4

Post-Positivist Philosophy of Science

THE EMERGENCE OF HISTORICALLY GROUNDED PHILOSOPHY OF SCIENCE

The most influential development in philosophy of science in the wake of the demise of Logical Positivism has been the introduction of an historical perspective into Philosophical thinking about science. This has resulted from a greater concern among philosophers of science to describe the actual character of scientific investigation and the ways in which scientists choose which theories to accept or at least to pursue in further research. Within the framework of the Positivists and most of their critics (Popper being a notable exception), it was generally accepted that the primary factor that should govern decisions about the acceptability of a theory was the degree to which it corresponded to the evidence.[1] Post-Positivist philosophers argue,

[1] Even the Positivists recognized that this was not a sufficient criterion. In principle, at least, theories are always *underdetermined* by the evidence. There are always a variety of possible theories that will accommodate any set of evidence. This, however, is not a terribly serious problem for practicing scientists because it is often very difficult to find even one theory that perfectly fits the evidence available. But even then there is an important underdetermination in science. When theories fail, there are a variety of ways to revise the theoretical framework to accommodate the apparently disconfirming evidence. The Positivists and their descendants have proposed a number of criteria for guiding such decisions. Quine and Ullian (1970), for example, propose five such criteria—conservatism, modesty, simplicity, generality, and refutability. According to these criteria, we should prefer new theories that require the least change in currently accepted views, that are least risky in their claims, that are most simple, but yet that have broad generality and for which it is conceivable to acquire falsifying evidence. Although recognizing some tension between their criteria, they argue for each in terms of how following them is likely to enhance our ability to develop true theories.

however, that these are not the primary factors governing scientists' decisions and that, by focusing on them, earlier philosophers have developed accounts of science that do not accurately describe real science. If philosophy of science is to be of any value, it must attend to those factors that govern real scientific practice and that have enabled science to become a valuable knowledge producing enterprise.

Having criticized the Positivists for failing to describe real science, post-Positivist philosophers of science propose to develop their analyses not from general logical considerations but on the basis of careful examinations of the actual processes of science, particularly as revealed through its history. This, however, gives rise to a crucial issue. By showing how theories could be logically justified or corroborated on the basis of objective evidence, the Positivists hoped to show why science was able to produce true accounts of nature and why it should be valued as an objective source of knowledge. Further, the Positivists offered their account as a normative guide for the conduct of good science. By focusing on how scientists actually make decisions and countenancing factors other than criteria that could improve the likelihood of giving true accounts, advocates of an historical approach seem to be foregoing any possibility of producing a normative account of science and are left with attempting to provide only a descriptive account.

Historically based philosophy of science also encounters objections from another direction. A number of contemporary historians and sociologists of science claim that their tools of analysis enable them to provide a more adequate descriptive account of how science actually operates than philosophers. In particular, they charge that philosophers are preoccupied with the reasons scientists give and the logic of their arguments, but that these are not the true determinants of scientific investigations. Rather, these historians and sociologists maintain that a variety of institutional and social factors are what really govern the conduct of science. This raises the issue of whether there is a distinctive task for philosophy of science once it surrenders its claim to being able to give, on the basis of logical analysis, a normative account of what science should be.

Many philosophers have tried to show that there is still a point to analyzing the reasoning of science and to addressing such questions as how science makes progress. In doing so, they recognize that social and historical factors do influence the actual process of science, but they maintain that the reasoning of scientists is also an important determinant. Some post-Positivist philosophers even hold that their endeavor may give rise to normative judgments about the best ways to pursue scientific inquiry. They maintain, however, that judgments cannot be grounded on a priori principles but must be pragmatic judgments based on what has been successful science.

The primary inspiration for the development of post-Positivist philosophy of science, as well as for many of the recent endeavors in the history and sociology of science, was the publication of Thomas Kuhn's (1962/1970) *Structure of Scientific Revolutions.* This book, although not establishing a new general theory of science, offered a radically new framework for thinking about the character of science. Hence, the first part of this chapter is devoted to an exposition of Kuhn's account of science. I then survey more recent developments within this general approach.

KUHN'S CHALLENGE: NORMAL AND REVOLUTIONARY SCIENCE

Kuhn challenged the assumption of many previous philosophers of science that science offered a steadily accumulating body of knowledge. In contrast, he claimed that scientific disciplines go through distinct stages and that the character of research in the discipline varies between stages. Kuhn differentiates five stages: (a) immature science, (b) normal mature science, (c) crisis science, (d) revolutionary science, (e) resolution, which is followed by a return to normal science. Thus, a closed loop results involving stages b, c, d, and e (see Fig. 4.1).

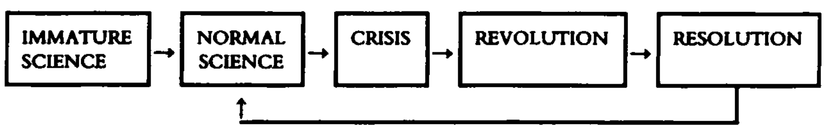

FIG. 4.1. Five stages in the history of scientific disciplines according to Kuhn.

Following the lead of the title of Kuhn's book, much of the discussion has focused on scientific revolutions. However, some of the most revolutionary claims Kuhn makes concern not revolutionary science but what he labels "normal science." Thus, I begin with his conception of normal science.

Normal science requires the establishment of what Kuhn has called variously a "paradigm" or "theoretical matrix." In his original book, Kuhn was not precise in his characterization of paradigms, which has resulted in numerous publications in which Kuhn and others (see Kuhn, 1970b) have tried to clarify that notion. The intuitive idea underlying Kuhn's term *paradigm*, however, can be explicated readily. A paradigm provides a framework for characterizing phenomena that a particular discipline takes as its subject matter. This might involve a basic model or a general theory. Kuhn has in mind items like Copernicus' model of the planets revolving around the sun or the theory of physical bodies attracting one another in accord with Newton's laws. In cognitive science, the idea of the mind as an information-processing system could constitute the paradigm.

A paradigm is not simply a model or a theory, but also includes instructions as to how such a theory or model is to be developed and applied in further research. These instructions may take the form of examples showing how derivations are to be made from the theories in the manner espoused by the Positivists, but they need not be so restricted. In describing these instructions, Kuhn appeals to science textbooks. The general theories of a discipline are commonly taught through a set of examples that show how to apply the theory to phenomena in its domain. For example, a textbook may illustrate Newton's principles by showing how they can be used to determine the gravitational attraction the earth exerts on a body on its surface. The examples taught often are the first applications that were developed in the discipline using the theory and they become the models for the new students. In cognitive psychology, the example could be the experiments used to establish the distinction between short-term and long-term memory and the use of chunking to explain how experts appear to exceed the normal limits on information in short-term memory. Kuhn refers to these standard applications of the basic framework as *exemplars*. (He also speaks of these exemplars themselves as *paradigms*, using the term now in a narrower sense.) One important function of an exemplar is to teach students by providing a model for them to imitate. Students begin by literally imitating the exemplar itself, for example by solving the same equation for different values. Later they learn to modify and extend the exemplar to solve new problems by analogy to the original problem solution.

Even though I have used the term *theory* to characterize one aspect of Kuhnian paradigms, Kuhn's understanding of theories is quite different from that of the Positivists or their critics considered so far. A theory need not be something that can be rendered into an axiomatic structure. It is not a set of postulates from which observations are deduced. Rather, it often consists in a rather fuzzy schema of how nature behaves—a schema that is imprecise and requires a great deal of further clarification (as does the idea that cognition involves processing information). Moreover, Kuhn offers a quite different account of how theories figure in the research of normal science than did the Positivists or Popper. A scientist's goal, according to Kuhn, is neither to confirm nor to falsify theories. Rather, it is to fit the theory to nature. The initial theory is incomplete. It offers a general account of how processes in nature work, but this general account needs to be embellished and filled in. Normal science must figure out what must be added to the general account so as to apply it to specific situations. This may involve figuring out ways of performing measurements in specific contexts or it may involve supplying additional assumptions that are needed to cover those contexts. Practitioners of normal science are continuing to do what they learned to do as students—imitate the exemplars they learned in school in new contexts. Through such activity they extend the applicability of the general theory or paradigm.

The process of further developing the paradigm by applying it to new cases is not necessarily an easy task. The historical record, Kuhn claims, shows that theories seldom fit nature precisely. Even during the period in which a particular theory or paradigm (e.g., Newtonian mechanics) is used and generally accepted, there are predictions that are not borne out in observations. During normal science these differences between theoretical predictions and empirical observations are not taken as falsifying the theory, but rather as creating further problems scientists must solve. Sometimes when the theory fails to fit nature the researcher will tinker with the theory, but that is not the only recourse. Kuhn contends that generally when experiments do not come out in accord with a theory, the problem is attributed to the experiments and not to the theory. The task of the experimenter, according to Kuhn, is to get experiments to produce results in accord with the accepted theory. In this endeavor, the theory is assumed to be largely correct and the task is to make it work. The success of a scientist during a period of normal science is judged in terms of whether he or she is able to resolve problems or puzzles that result from experimental failures and to demonstrate new applications of the already accepted theory. Thus, we see a key difference between Kuhn's conception of science and that of the Positivists. For the Positivists the task of science was to evaluate theories whereas for Kuhn it is to work out the details and develop experimental applications of theories.

Because of the nature of the task during normal science, the activities are clearly delineated and success is easily evaluated. The result is that there is general agreement and harmony within the scientific community. Such agreement and harmony is a sign for Kuhn that the community is engaged in normal science and that a paradigm is in place. Prior to the acceptance of such a paradigm, such agreement will be rare. Activities will not be guided by a generally accepted paradigm. Rather, there will be competing schools, each with its own view of how the discipline should develop and each trying to establish hegemony over the discipline for its approach. Such conflict between schools characterizes Kuhn's first period, the period of *immature science*. According to Kuhn, many of the social sciences are still in this kind of period, waiting for the first paradigm to be established. Rather than setting out to make the paradigm work, researchers spend most of their time battling over what general approach should be adopted. It is only once such a paradigm is established that the discipline will start to make progress because progress involves fitting an already accepted paradigm to nature. (The battles between various schools of psychoanalysis give a clear model of what Kuhn has in mind. Experimental psychology was presumably in this situation at the end of the 19th century, when associationists, functionalists, and structuralists all propounded their vision of what the science of psychology should be like.)

It is not necessary for researchers to prove that their paradigm is correct in order to reach the stage of normal science. Rather, a paradigm is adopted because it seems to offer potential for explaining a particular domain of phenomena and suggests a research program that various scientists can work on together. Probably the first clear case of normal science in psychology emerged with the rise of behaviorism. Although it began as simply another school, it offered a variety of techniques (e.g., classical conditioning, later operant conditioning) and an agenda (to show how much of behavior could be accounted for by behavioral laws) that generated a sustained attempt to force nature to conform to the mold of the paradigm.

Although generally such paradigms exhibit significant successes at the outset, Kuhn maintains that most paradigms eventually reach a juncture where unsolved problems begin to build up and success in solving problems slows significantly. Instead of new solved problems, failures, or what Kuhn calls "anomalies," begin to occur. This produces the third of Kuhn's stages, the stage of *crisis*. In response to the reduced rate of progress and the amassing of unsolved problems, the rules of research that are generated during the period in which the paradigm was successful are relaxed and researchers become more imaginative in considering ways in which the paradigm itself may have to be modified. Thus, to continue with the example of behaviorism, after its initial successes there emerged a variety of problems in which researchers were interested but which seemed beyond the reach of behaviorist explanations. These included language, rational planning, and problem solving. Skinner's (1957) *Verbal Behavior*, although intended to extend the behaviorist program into the domain of language, convinced many others of the poverty of the behaviorist approach. Behaviorism was then in the crisis period.

Sometimes the imagination shown during a scientific crisis will, Kuhn claims, result in the development of new, alternative paradigms that work with different fundamental principles and models and offer their own promise of creating a problem-solving tradition. If such new paradigms begin to develop, the discipline enters Kuhn's period of *revolutionary science*. The use of the political term *revolution* is clearly intended by Kuhn. A revolution is a period of active struggle between defenders of the old paradigm and the proponents of the new one. As in a political revolution, the rules that govern during normal times are no longer applicable (for these rules depended on the paradigm that is now under assault) and so part of what is at issue is what rules will be used to decide between the opposing paradigms. It is here that Kuhn makes his most contentious claims. He contends that competing paradigms are incommensurable in that they cannot be compared and evaluated on rational grounds.

In arguing for incommensurability, Kuhn draws upon the claimed theory-ladenness of observation discussed in the previous chapter. In accord with Hanson, Kuhn denies that there is a neutral observation language and claims that practitioners of a paradigm learn to report their observations in a theoryladen (or paradigm-laden) manner. Because each paradigm will have its own way of reporting observations, advocates of competing paradigms will not characterize what they see in the world in the same way. At one juncture Kuhn speaks of scientists working from different paradigms as living in different worlds. Kuhn also concurs with Quine (see previous chapter) in repudiating the analytic–synthetic distinction and the idea that vocabulary of a language can be assigned meanings independent of the theories presented in the language. Hence, for Kuhn there is no neutral language in which one can compare paradigms. The result is that the competing parties in a scientific revolution must resort to extra-rational means to settle their dispute. Fundamentally, this involves the proponents of one paradigm convincing significant numbers of scientists to adopt their paradigm.

During the revolution, researchers return to a situation much like that found in immature science. Practitioners of competing programs bicker amongst themselves just like the proponents of different schools do in an immature science. Thus, in the late 1950s and through the 1960s there were active debates between behaviorists and those advocating the new cognitivist approach (for examples of the cognitivists' attacks, see Chomsky, 1959 and Miller, Galanter, & Pribram, 1960). The two groups waged bitter arguments about the legitimacy of positing mental states and using them in the course of explanations. Behaviorists and cognitivists defined the goal of psychology differently and offered different criteria as to what would constitute an adequate psychological explanation. Given the nature of these differences, they could not be resolved in the same manner as differences would be resolved in normal science, for there was no agreement on what counted as an adequate psychological explanation. That, in fact, was what was in dispute. Eventually cognitivists generally gained the ascendancy, largely by attracting new researchers to their approach and showing that a successful research program was possible. In some departments behaviorists remained prominent and carried on their research, whereas other behaviorists quietly adopted some elements of the cognitivist approach. Overall, however, cognitivism supplanted behaviorism.

What has happened accords with Kuhn's characterization of the typical outcome of a revolution—a new group of researchers has gained ascendancy (controls the awarding of degrees, access to journals, etc.). It has pushed its paradigm on the discipline and created a new period of normal science. This process is typical of what occurs in Kuhn's final stage of *resolution* during which one school succeeds in making its paradigm dominant. The resolution generates a new period of normal science and a repeat of the cycle. There is some evidence that the cycle is about to repeat again in cognitive science. Cognitive scientists who take the name "new connectionists" have pointed out shortcomings in the traditional cognitivists' program and the past few years have witnessed vocal and sometimes acrimonious debates between traditional cognitivists and new connectionists. The emergence of connectionist research programs suggest that the alternative paradigm is developing and may succeed in developing its own normal science.

Kuhn's claim that competing paradigms are incommensurable so that scientific revolutions can only be settled through extra rational means has been the focal point of much subsequent controversy (see Gutting, 1980; Scheffler, 1967; Shapere, 1966). The reason is obvious—Kuhn's account seems to undermine the validity of science as a rational enterprise. Instead of basing the acceptance of a paradigm or theory on good rational justification (i.e., on sound logical arguments), Kuhn's view is that the decision to accept a paradigm is a matter of taste or persuasion. Kuhn does offer one extra-paradigmatic criterion for evaluating paradigms—their progressiveness. He characterizes scientists as choosing a paradigm because of its potential to solve puzzles and extend its range of applicability. But many philosophers have found this to be inadequate. If using different paradigms results in such drastic differences as Kuhn portrays, the puzzles confronted by two paradigms will not be the same. The identity of problems will be paradigm bound. More drastically, we will not even be able to determine when two paradigms are competing in the same domain because they will offer radically different accounts of the domain for which they are a paradigm.[2] As we see, this is one of the issues that many post-Kuhnian philosophers of science have tried to address.

Through his account of normal science and revolutionary science Kuhn has transformed philosophical thinking about science, but has not succeeded in creating a new orthodoxy. He refocused philosophical thinking on the actual dynamics of scientific activities and away from the abstract logic of confirmation and falsification of scientific theories. Kuhn's accounts have opened up a new kind of criticism of philosophical theorizing—a criticism that charges that the philosophical theory in question does not accurately reflect the processes that govern actual scientific investigation. Such a charge has been leveled against Kuhn's own theory by philosophers who have followed in his footsteps and tried to study the actual dynamics of scientific investigations. Although I cannot hope to discuss all the competing views put forward since Kuhn, I highlight some of the points where these views have differed from Kuhn's so as to reveal the character of contemporary thinking about the nature of the scientific enterprise.

[2] In the case of cognitive psychology and behaviorism, there is a case to be made that the two enterprises are not really concerned with the same thing and should not be viewed as competitors. Behaviorism is concerned to identify the external factors governing behavior while cognitivism is concerned with the mediating structures inside the mind. One could readily grant that both are important and look for ways to incorporate insights from both approaches (see Bechtel, in press b; Schnaitter, 1987).

FEYERABEND'S ATTACK ON METHOD

Although most of Kuhn's critics have objected that he sacrificed too much of the rational or logical character of science in his account of how paradigms succeed each other in revolutions, at least one critic, Paul Feyerabend, has pushed to an even more radical position. In his earlier writings (Feyerabend, 1962, 1963a, 1965) he argued against two features of Positivistic philosophy of science that he calls the *consistency condition* and the *condition of meaning invariance*. The consistency condition holds that new theories should be consistent with currently held theories. The condition of meaning invariance holds that the meanings of terms should be held constant across theories (e.g., by something like the verificationist theory of meaning). Feyerabend's objections to these two conditions rest on examination of actual scientific practice and on demonstrations that cases of major advance in science have not adhered to them. For example, Newton's laws often are portrayed as subsuming Galileo's law of free fall and Kepler's law of planetary motion, but Feyerabend argues that Newton's law is actually inconsistent with both of them (e.g., Galileo posits a constant rate of acceleration in free fall, whereas Newton's laws predict a decreasing acceleration). To make his argument against meaning invariance, Feyerabend argues first that the meaning of terms depends on the theoretical context in which they are used, and then shows that critical terms (e.g., mass) change their meaning from one theory (Newton's) to another (Einstein's). Both the consistency principle and the condition of meaning invariance impose, Feyerabend contends, a destructive conservatism on science that would paralyze it.

Feyerabend finds an unhealthy conservatism built into any attempt to specify a methodology for science. In particular, he rejects the idea that researchers should continue to accept a theory until it has been falsified. He contends that we need to consider alternative theories in order to discover the data that might falsify a theory. Each theory we pursue will bring to light new data, and it is these data that may serve to falsify our preceding theories. Again, Feyerabend argues through examples, pointing to the example of Brownian motion, which would not have been discovered simply by those trying to test the second law of phenomenological thermodynamics. It was only discovered by those investigating the kinetic theory of gases, which is inconsistent with the phenomenological second law (Feyerabend, 1965). Feyerabend thus rejects Kuhn's view of normal science, claiming both that such periods of research done totally within the framework of a single paradigm are not common in science and that they would be destructive of science. Science, Feyerabend contends, must depend on maintaining a plurality of pursuits.

Having rejected both Positivist and Kuhnian accounts of scientific methodology, Feyerabend (1970, 1975), advanced a principle of *methodological anarchism* that denies that there are any sound methodological principles that should be imposed on science. He claims that any principle we might propose has been violated by good scientists and had to be violated for science to progress. He (1970) concludes that "there is only one principle that can be defended under all circumstances, and in all stages of human development. It is the principle: *anything goes*" (p. 22). Although in many other respects, Feyerabend aligns himself with Popper, in this he challenges Popper as well, arguing that we must even pursue theories that have been amply falsified. Pursuing even these falsified theories may reveal new information, which may serve to invalidate the supposedly falsifying data. In particular, he maintains that new theories will, almost inevitably, be falsified by data produced by older theories but may themselves reveal new data which favors them. In order to break the hegemony of old theories, and bring new data to light that may in turn defeat those theories and support the new idea, Feyerabend calls upon scientists to proceed "counterinductively." This involves producing and defending theories that seem already to have been effectively refuted by current evidence. His contention is that it has been the rule violators who have made the most progress.

To defend counterinduction, Feyerabend points to Galileo, whom he portrays as succeeding only by sabotaging the enterprise of the dominant Aristotelian physics through effective propaganda and circular arguments. When Galileo first offered his new theories of motion, the Aristotelian establishment could readily offer counterevidence. For example, against the theory that the earth was in motion it was seemingly sufficient to drop an object and note that it fell directly to the spot below it, not to a spot behind, as it would if the earth had moved during the fall. To undercut this evidence, Galileo had to argue circularly. The circular arguments would employ unorthodox research methods to establish unorthodox theoretical claims and then use those results to justify the use of the methods themselves. For example, Galileo sought to provide evidence for the new Copernican astronomy according to which the earth orbited the sun. One of the keys to his claim was the assertion that celestial bodies like the moon were in fact physical objects like the earth. To establish this he invoked the telescope, through which one could detect the mountainous lunar landscape. However, the Aristotelians dismissed the use of the telescope on the grounds that when it was used to look into the heavens it would produce distortions because of the quite different etherial medium through which light was passing. So Galileo had to invoke a new optical theory. Only by packaging his alternative view as a whole and then insisting on answering all objections on grounds internal to his new conception, was Galileo able to establish his new physics.

On the basis of such historical analyses, Feyerabend contends that attempts to prescribe particular methodologies for scientific research are used primarily to protect vested interests and to prevent new approaches from developing.

He warns against attempting to devise rational criteria for deciding between theories, including criteria that appeal to how progressive a theory has been and seems likely to be in the future. New ideas that have not had a chance to devise their own methods of support need to be protected against premature dismissal. He even maintains that we should hold on to long-tried ideas that have failed, such as alchemy, because we can never tell when these old ideas might produce new insight and show us errors in our current theories. Feyerabend carries this point to what many would consider an extreme. He maintains that any account, no matter how absurd it seems to those who think in terms of contemporary science (he cites creationism, astrology, and parapsychology), might prove instructive.

Generally, Feyerabend's views have been deemed too extreme to be worth serious treatment, and so he has lost credibility. Yet, he has provided a useful service in showing how conservative established science can be, and how generally recognized progress in science sometimes requires going outside the established order. Most philosophers, however, are interested in what rational strategies are available to maximize scientists' endeavors to improve their science. Hence, unlike Feyerabend, most post-Kuhnian philosophers of science have tried to show how rational considerations can provide useful guidance.

LAKATOSIAN RESEARCH PROGRAMMES

In trying to explicate how science is a rational enterprise, a number of philosophers, who generally accept the importance of adopting an historical perspective, have sought to bring logical analysis back into the philosophy of science and to rekindle some of the same interests that inspired the Logical Positivists—that of evaluating and justifying the scientific enterprise. This is clearly seen in the work of Lakatos (1970, 1978) who analyzes cases from the history of science, but freely adopts the Positivists' strategy of reconstructing these episodes so as to show how they could have progressed by adhering to rational canons. He argues that for philosophical purposes we can overlook some of the details about how a science actually progressed and develop an alternative account about how it could have progressed in a rational manner.[3] Unlike the Positivists, however, Lakatos is interested in both discovery and justification, and, in particular, in the ways in which a science may develop over time. Lakatos' account of the character of science can thus be seen as an attempt to recast Kuhn's insights about the nature of actual science as an historical process into a perspective that can explicate its rational import.

[3] This approach of Lakatos' has aroused the wrath of a number of historians and sociologists of science, and has been disavowed by many subsequent philosophers of science.

Lakatos begins by taking issue with Kuhn's claim that we can differentiate distinct stages of normal and revolutionary science. Rather, Lakatos contends that science is seldom dominated by just one paradigm, as Kuhn claims in his account of normal science, but rather that competition between paradigms generally co-occurs with processes of development within a paradigm. Lakatos also takes issue with Kuhn's conception of normal science as filling in and further applying a single paradigm. He contends that research often consists in developing a succession of theories, in which new theories replace older ones while preserving important features of the older theories. To allow for this idea of a succession of theories, Lakatos replaces Kuhn's term *paradigm* with the term *research programme*. The common thread linking different theories into a common research programme is a "hard core" of basic assumptions shared by all investigators. This core is surrounded by a "protective belt" of auxiliary assumptions. The hard core, which may consist of assumptions such as "No action at a distance," remains intact as long as the research programme continues, but researchers can change the auxiliary assumptions in the protective belt to accommodate evidence that either has accumulated or is developed in the course of research.

For Lakatos, the ultimate measure of a research program is whether it is progressive. Progress consists in developing new theories in the protective belt around the core. Lakatos distinguishes two kinds of progress—theoretical and empirical progress. *Theoretical progress* consists in extending the empirical scope of a theory by applying it to new empirical domains. *Empirical progress* consists in corroborating empirically the new claims that are made in the course of theoretical progress. (In these notions of theoretical and empirical progress we can see the influence of Popperian principles on Lakatos' thinking.) To be counted as progressive, the research tradition must be making both theoretical and empirical progress, although Lakatos allows that the empirical progress may be more intermittent than theoretical progress. If it is not progressing, Lakatos speaks of the research programme as degenerating. Unlike Popper, however, Lakatos does not see degeneration, even when the new theories in the research tradition seem to be falsified, as providing a reason to give up the research programme. A verdict that a research programme is degenerating is not final and should not, Lakatos maintains, lead to the total rejection of the degenerating research programme. A programme may be very progressive for a period, degenerate for a while, and return as a progressive programme again. Within cognitive science, connectionism might be such a research tradition which, after languishing, has once again become progressive. In the guise of perceptrons (Rosenblatt, 1962) and Hebb's (1949) account of reverberation of activation in neural networks, the connectionist model initially seemed to offer great hope. Partly as a result of criticisms (e.g., Minsky & Papert's, 1969, criticism of the perceptron model) and partly due to lack of major successes, the connectionist model languished until it was rekindled in the current connectionist literature. Now once again the connectionist program seems to be advancing due to enhancements in the types of models used, and connectionists are developing models that give realistic performance on numerous cognitive tasks.

For Lakatos, the way a research programme makes progress is not totally random but is guided by heuristics. He distinguishes negative and positive heuristics. The negative heuristic of a programme is simply the injunction not to modify the core principles of the program. More significant is the positive heuristic, which "consists of a partially articulated set of suggestions or hints on how to change, develop the 'refutable variants' of the research programme, [and] how to modify, sophisticate, the 'refutable protective belt' " (Lakatos, 1970, p. 135). These heuristics play a critical role in the development of a research programme, for they are what determine whether a research programme will be progressive. In successful research programmes, these heuristics will enjoy a period of success as they guide researchers to the development of viable new theories, but almost inevitably they will become exhausted. Then the programme slips into degeneracy, although it may revive when someone develops new heuristics to restart the programme or finds a way to graft the old programme onto a new endeavor that itself is progressing. Such a graft may not be totally consistent, as the assumptions of the new and old theories may disagree. If, however, the graft suggests ways of solving problems within the old research programme, the inconsistency will not be critical.

Lakatos' conceptions of science brings together Popperian considerations about what makes scientific investigation rational with the Kuhnian perspective of looking at larger scale units as the effective units in scientific progress. One of Lakatos' important insights, which constitutes a departure from Popper in the direction of Feyerabend, is to recognize that in its early days a new research programme will not yet have achieved as much success as older programmes. Moreover, given the initial, simplified version of the early theories offered in the programme, there will be many phenomena that seem to falsify them, but that subsequent theories in the program will nonetheless be able to handle if they are allowed to develop. Thus, Lakatos contends that new research programmes, which show potential, need protection at the outset until they have an opportunity to develop. From Lakatos' perspective, both early cognitivism and contemporary connectionism can claim a grace period before they should be expected to compete as equals with the preceding research programmes.

By rejecting Kuhn's notion of normal science as involving the dominance of just one research program and proposing that competing research programmes are always being pursued, Lakatos introduces a new kind of theory evaluation—a relative evaluation as to which research programme is best. For Lakatos, the measure is progress, measured in terms of extending the

theoretical and empirical scope of the theories in the program. One of the more interesting developments in the wake of Lakatos' work is an attempt to develop a more complete account of what such progress consists in and how it can be measured.

LAUDAN'S RESEARCH TRADITIONS

Laudan (1977) proposes an account of scientific activities designed to capture some of the strengths of Kuhn's and Lakatos' account and yet overcome shortcomings in them. Laudan agrees with Kuhn and Lakatos that the main activity of scientists is problem solving. But he offers a more complete account of the kinds of problems scientists encounter and also a finer grained analysis of how scientists may evaluate the seriousness of problems and the importance of problem solutions. In developing his account of science as problem solving, Laudan invokes the idea of a large-scale unit in science that he calls a "research tradition." Like Lakatos' research programmes, research traditions for Laudan consist of a sequence of theories, but they lack a common core that is immune to revision. What holds a research tradition together are simply common ontological assumptions about the nature of the world and methodological principles about how to revise theories and develop new theories.

Laudan distinguishes two kinds of problems that can confront a research tradition—empirical inadequacies of the current theories and conceptual problems with the theories comprising the tradition. Laudan's treatment of empirical problems is generally consistent with that of Kuhn and Lakatos—scientists face empirical problems when expectations based on the theories within their research tradition fail. It is by introducing conceptual problems as important kinds of problems driving scientific research that Laudan claims to be making a new contribution. Conceptual problems do not result from failure of empirical fit, for they could arise even if the theories were totally adequate empirically. The simplest kind of conceptual problem would be an inconsistency within a theory or an inconsistency between two theories that are conjoined in the research tradition. But Laudan claims that there are other sorts of theoretical problems as well: inconsistent ontological assumptions; conflicts between expressed ontological views and theories; and conflicts between the claims of theories and the broader world views of society, including religious or political views. Laudan has in mind such conflicts as those between the assumption that light is a unitary phenomena and the light-wave duality thesis, and those between evolutionary theory and religious fundamentalism. For cognitive science, such a problem might be the tension created by assuming that the mind is an information processing machine that works by deterministic principles and moral principles that seem to require free will.

Laudan not only broadens our conception of what counts as a problem that a research tradition must address, but also tries to develop a finer grained analysis of what problems are important to solve. Although any empirical or theoretical problem would, in principle, be worth solving, many are set aside as less important. Empirical problems, he proposes, are evaluated in terms of how important a challenge they seem offer to the research tradition and whether other competing research traditions have been able to solve them. Consider the case of an anomalous empirical finding that neither results from explicit tests of major theoretical assumptions nor has been explained by any other research tradition. This does not pose nearly as serious a problem as a research finding that directly confutes a major tenet of the research tradition and has already by accounted for by theories in a competing tradition. Conceptual problems can range from a serious problem as when there is a clear logical inconsistency between two accepted propositions, to a moderate problem when it seems implausible that two accepted propositions case are actually both true, to a relatively minor problem when two logically consistent propositions fail to support each other in a way that might be desired.

The task of scientists is to solve both the empirical and conceptual problems they encounter. In this effort it becomes important for them to evaluate competing research traditions. The capacity for such evaluation is what makes science rational. Like Lakatos, Laudan argues that the basis for evaluation is the progressiveness of the tradition. But just as Laudan broadens the conception of what counts as a problem, he also proposes that there are different standards one may use to measure progress in different contexts. One context is where we have to take practical action, for example, to use the results of scientific investigations in treating a disease. Here it is the overall progress that has been made by the research tradition that is relevant. However, when deciding what research tradition is likely to bear the most fruits in the future (e.g., in deciding which one to fund) we will be concerned with the promise of the tradition to solve future problems. We may use as a guide how rapidly research is progressing in the tradition, and choose a tradition which has not solved as many problems as another but is making rapid progress over one that was previously very successful but seems now to be making little progress. Thus, Laudan maintains that we might accept one research tradition to guide our actions while pursuing another in our research.

There is a distinctive aspect of Laudan's account of science that I have not yet touched on. Most of the philosophers who have adopted a historical perspective on science and who have concerned themselves with the progressiveness of science have assumed that science is progressing toward a true conception of nature. Laudan, however, although employing the idea of progress, does not think science is getting any closer to the truth. He argues for this partly by observing how frequently scientists have repudiated previous theories and replaced them with others that are radically different. He denies

that there is any metric upon which we can evaluate later theories as being closer to the truth than earlier ones or judge ourselves to be closer to the truth than our predecessors. Laudan maintains that the idea of truth as a goal for scientific investigation is peripheral. Progress is sufficient.

Laudan's account of scientific change, like the others, has attracted critics. On the one hand, insofar as Laudan permits such things as inconsistencies between a particular scientific theory and a religious tradition to matter in evaluating a science, he seems to be allowing nonscientific factors to enter into the evaluation of science. Although it is unproblematic that such external considerations do affect the direction of science, many question whether they should have such a role in a rationally developed science. Laudan however defends using these external considerations in evaluating a research program because he views them as also a part of our cognitive apparatus. He in fact maintains that even in nonscientific domains such as politics and theology the same conception of rationality measured in terms of problem solving applies so that these are equally a part of our rational cognitive endeavor. On the other hand, Laudan has also been attacked for not giving sufficient weight to social factors. Laudan wants to restrict the role played by such social factors in developing science and to emphasize rational considerations. He thus resists the move by some sociologists of science to treat scientific developments as purely social developments and argues for maintaining a privileged role for philosophical analyses of the reasoning of scientists, a contention that many philosophers as well as nonphilosophers regard as problematic.[4] Despite these and other criticisms, Laudan's account at present is one of the most fully developed philosophical analyses of the character of scientific research. Although it does not offer anything as comprehensive or as precise as the positivists' analyses, his analysis of problems and the evaluation of problems offers a promising route to further development.

STUDIES OF SCIENTIFIC DISCOVERY

One consequence of philosophers' growing interest in the history of science has been a growing interest in scientific discovery. As explained in chapter

[4] Laudan actually takes an extremely strong position with regard to sociological analyses, arguing that we only need to invoke sociological analyses in cases where analyses in terms of rational decision making fail to explain the behavior of scientists (Laudan, 1981). In part, Laudan is reacting to the "strong program in the sociology of knowledge," which argues that all beliefs, rational as well as irrational, should be explained in the same manner and, because rational factors can never fully determine what should be believed (because of the underdetermination of theories, etc.), social factors are always relevant to explaining scientific practice (see Barnes, 1977; Bloor, 1976, 1981). Not all philosophers have found the strong program inimicable to philosophical interests (see Hesse, 1980). There are, moreover, other programs in sociology of science (for example, the laboratory studies of Latour & Woolgar, 1979, and Knoor, 1981, and studies of scientific institutions by Whitley, 1980, 1982) that can be extremely informative for those interested in developing philosophical analyses of scientific decision making.

1, the Logical Positivists made a sharp distinction between the contexts of discovery and of justification. Discovery was assumed to be a nonrational process and so philosophical attention was directed at the question of how theories could be justified, not how they were initially developed. Hanson (1958, 1960, 1967) was one of the first to urge philosophers to redirect attention to discovery. Hanson's own approach, however, constituted only a modest step in the direction of studying discovery. His proposal was to pursue what the 19th century American Pragmatist, Charles Peirce, called *abductive inference*. According to this procedure, you start with a surprising phenomenon, identify the kind of hypothesis (partly on the basis of past experience) that would explain that phenomenon, and then pursue development of that kind of hypothesis.

Hanson's call for studying discovery was answered only gradually. Part of the reason philosophers were reluctant to pursue the study of discovery was that it did not seem to lend itself to logical analysis. Discovery was not a deductive process, because deductive reasoning could only lead one from premises to conclusions that would have to be true when the premises were. Quite clearly there are no such rules that can guarantee that the hypotheses developed in scientific investigations will be correct. Discovery procedures are, at best, fallible. The alternative to deductive reasoning is generally taken to be induction. The most commonly discussed type of induction is enumerative induction in which one proceeds from examples to general statements. (For example, one might, mistakenly, infer from seeing a hundred white swans and no black ones that all swans are white.) But despite Bacon's (1620) claim that such induction could lead to basic theoretical principles, it is generally recognized that such induction cannot generate the kinds of theoretical principles that figure in scientific explanations. For example, it does not lead one to understand the causal processes behind phenomena.

One of the things that has brought about renewed interest in discovery is the recognition, partly motivated by work in empirical psychology, that human reasoning involves additional modes of reasoning than deductive logic and enumerative induction. People invoke strategies of reasoning to deal with problems that may work perfectly well in most contexts, but violate the norms of formal logic. (See papers in Kahneman, Slovic, & Tversky, 1982.) Because scientific reasoning is simply an extension of ordinary human reasoning, there is reason to think that such strategies figure also in science and that a detailed study of the history of science may permit us to identify some of these strategies. (For a useful collection of papers, see Tweney, Doherty, & Mynatt, 1981.)

Simon (Newell, Shaw, & Simon, 1962; Simon, 1980) popularized the idea that in solving complex problems we rely on heuristic principles that simplify the process through which we search for a solution. It is useful to contrast heuristics with algorithms. Although both often can be stated as explicit rules, and so can be implemented in a computer, Wimsatt (1980) identifies three distinctive characteristics of heuristics: (a) they simplify the problem and so are "cost-effective" ways of reaching solutions; (b) they do not guarantee that a solution will be reached, or that the solution arrived at will be correct; and (c) the errors that result will be systematic so that it is possible to devise situations in which any given heuristic will fail. If scientists do reason using heuristics in discovery, then it is reasonable for anyone trying to studying scientific discovery to try to identify the heuristics they use. The systematic errors that result from such heuristics provide a tool for identifying them (see Wimsatt, 1980, for such an attempt to identify heuristics through studying the reasoning of different model builders in population genetics). Moreover, through identifying these heuristics and through identifying the circumstances under which they may fail, philosophers can once again play a normative role of evaluating the practice of science (Bechtel, 1982).

Recently there has been considerable interest by both philosophers and those in artificial intelligence (AI) in using AI as a tool for studying scientific reasoning. Simon and his colleagues have developed a number of programs designed to discover patterns in numerical data (BACON) and certain types of qualitative laws (GLAUBER; Langley, Simon, Bradshaw, & Zytkow, 1987). Currently, Simon is collaborating with an historian of science, Frederic L. Holmes, to develop a program that captures the details of Hans Krebs' discovery of the citric acid cycle in biochemistry. Two philosophers have also made contributions to this endeavor of using AI to understand discovery. Thagard, in collaboration with other cognitive scientists (Holland, Holyoak, Nisbett, & Thagard, 1986), has developed a computer simulation to capture the process through which the wave theory of sound was discovered. Darden, together with Rada (Darden & Rada, in press), has developed a computer program that discovers part-whole relationships (such as the discovery that genes are parts of chromosomes). Darden (1987) and Thagard (in press) both advocate the use of AI reasoning processes as a strategic tool for future studies in philosophy of science. One difficulty confronting this approach is the tremendous diversity of reasoning patterns that are exhibited in cases of scientific discovery. Even if we can develop reasoning strategies that account for particular cases of discovery, there remains a major problem of determining which procedure is appropriate for a particular circumstance. Nonetheless, the introduction of computer simulations into studies of scientific reasoning has introduced a new rigor to the enterprise (because procedures must be explicitly stated) and provided a means of studying these procedures empirically.

Computer simulations are not the only means for studying discovery. A number of philosophers have engaged themselves in the task of trying to extract from cases of science basic principles governing the discovery process. Nickles (1978, 1980a), for example, beginning with Laudan's conception of scientific problems (see preceeding), has explored, through a variety of historical cases, how such problems are constrained by known or accepted information and how these constraints may figure in devising solutions to the problems. (For additional case studies of discovery processes, see the papers in Nickles, 1980b.) One thing that has emerged from these case studies is that the context of discovery is not uniform. Some have argued that there are a variety of stages in scientific discovery, including stages of generation and of pursuit, with different strategies appropriate to each (see Nickles, 1980c, for discussion).

Philosophers have just begun again to study scientific discovery and there is considerable disagreement over how to proceed. Many philosophers remain skeptical that anything of value will result from the endeavor. Others, however, maintain that if philosophy of science is to remain faithful to actual science, then it must develop accurate accounts of a major scientific activity, that of reasoning through problems to new solutions. This is one area where philosophy of science promises to develop in subsequent years and to engage in fruitful collaborations with other disciplines of cognitive science.

SUMMARY OF POST-POSITIVIST PHILOSOPHY OF SCIENCE

In this chapter I discussed some of the prominent viewpoints developed by philosophers of science since Kuhn's (1962/1970a) publication of *The Structure of Scientific Revolutions*. Although none of these approaches has yet obtained the status once enjoyed by Positivism, they have yielded useful insights into the character of science and have also made philosophers much more concerned with the actual character of scientific research than they were previously. As a result of these analyses, new issues have emerged as principal topics of discussion in the philosophy of science. Shapere (1984), for example, has opened up a new line of inquiry into what the units of scientific investigation are, and has himself argued that the way in which the domain of scientific inquiry is defined is a prominent factor in governing developments in science. Another issue, which is perhaps the most widely discussed current issue, is whether scientific theories must be treated as real descriptions of phenomena in nature, or, in the manner of Laudan, as simply vehicles for characterizing phenomena we experience (see Leplin, 1984; Churchland & Hooker, 1985, for introductions to this debate).

The historical analyses of science that I have focused on in this section have attracted attention from a number of practitioners of cognitive science who have been interested in understanding the development of this field of inquiry. I have already noted that early practitioners of modern cognitive science viewed themselves self-consciously as revolting from the domination of behaviorism and many subsequent writers have used Kuhn's conception of a scientific revolution to characterize the development of the cognitive orientation (Reese & Overton, 1970; Weimer & Palermo, 1973). The development of alternative analyses of scientific development by Lakatos and Laudan, however, has led some investigators to inquire whether the history of psychology and other cognitive sciences might better be characterized using one of these frameworks. Gholson and Barker (1985) argue that Lakatos' account is far more compatible with the history of psychology than Kuhn's. In conformity to Lakatos' views, they argue that the competing traditions of cognitivism and behavioristic learning theory have experienced a long history of competition. One did not simply supercede the other in a Kuhnian revolution. Further, there was interaction between the two traditions, not just competition. As a result, some features of a cognitive perspective were adopted by some researchers in the learning theory tradition while cognitivists themselves made use of some contributions of learning theorists in developing their accounts. Moreover, Gholson and Barker argue that, in conformity with Lakatos' analysis, these research programmes each had periods of progress followed by periods of stasis or degeneracy before becoming progressive again. If one construes late 19th century endeavors by Wundt, James, and others as ancestors of modern cognitivism, one might argue that cognitivism itself was in a period of degeneracy during much of the period when behavioristic analysis dominated, but has reemerged in recent decades as a progressive programme.

Although they see Lakatos as providing a more accurate account of the cognitive sciences than Kuhn, ultimately Gholson and Barker advocate Laudan's approach. A primary advantage they see is that Laudan's account allows for changing core commitments, a phenomenon they attempt to trace through the history of learning theory. They find the idea of a cluster of theories, those held contemporaneously and those held successively, as providing a better account of what actually occurred in the development of psychology. Finally, they argue that conceptual factors as well as empirical factors have shaped psychology, again favoring Laudan's analysis over Lakatos'. For example, in the arguments between behaviorists and cognitivists, much of the controversy revolved around such issues as whether the digital computer provided a useful model of the cognitive system. That issue has recently reemerged with the development of connectionism. One argument advanced for connectionism is that a network system provides a more biologically realistic model of cognition than a digital computer.

These post-Positivist accounts of science do seem to provide convenient tools for explicating developments in a variety of scientific research domains, including cognitive science, but one must be cautious not to view the ease of application as clear evidence of the correctness of one account. At present, there are a variety of competing philosophical accounts of how science works, each of which seem to apply to a number of cases. But to appraise the correctness of any one of these far more rigorous investigations are needed that test philosophical theories against actual history. This process has begun (see Laudan et al., 1986; Laudan, Donovan, & Laudan, in press), but it faces an interesting reflexive problem. In developing historically adequate philosophical analyses of science, philosophers of science have treated their enterprise as itself scientific. Now they face the problem of deciding which model of science to appeal to in adjudicating the battle between competing models of science. For this and other reasons, philosophy of science can still be viewed as a discipline in flux. The post-Positivistic analyses have raised new issues and introduced new ideas that seem potentially fruitful, but as yet there still is no widely accepted, clear understanding of the nature of science.

5

Theory Reduction as a Model for Relating Disciplines

INTRODUCTION: RELATING DISCIPLINES BY RELATING THEORIES

In the last three chapters I focused on the general character of scientific inquiry and scientific theories. I now turn to a more specific issue, the question of how disciplines in science relate or should relate to one another. Many scientists are interested in these questions, but especially those involved in what I term *cross-disciplinary research clusters* such as cognitive science. In these clusters the avowed objective is to integrate contributions of various disciplines to deal with a common problem. One of the last legacies of Logical Positivism has been a very influential model of how to unify disciplines, known as the *Theory Reduction Model*. This model is a natural outgrowth of the Positivist's deductive-nomological model of explanation (see chapter 2). Like other parts of their account, the Positivists' account of reduction is both clear and precise, which partly accounts for its continued influence even after many of the Positivists' doctrines have been rejected. Recently, however, there has been growing dissatisfaction with the theory reduction model and alternative accounts of how to unify science have emerged. In the next chapter, I turn to the points of dissatisfaction with reduction and the attempts to develop alternative perspectives on the relationship between disciplines. In this chapter, however, I focus exclusively on the theory reduction model.

Reduction, like many other terms, carries a special meaning for philosophers. Whereas many scientists employ the term *reduction* for any attempt to invoke