developing a science, or in articulating the philosophical foundations of science, without any risk of introducing error. Thus, as we have seen, the Positivists freely employed symbolic logic in their analysis of science, assuming that these logical underpinnings were not themselves dependent for their truth upon the substantive empirical claims of the science.

In the early 1950s, however, W. V. O. Quine launched an attack on the distinction between analytic statements and synthetic statements that has had quite broad ramifications both for philosophy of science and analytic philosophy generally. Quine's strategy is to show that the term *analyticity* can only be defined in terms of other concepts like *meaning*, which in turn can only be defined in terms of analyticity. The result is a vicious circularity. Thus, terms like *analyticity* and *meaning* fail to meet the Positivists' own standard of meaningfulness according to which only terms which have clear standards for verification count as meaningful (see the first section of chapter 2). Quine's (1953/1961b) conclusion is that insisting on a distinction between analytic statements and other statements in any language, formal or natural, is "an unempirical dogma of empiricists, a metaphysical article of faith" (p. 37).

Quine draws upon this rejection of the analytic–synthetic distinction to launch a full attack on the general program of the verificationist theory of meaning. He treats the assumption that all sentences of a language should be reduced individually to experience as a second unempirical dogma of empiricism and proposes to replace it with a view according to which "our statements about the external world face the tribunal of sense experience not individually but only as a corporate body" (p. 41). He contends that the terms of our language are interconnected with one another in a vast network, so that we cannot differentiate between those connections in the network that establish the meanings of theoretical terms from those that present empirical findings.

The conclusion Quine draws from this is quite startling and revolutionary. It goes far beyond attacking just the verificationist theory of meaning. He contends that we must give up the idea that we can use experience either to confirm or to falsify particular scientific hypotheses. When experience contradicts our science, we must modify our science, but Quine's contention is that we can do this in a wide variety of ways. We can always protect particular hypotheses by modifying others. For example, if someone who believes that only humans can use languages is confronted with the reports of primate linguistic capacity, he or she can either take these reports as refuting evidence or narrow the conception of language so as to protect the original claim. Without a fixed analysis of meaning, the fundamental idea of either confirming or falsifying scientific hypothesis on the basis of evidence is called into question. (This claim that evidence does not itself determine our evaluation of hypotheses is commonly referred to as the "Quine-Duhem thesis" since the physicist Pierre Duhem, 1906/1954, had advanced a similar claim a half century earlier.)

As I noted in the previous chapter, the concept of an operational definition was an extension of the Positivist's verificationist theory of meaning. The demand to define terms operationally has been prominent in the social sciences, including psychology. One benefit of this demand has been to tie

psychological theorizing to experimental results and avoid speculative theorizing that borders on meaninglessness. But if Quine's challenge is correct, then these operational definitions must be viewed in a different light than they commonly are. They should not be taken as specifications of meaning that are absolute and unmodifiable, but rather they should be seen as synthetic claims, subject to revision in the course of inquiry. That is, when a theoretical claim is jeopardized by an experimental result that is based upon operational definitions of theoretical terms, one option that should be considered is revising the operational definition to save the theoretical claim. One result of opening this option is that judgments about theoretical claims will not be precisely determined by experimental results in the manner that seemed possible given the use of operational definitions, but Quine would maintain that this kind of ambiguity is something we must accept, and not try to legislate against.

One consequence of Quine's attack on the analytic–synthetic distinction is that decisions about modifying scientific claims must be treated as pragmatic, not logical decisions. A further result is that, in deciding where to modify our theoretical structure in the face of negative evidence, we may choose to modify the propositions of logic and mathematics as well as those more generally thought of as part of empirical science. Thus, these principles lack the privileged status that the Logical Positivists assumed they possessed. Quine acknowledges that typically we will be loath to modify the principles of logic, but he claims this is not due to any privileged status they enjoy. The reason is rather that they are so "central" to our network of scientific claims that modifying them will have enormous consequences throughout our science and our conceptual systems generally. Thus, it is the pragmatic principle of conservatism, not privileged status, that explains why we find rejecting the principles of logic and mathematics nearly inconceivable.

The consequence of depriving logic of any privileged status, however, is enormous. The Positivists' program is built on the assumption that logical analysis is inviolable and that philosophical analyses that depended on logic alone could not therefore be supported with empirical investigation. Without such a foundation, philosophy loses its claim to be able to inform us simply through the vehicle of logical analysis. In light of this, Quine claims that we must forego any hope for a "first philosophy" that is independent of and established prior to experience. Rather, we must treat philosophical discussions of science as themselves a part of empirical inquiry (Quine, 1969b, 1975). Thus, Quine advocates a "naturalized epistemology" that is integrated with work in psychology. (For Quine, however, this psychology is to be behavioristic, not cognitive. See Bechtel, in press a, chapter 3.) As part of psychology, philosophical claims about how we can have knowledge of the world and how science can produce knowledge are just as open to revision as any part of science.

Needless to say, Quine's arguments have not been universally endorsed by analytic philosophers whose *modus operandi* is challenged by these arguments (see Grice & Strawson, 1956; Katz, 1964; Putnam, 1962). However, a number of philosophers have taken Quine's challenge to heart and have come to accept a need to base philosophical analysis on empirical findings. A generation of philosophers of science have accepted the importance of studying both the history of science and contemporary science and to develop accounts of science that fit the actual mode of scientific inquiry. One cost of working within such a naturalized framework is that these philosophers can no longer claim to be strictly normative, prescribing how science ought to be done if it is to have logical validity. Instead, they must settle for being pragmatic commentators on science, explicating what is revealed by empirical inquiries about science and making pragmatic recommendations as to how scientists might learn from their past practices and better realize their objectives.

A CHALLENGE TO THE OBSERVATION-THEORY DISTINCTION

Yet another critical assumption made by the Positivists was that through empirical experience we had an objective basis for evaluating scientific claims. Such empirical experience was acquired through observation. Because the Positivists characterized the evidence relationship in terms of relations between sentences, they required that the evidence acquired from experience be coded in a set of sentences. This set of sentences had to be distinct from the set of sentences stating the laws or theories to be tested so that they could be appealed to in the course of such tests. Quine's attack on the analytic-synthetic distinction began to call this assumption into question, for he challenged the idea that we could clearly differentiate between the empirical claims of a science and the meanings of the terms used to present these claims. He proposed that science constituted a network that could be adjusted at various points, with no point being privileged. Yet, Quine remained sympathetic to the basic empiricist view that our theories of nature must be grounded in sensory experience and he offered criteria to differentiate a class of observation sentences from the theoretical networks built upon them. These sentences, he claimed, could provide a neutral basis for testing different theoretical claims. However, the assumption that there is an objective set of observations that are neutral between theories has increasingly come under attack as a number of authors have argued that observation itself is theory-laden.

One of the strongest proponents of the theory-ladenness position was Hanson (1958), who maintained that what we perceive is influenced by what we know, believe, or are familiar with. He offered examples where we directly

perceive certain objects that would not have been recognized at all by members of previous generations because they were unfamiliar with such objects. A simple contemporary example would be a microcomputer, which all of us recognize immediately but which would have been unrecognizable to anyone 50 years ago. We directly see microcomputers because of the background knowledge we bring to the perceptual context. When Hanson argues that we *directly* perceive such objects, he is rejecting the claim that we are seeing a neutrally characterizable object and then making an inference about what the object is. Sometimes we do make inferences based on what we see, but this is not the usual situation when we see ordinary objects or even special laboratory objects with which we are familiar. Rather, we see what our knowledge and training equips us to see. Hanson's position here is, by now, one quite congenial to many cognitive scientists who argue for some degree of top-down processing through which knowledge enters into observation. In the classical experimental paradigm, the stimulus is presented in a situation where contextual information directs subjects to one particular interpretation of the stimulus, even though others were possible. For example, in the display in Fig. 3.1, the middle letters of the two words are drawn identically, but nearly everyone sees the first as an "H" and the second as an "A".[5]

FIG. 3.1. The middle letter in each word is printed the same and without context would be ambiguous. In context, however, it is viewed as an "H" in the first word and an "A" in the second (figure after Selfridge, 1955).

For Hanson, part of what is involved in learning a particular science is learning to see the world in a particular way. Hanson proposes that the difference between the trained observer and the untrained observer is similar to the gestalt shifts that any of us can experience when we look at ambiguous figures. Figure 3.2 provides a classic example. Many of you can see both an old woman and a young woman in this figure. But you cannot see both figures at the same time. Looking at it one way you see it *as* a young woman, looking at it in a different manner you see it *as* an old woman. The expression "see it as" suggests that you are making an inference, but this is what Hanson resists. We do not see the figure neutrally as a set of lines and then infer that one curve represents the nose, and so forth. To see the figure as a set

[5] Hanson, however, would resist the temptation to characterize the information processing that goes into recognizing objects in terms of inferences and problem solving.

of lines in fact takes sophisticated training, which most of us lack. Rather, we see the figure of the young woman or the old woman.[6]

FIG. 3.2. An ambiguous figure that can be seen either as an old woman facing forwards and to the left or as young women looking away to the left. (figure after Boring, 1930).

Hanson, however, recognized that the consequences of this view for the Positivists' conception of science are quite serious. It shows that observation does not offer a neutral basis for evaluating theoretical frameworks, but is itself influenced by the theoretical framework a scientist brings to the situation. Those who bring different theoretical frameworks to the same situation will see the world differently. Hanson points to the example of Tycho Brahe (who still believed in a form of earth-centered astronomy) and Copernicus (who introduced the sun-centered astronomy) viewing a sunrise. Tycho, according to Hanson, sees the sun rising, whereas Copernicus sees the earth turning toward the sun. Recognizing such evidence involves, he claims, being trained to see certain features of the world in a certain way (i.e., in accord with a particular theory). This is even true in naturalistic settings, such as in the observations of Tycho and Copernicus, where no sophisticated experimental equipment is employed. It arises even when one is not adjudicating between competing theories. It would arise in a situation where a therapist had to evaluate whether a particular mode of therapy has changed a person's behavior. This requires the ability to identify the behavior and recognize specific changes in it, something a lay person may be ill equipped to do. But the influence of theories is even more clear in the case of experiments involving elaborate equipment to report the data. The tracing on an EKG machine will be immediately recognized and understood in terms of heart

[6] The difficulty with the inference proposal can be appreciated when you try to help someone who cannot see the figure in one of the alternative ways and you try to point out features. Although that may help, it does not guarantee success. For many years I could not see the young woman, although numerous people had tried to point out the features. Then in frustration I threw a copy of the diagram upside down in front of a colleague and demanded to know again where this obscure image was. As the drawing hit the table, the picture of the young woman clearly came into view—I simply saw her, I did not infer anything.

activity by a trained technician whereas it will just be a pattern of lines to the lay person.

This challenge to the theory-observation distinction has consequences for cognitive science as well. The data for cognitive science is generally assumed to be behavior. The task of theories is to account for this behavior. Behavior is taken to be something relatively objective, against which we can test the implications of different theories about how the mind works. But the theory-ladenness objection maintains that behavior is not so clearly objective. How we classify behavior may depend upon the theory we are using to try to understand the behavior. We can capture part of the difficulty by considering how we might distinguish between action and mere bodily movement. For a cognitivist, this distinction presumably depends on the mental states we are attributing to the person. If we attribute appropriate desires and beliefs to a person, then the act of raising an arm may appear to be an action—something done in order to achieve the desire. But if we have a quite different theory about what is going on in the person's head, the same bodily movement may be ascribed a quite different significance, either as part of a different action, or as mere behavior that is not itself an action. Although such difficulties are most serious when dealing cross-culturally, where we lack an appropriate understanding of the cognitive perspective of the other person, it arises also in standard experimental designs. A common problem is to determine how the subject understood the task he or she was performing. Given a different understanding of the task, the behavior of the subject, which provides our observational evidence, may be appropriately described much differently.

The claim that all perception is theory-laden has been viewed by some philosophers (Scheffler, 1967; Shapere, 1966) as having grave implications for the appraisal of science. They have portrayed it as undercutting objectivity and rendering science totally subjective. The objectivity is undercut by the fact that we must already accept part of the theory in order to make the requisite observations (e.g. , we must accept the classification of behavior or the theory behind the EKG machine in the examples in the previous paragraph). If we lack an objective, theory-neutral reference point, it is claimed, scientists who hold competing theories will simply see what they are prepared to see by their theory and there will be no theory-free reference point to which we can refer to settle disputes.

These fears are, however, almost certainly exaggerated. Those who hold that all perception is theory-laden are not discounting that the stimuli from the external world play a critical role in determining what we see. These stimuli are factors constraining perception and mark a critical difference between imagination and perception. Theory-ladenness does not entail that we can see whatever we want to. Given the way we have been trained to see, what we see is determined by what is there to be seen. For example, given that we has been trained to read an EKG machine, we read the pattern that is there. We may see things that go against our theories and be forced to revise them. For example, you may have formed an hypothesis that on a particular reaction time measure two tasks would require identical times. When you carry out the experiment, however, you may find that the reaction times for one task are reliably longer than those for another. Then (ignoring the problem raised by the Quine-Duhem thesis) you will need to revise your theories.

When you revise your theories in light of experience, these theories may lead you to carry out different kinds of observations. You may develop different apparatus and use different techniques for reporting what you observe. Once again, however, nature may not fully cooperate, What you see may not correspond to what you expect even though the way you see is influenced by your theory. Nature may contravene your expectations, forcing yet another revision in the theory. Thus, even if observation is theory-laden, it is possible to discover that your theory is wrong. Theory-ladenness does not render science totally subjective, as the critics of theory-ladenness seem to fear. Objectivity remains because it is still possible (and happens frequently) that we make observations that contradict our theoretical predictions and thus show us that our theories are wrong.

This is not to say that the theory-ladenness of observation does not produce real problems. These problems arise when those who use different theoretical perspectives try to produce evidence to show the other that their way of interpreting the world is correct. Such people may well disagree about what it is they see and both may take what they see to support their theory. In such a case, without a neutral set of observations, it may not be possible to adjudicate between competing theoretical disputes simply on observational grounds. This, however, may not be too severe a problem if there are other ways of settling such disputes. A potentially more serious difficulty is that one may even lack ways of comparing competing theoretical frameworks. For the Positivists, the observational framework provides a basis for explicating the meaning of scientific claims, and if there is no common observational framework, such comparisons between competing frameworks may not be possible. This problem is commonly referred to as the *incommensurability of theories*. The incommensurability problem has been developed by those philosophers who have focused on the historical development of science, to which I turn in the next chapter. There I show how these philosophers view the theory-ladenness of observation as just a piece of a story that ultimately requires us to adopt a quite different view of science.

OVERVIEW OF THE CRITICISMS OF POSITIVISM

In this chapter I discussed four basic objections that have been raised against

Positivism. The Positivists sometimes adopted the metaphor of a science as an architectural structure built up from its foundations. I have presented these criticisms in an order that takes the structure apart from the top. The second and third sections reviewed objections that attack the superstructure—the hypothetico-deductive method of theory development, the deductive-nomological model of explanation, and the axiomatic view of theories. The objections discussed in the last two sections focused on the foundational assumptions of the verificationist theory of meaning and the idea of a neutral foundation in observation. (For further discussion of these and other criticisms of Logical Positivism, see Suppe, 1977, and Brown, 1979.) Altogether, the criticisms sketched in this chapter have proven so destructive that only a few contemporary philosophers still affirm allegiance to the Positivists' position in its original form.[7] Yet, the Positivists' picture of science remains the most comprehensive we have. The failure of Logical Positivism, if indeed it has failed, is, therefore, all the more noble and it leaves a legacy. Most philosophers of science find it impossible to dispense totally with the Positivist heritage even while recognizing various shortcomings. In the following chapter, however, I describe an alternative approach to philosophy of science that seems to many philosophers to provide the beginnings of an alternative.

[7] One important element of the positivists' tradition, a focus on a formal analysis of science, however, is affirmed by a number of contemporary philosophers. Prominent examples include the development of the semantic view of theories by Van Fraassen (1980) and defended by Giere (1979), Lloyd (1984), and others; a variety of probabilistic interpretations of science, many deriving from Bayes Theorem (Levi, 1967); and Glymour's (1980) bootstrapping analysis.