

Published by the Press Syndicate of the University of Cambridge
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© Cambridge University Press 1996

First published 1996

Library of Congress Cataloging-in-Publication Data
Kincaid, Harold, 1952-

Philosophical foundations of the social sciences : analyzing
controversies in social research / Harold Kincaid.
p. cm.

ISBN 0-521-48268-2. — ISBN 0-521-55891-3 (pbk.)

I. Social sciences—Research. 2. Social sciences—Philosophy.

I. Title.
H62.K515 1996
300'.1—dc20

95-13774
CIP

A catalog record for this book is available from the British Library

ISBN 0-521-48268-2 Hardback

0-521-55891-3 Paperback

Transferred to digital printing 2001

In memory of the generation that went before

Harold Wilson Kincaid
1920-1943

Earl Alexander Kincaid, Jr.
1926-1970

Katherine Kincaid Brooks
1919-1987

2

Challenges to scientific rationality

Arguing for naturalism would be pointless if the natural sciences were not a paradigm of rational investigation. However, many philosophers and social scientists now argue that social forces drive the natural sciences. According to Kuhn, incommensurable paradigms dominate science, making scientific change a social process rather than a rational one driven by decisive tests. According to the social constructivists like Latour (1987), appeals to "evidence" and other epistemic values are really just a cover for negotiation and network building. These irrationalist views have enormous implications if true. Not only would naturalism be a non-starter, the social sciences would also have to be drastically reconceived. The drive for better data and more careful tests would be misguided. Much that social scientists do and struggle for would be without foundation. Moreover, social scientists would lose any claim to expertise on policy issues if their results are mere social constructs or one of many incommensurable paradigms. Social science would be just one more human conversation but with pretensions to be something it could not be.

This chapter tries to show that social scientists need not adopt these irrationalist doctrines. In what follows I try to do three basic things: answer the many irrationalist critics of natural science, identify what I take to be the basic core practices defining good science, and sketch the recent developments in philosophy of science that later chapters presuppose. The first two tasks are essential if we are to take naturalism seriously; the third, informal task will emerge in the process. (Readers familiar with recent philosophy of science or unconcerned with challenges to scientific rationality can skip all but Section 2.6 without detriment to later arguments.)

The argument proceeds as follows. Section 2.1 sets the stage by discussing Quine's (1953) path-breaking attack on positivism. Quine's criticisms made possible much of the currently trendy irrationalism, but his

views also ground our best current answer to relativism. Section 2.2 clarifies issues. Numerous different theses are at work in debates over scientific rationality; discussing them undistinguished has caused considerable confusion. Section 2.3 takes up Kuhn's (1970) influential antirationalist views. He has convinced many social scientists that data cannot decide between theories, for data are theory laden and theories embody incommensurable paradigms. Section 2.3 explains why these claims are untenable.

Section 2.4 takes up more recent sociological approaches: social constructivism and Rorty's "post-modernist" pragmatism.¹ Social constructivists make the obvious, albeit long-ignored point that scientific institutions are social institutions. Accordingly, scientific beliefs must be formed by a social process. So, it is social negotiation, not evidence and the structure of reality, that explains scientific practices. Post-modernism presents an equally serious challenge. Rorty's claim that science is just another form of non-coercive persuasion has caught on among social scientists (for example, McCloskey 1985; Weintraub 1991). Section 2.4 examines these doctrines. We should, I argue, take the social nature of science seriously. But we can do so without committing ourselves to the dismal conclusion that all science is politics by other means.

Another threat to scientific rationality comes from the claim that science is inherently value laden. If normative elements constantly intrude in scientific practice, then we can hardly claim that science has any special lock on objective, impartial pursuit of the truth. Though there are important reminders here, I argue in 2.5 they are no serious threat to scientific rationality.

A general picture of good science will slowly emerge as we answer the irrationalists. Section 2.6 makes that picture more explicit. It also explains how we can say something useful about scientific adequacy while recognizing the contextual, contingent, and empirical nature of claims about scientific method. This last section thus more precisely formulates the initial naturalist theses identified in the last chapter. The general upshot is then twofold: first, we can defend the rationality of science without falling into simplistic positivist assumptions, and thus, second, naturalism in the social sciences is a live and fundamental issue.

2.1 Quine and the demise of positivism

Current philosophy of science owes its intellectual origins in large part to Quine. Quine's influence comes from the role he played in the

¹ Seminal works include Bloor (1976), Barnes (1982), Latour and Woolgar (1979), and Latour (1987). Rorty (1979, 1982b) are important works for the rhetoric of science.

downfall of logical positivism. The common thread to nearly all current philosophy of science is what it is not: positivist. And it was Quine's criticisms that led to the current post-positivist outlook. In this section I explain those criticisms and the basic framework for philosophy of science set by Quine.

Positivism, like other bygone intellectual movements, is probably inaccurately treated by its current detractors. Some or many positivists may not have held all the views currently falling under that rubric, and the positivists also changed their views over time (see Michael Friedman 1992). Nonetheless, the "positivist" program loosely construed has been an important factor in the way philosophers and social scientists have understood their work. The positivists believed that there are two kinds of truths: empirical, factual truths of observations and truths based on the meaning of words. Good science was the paradigm of the former; mathematics and logic, the paradigm of the latter. Like Hume before them, the positivists used this division to draw some striking conclusions. Traditional philosophy clearly was not based on facts of experience; yet, its metaphysical claims about substance, soul, and the like were not true simply by definition, and the positivists thus recommended – in Hume's words – that such doctrines should be "committed to the flames."

What then was philosophy supposed to do? Philosophers were not scientists and thus were not in the business of producing and testing empirical claims. Thus the positivists claimed that they were in the business of analyzing meanings (going bankrupt was not considered a live option). Philosophers should use their skills to analyze the concepts of science and common sense. Philosophy's job was to clean up conceptual messes.

Philosophy of science thus became primarily analyzing basic scientific concepts. Those concepts could either be theory specific – like the concept of mass in physics – or be metascientific notions like "is a law," "is confirmed," or "x is explained by y." So the positivists left much for philosophers of science to do, and what they left was something only philosophical analysis could provide. Analyzing concepts was not an empirical enterprise. For example, Carnap (1950, p. v) claimed that the concept of confirmation concerned "a logical relation between two statements" that "is not dependent on any synthetic statements" and is "analytic." Philosophy thus could give us something gotten nowhere else: the meaning and logic of science's fundamental concepts.

While the positivists were philosophical radicals in many ways, they were nonetheless quite traditional in others. Epistemology in its traditional form remained healthy in positivist practice. For example, some positivists were committed to "the given," the doctrine that sensory experience directly confronts us with information that is self-evident, relying

upon no further inferences or theory. That idea had a long philosophical history and was far from radical. In the positivist's philosophy of science, the given appeared as "protocol sentences" or "observation reports" – the empirical bedrock of experience which is certain and from which all theories are derived and confirmed. Quoting Carnap (1934, p. 45) again, protocol sentences "refer to the given, and describe directly given experience"; they are "statements needing no justification and serving as the foundation for the remaining statements of science."

The positivists were traditional in yet another way. Although they ridiculed the idea that philosophers could produce a special kind of speculative knowledge, they did not practice what they preached. Under the guise of conceptual analysis, the positivists did not hesitate to rule on what constituted knowledge – on the requirements for justification, evidence, confirmation and other normative epistemological concepts. Moreover, since philosophy is not an empirical discipline, the positivists implicitly assumed that they had special methods that could distinguish good from bad science. However, this foundationalist enterprise was one Descartes, for example, could have been proud of.

There are, of course, numerous other ideas associated with positivism. But these core notions – that there is a clear distinction between empirical truths and conceptual ones, that science provides the former by beginning with the certainty of direct experience, that philosophy can, via conceptual analysis, tell us what explanation, evidence, and other scientific concepts require – formed a lasting legacy. That legacy has influenced how philosophers and scientists view the scientific enterprise and that legacy continues. Later chapters will find that unnoticed positivist assumptions still play an important part in debates over naturalism.

Positivism was officially put to rest in the 1950s. No doubt many factors were involved. The positivists themselves sometimes abandoned their most stringent views in the face of difficulties. However, the key intellectual factors were Quine's (1953) criticisms, for they sketched a broad philosophical framework that undergirds much contemporary philosophy of science. Quine's attack turned on challenging a crucial positivist distinction: that between truths based on observation and truths based on meanings. The argument worked on two fronts. Quine first argued that we have no clear notion of truth by definition or analyticity. Explaining analyticity in terms of synonymy, for example, goes nowhere. Synonymy is itself not a particularly clear notion and is most naturally defined by analyticity. Other attempts to clarify analyticity run into the same problem: either obscurity or circularity. Thus we seem to have no criteria for drawing a sharp distinction between analytic and synthetic truths.

Quine also raised a second, more systematic criticism. The analytic-

synthetic distinction tries to separate the linguistic and factual components behind our beliefs. Some statements are directly tied to confirming evidence or experience; they are synthetic. Other statements gain their credibility entirely from linguistic conventions and thus the empirical data can never refute them; they are accordingly analytic and *a priori*. Quine, however, denied that we could sharply divide evidence this way, because testing is a holistic affair. Following Duhem (1954), Quine argued that hypotheses do not confront experience or evidence one by one. Rather, testing a single hypothesis requires a host of background theory about the experimental apparatus, measurement theory, what data are relevant, what must be controlled for, and so on. So, when experiments fail, they only tell us something is wrong somewhere. We can save any hypothesis from doubt by changing our background assumptions. Theories face the test of evidence as wholes.

The moral is that evidence does not fall into two neat categories, one linguistic and the other empirical. We are faced, to use Quine's metaphor, with a "web of belief" (Quine and Ullian 1970). The evidence for any particular hypothesis depends on how well it fits with experience and our background theory. Individual observations that conflict with fundamental theoretical postulates lose credibility. Fundamental theoretical axioms that conflict with a host of empirical and theoretical data may be given up, even though earlier they looked to be true by definition. All parts of the web are indirectly relevant to all others. There is no absolute way to isolate the analytic, necessary truths from the merely empirical. In the end there are no *a priori* truths.

By denying a sharp conceptual-empirical distinction and pointing out the holistic nature of testing, Quine provided the intellectual foundations for a broad change in the philosophy of science. These two seemingly abstract philosophical ideas have wide-ranging ramifications, ramifications that are still being explored. Let me sketch those that matter most directly to the philosophy of science.

Since conceptual matters are not entirely distinct from empirical ones, philosophy of science can no longer be a purely conceptual enterprise. No doubt philosophers can still try to analyze confirmation, explanation, and the like. But those accounts ultimately will be empirical claims, not *a priori* conceptual truths. We thus have to rethink the relation between philosophy of science and scientific practice. Philosophers have no special place outside or prior to science itself. Foundationalism — as the idea that philosophers can describe on *a priori* grounds the standards for real scientific knowledge — has to go.

How then should philosophy of science proceed? Obviously it has to be tied to real scientific practice. That requirement is, however, quite weak.

Much current epistemology, for example, pays lip service to Quine's dictum. Yet it still tries to describe *a priori* conceptual truths about justification and other epistemological concepts (cf. Goldman 1986). Moreover, philosophy can be "tied to" scientific practice in different ways. Philosophers of science are still debating what the idea entails.² For Quine, philosophy of science must be "naturalized." Like the advocate of naturalism in the social sciences, the advocate of naturalized epistemology believes that human knowers are part of the natural world. They are thus best studied by the broad methods of the natural sciences. Epistemology must be an empirical discipline, one in the end that is continuous with science itself. This was Quine's message, though he did not always practice what he preached (see Chapter 6).

The idea that the philosophy of science must be continuous with science itself is widely accepted. Again, philosophers do not agree on what this stricture requires. They disagree both on what kinds of evidence are relevant to philosophy of science and on where that evidence is to be found. Some (Doppelt 1988) see philosophy proceeding roughly by the method Rawls described for political theory: by balancing intuitions (about good science) with philosophic accounts of scientific methodology until a "reflective equilibrium" is reached. Others (Larry Laudan 1990) see philosophy of science as largely making means-ends claims — judgments about methodology are judgments that particular practices best promote scientific goals. Philosophers also disagree on which empirical disciplines are most directly relevant. Quine thought naturalized epistemology was just a branch of psychology; the science that investigates our perceptual apparatus and thus our sources of knowledge. Cognitive science approaches to scientific reasoning are of a similar mind. Others see an essential role for the history of science, the sociology of science, or the numerous methodological studies inside the natural sciences themselves.

I shall not try to sort out these issues with any care. The naturalized philosophy of science employed here takes good methods to be those that promote scientific ends, above all knowledge or truth. However, making such judgments will inevitably rest on some prior sense about paradigm cases of good science; we cannot entirely avoid balancing intuitions and principles. It is also clear that psychology cannot tell us everything we want to know. The history of science and the sociology of science can help us see how science works and which scientific practices achieve their aims. Similarly, the methodological studies inside various natural sciences illustrate how to investigate empirically what promotes scientific goals. So,

² See, for example, Laudan (1984, 1987, 1990), Garber (1986), Doppelt (1990), Lepin (1990), Rosenberg (1990), Giere (1988), and Thagard (1988).

nineteenth century biology may help us evaluate reductionistic structures on good science; practical experience across numerous fields shows that double-blind experiments may be necessary to get reliable results. Thus construed, naturalized approaches do allow a place for normative evaluations – we are not stuck with simply describing what scientists do.³ Although I find this general picture of naturalism most plausible, it is not essential: the symptoms of good science I outline below and use throughout the book do not presuppose any very specific version of naturalized philosophy of science.

Quine's naturalized epistemology suggests that the positivist ideal of a universal and substantive "logic of science" is misguided. The quest for such a logic has dominated twentieth century philosophy of science. Seeing why that quest is misplaced will help make it clear just how deep the Quinean criticisms go.

For the positivists, science was a rational and objective enterprise because it had a method. By method, they had in mind a set of universally valid rules much like those of formal logic. We can tell logically valid arguments simply by their form without attending to their content. *Modus ponens* (If *P*, then *Q*; *P*; therefore *Q*) holds for whatever sentences we plug in for *P* and *Q*. We know this inference is valid without making any specific empirical assumptions. The positivists thought that similar relations should hold between data and theory: given the data, we can apply the formal and universal rules of scientific inference to determine if the theory is confirmed. In short, the scientific method was the logic of science.

A logic of science thus has at least three basic requirements. It must be *universal*. Just as the laws of logic hold whatever the subject matter or time, so too good scientific inference is universal. A logic of science also requires that scientific method be *formal*; it should rest on no specific assumptions about the way the world is. Finally, a logic of science must be *sufficient*. Given a set of premises, rules of logic suffice to decide if a conclusion follows; we need no other information. A logic of science should allow us to do something similar – to decide whether a hypothesis is confirmed given the data. While the positivists did not always explicitly endorse these requirements, they are certainly implicit in their views.

A Quinean naturalized philosophy of science makes such a scientific method unlikely for several reasons. For one, any claims about the logic of science will be generalizations from scientific practice. Scientific prac-

³ Critics also worry that naturalized epistemology rules out skepticism by fiat (see Stroud 1985). For my purposes, however, that problem can be put aside, since philosophical skepticism is a doctrine that will draw no wedge between the natural and social sciences.

tice is, however, diverse, both across time and across fields. That means finding *universal* rules will be difficult. Any rules that do describe all science must abstract from detail. However, abstract rules are unlikely to be *sufficient* because of the information they ignore. Moreover, the holistic nature of testing also makes sufficient rules unlikely. If testing is holistic, then we can apply criteria for good science only with the help of specific background assumptions. The more we abstract from concrete practice to devise universal standards, the more we will need domain-specific knowledge to apply those standards. Universal rules are thus unlikely to be sufficient.

Even if there is some set of scientific virtues that is universal and sufficient, they are unlikely to be *formal*. The positivists, remember, wanted formal rules that depend on no empirical assumptions about the world. If claims about good science are (1) ultimately empirical in nature and if (2) empirical evidence is holistic, then purely formal rules are unlikely. On the Quinean account, the rules of good scientific inference are just high-level empirical claims. What methods do promote scientific goals will depend on substantive facts about the way the world is, about the psychology of scientists, about the sociology of scientific institutions, and so on. Similarly, claims about good methodology will rely upon the relevant background assumptions. Those background assumptions will, however, be empirical in nature. Consequently, scientific methodology will rest on empirical assumptions.

Of course, the above reasoning does not *prove* that there is no logic of science. It does, however, show that the three criteria of formality, universality, and sufficiency are unlikely to be satisfied jointly. The more an account of method ignores empirical detail and domain-specific assumptions, the more abstract it becomes; the more abstract and simplified an account of method becomes, the less work it can do by itself in evaluating specific pieces of science.⁴

Classic attempts to find the logic of science have run into just the problems sketched above. I want to look briefly at why two classic accounts – hypothetical-deductivism and falsificationism – failed and how those failures have a natural Quinean diagnosis.

Hypothetical-deductivism (H-D) describes a familiar part of scientific practice. Here is how Richard Feynman (1965, p. 156) describes it:

In general we look for a new law by the following process. We guess it. Then we compute the consequences of the guess to see what would be implied if this law

⁴ Siegel's (1985) defense of universal scientific method illustrates exactly this dilemma. For Siegel, the universal scientific method consists in "a commitment to evidence." While that seems right, it is so lacking in content as to be of little help in deciding substantive questions.

that we guessed is right. Then we compare the result of the computation with nature . . . to see if it works.

Carl Hempel (1965) has made admirable attempts to turn this simple idea into a sophisticated logic of science. The general consensus, however, is that the project has failed (Glymour 1980, Achinstein 1985). Suppose our hypothesis does predict what we subsequently observe. That agreement will constitute good evidence only if we know that there is not a more reasonable rival that also predicts what we observe. However, we can always construct rival hypotheses that fit the data. So agreement of hypothesis and observation is only convincing evidence if all rivals are inadequate. Furthermore, if a hypothesis H entails our data, then so does H conjoined with any other statement. Thus the bending of light around the sun would confirm the hypothesis "general relativity is true and Ronald Reagan is a transvestite." Of course, we know that the latter is irrelevant. However, hypothetical-deductivism gives us no formal way to rule out such cases. So the moral is that deducing observations is not good enough for confirmation – we have to build in a great deal of other empirical information, information that varies from context to context. Hypothetical-deductivism thus cannot be a purely formal account of confirmation.

Furthermore, *confirmation* is only one part of theory acceptance. Accepting a theory requires more than knowing whether a specific batch of data supports a particular hypothesis. It also requires that we factor in multiple tests, the scope of the data, the logical and evidential ties with other hypotheses, and so on. Even if H-D succeeded as a logic of testing, it would fall short of being the logic of science.

Popper's falsificationism fares no better. Hypothetical-deductivism fails because logically entailing the evidence is not enough to confirm. Yet, Popper argued, logically conflicting with the evidence is enough to *disconfirm*, for we need not know about competing hypotheses to know that the evidence rejects the hypothesis at hand. Of course Popper is right about this narrow point. Nonetheless, this superficial asymmetry between confirming and falsifying does not warrant the much more ambitious claims that Popper makes about methodology.

Three key elements in Popper's picture of scientific logic are the following claims:

- (1) Falsifiability separates science from pseudoscience.
- (2) The essence of scientific testing is the attempt at falsification.
- (3) Good science is that which has survived attempts at falsification.

"Falsifiability" in the first claim requires that a particular hypothesis or theory be *logically* inconsistent with some potential data. Understood that

way, falsifiability is a purely logical matter and thus is in principle suitable for a logic of science. However, falsifiability in this sense fails to distinguish science from pseudoscience. Hypotheses *by themselves* will generally be neither consistent nor inconsistent with the evidence. Only by bringing in what we might call a "theory of the test" will the hypothesis conflict or cohere with the data. A theory of the test includes all those background or auxiliary assumptions that tie hypotheses to data. Assumptions about what the data measure, about potential confounding factors, about the experimental setup, and so on are essential before the data tell us anything about theory. Individual hypotheses are seldom falsifiable.

Of course, we can shift from the falsifiability of single hypotheses to the falsifiability of a hypothesis plus auxiliary assumptions. However, this modified criterion will not give us what Popper wants for three reasons. First, we can make any piece of nonsense falsifiable if we join it with the suitable background assumptions. For any hypothesis H , we can just add the further assumption that H entails P , where P is a description of some data. H will then be falsifiable. Surely this violates the *spirit* of Popper's criterion. However, he has no purely formal way to rule it out, as should be the case if falsifiability is part of the logic of science. Second, the broader falsifiability criterion no longer separates falsification from confirmation, though Popper thought confirmatory evidence totally inadequate. A hypothesis that entails the data is well confirmed if we know that no other hypothesis does so. For example, consider the symptoms that are the differential diagnosis for a disease. If they really are differential symptoms, then they are compatible with only one diagnosis. So a diagnosis that is consistent with those symptoms plus the background knowledge that no other diagnosis is consistent conclusively confirms the diagnosis. So falsification and disconfirmation turn out to be two sides of the same coin. Finally, allowing in the theory of test means that falsifiability is no longer a purely logical matter. Rather, to determine how data bears on a particular hypothesis, we have to bring in substantive, domain-specific background knowledge. This means falsifiability is no longer a formal criterion and thus does not give us a logic of science, however useful it may otherwise be.

Once we see the difficulties with falsifiability, we can also see that the other components of Popper's methodology – testing as attempts at falsification and surviving such attempts as a criterion for good science – are no logic of science either. A good test is not simply one where the hypothesis had a chance of being disconfirmed, nor is good science simply that which survives attempts at falsification. Failing a test sometimes reflects on the theory of the test, not the hypothesis at issue. And scientists on occasion quite reasonably ignore failed tests for just this reason. Moreover, tests with controls – the prime way we ensure a chance of discon-

firmation – can be lousy tests, if the test does not rule out competing explanations. For example, the hypothesis that oat bran lowers serum cholesterol initially was tested by comparing those who ate oat bran with those who ate a regular diet. Yet that test was a poor test, for it did not eliminate a plausible competing hypothesis, namely, that eating large quantities of grain meant eating lower quantities of fatty food. So searching for falsification provides neither a sufficient criterion nor one that works independently of substantive empirical background knowledge. As a logic of science, Popper's falsificationism fails.

We predicted on Quinean grounds that there would be no logic of science. The empirical nature of methodological claims and the holism of testing would combine to make any informative formal criteria of good science unlikely. That is exactly what we have found. Neither deductivism nor falsificationism provides a *substantive* logic of science, for each depends extensively on contingent, contextual, and substantive empirical assumptions for their content.⁵ The positivists wanted much more.

However, the moral of our Quinean story is not simply destructive. Surely hypothetical-deductivism and falsificationism get at something important. For example, the search for sharp, falsifiable hypotheses no doubt plays an important role in good science as does the process of comparing implications to data. However, if there is no simple logic of science and if ultimately claims about scientific methods are empirical claims, we cannot evaluate social science by looking at simple formal traits. We must instead look in detail at the various different claims, kinds of evidence, and concrete methods that are used in the social sciences. It is those factors that make or break a scientific enterprise, and it is on those grounds that the

⁵ The Bayesian approach can provide a very useful framework for thinking about methodological issues, as is ably demonstrated by Howson and Urbach (1989). However, it seems clear that Bayesian approaches do not provide the kind of logic of science that philosophers have wanted, for they are entirely parasitic on prior information about important methodological questions. For example, Bayes's theorem is compatible with entirely different views on such fundamental normative questions as what constitutes a good explanation, what role explanatory information plays in confirmation, whether data used in constructing a hypothesis count less than novel data, whether simplicity has probative force, and so on. In each of these cases, we could have empirical information that led to one answer or the other, and Bayes's theorem must await the result of those investigations, e.g. we might discover that novel data indicate a reliable process and thus count for more (see Maher 1988), that particular versions of simplicity do up the odds of truth, that appeals to explanatory power generally are less reliable than other empirical virtues, and that specific constraints on explanation are incompatible with our best theory of the world. All of these results could be plugged into Bayes's theorem, but only after we had decided these fundamental issues.

prospects for a science of society must be decided. The Quinean revolution tells us that there is no real alternative.

So the Quinean transformation implies much, for both how and how not to do philosophy of science. Though the implications I have outlined seem far-reaching, some think even more radical conclusions follow. The most serious challenges to scientific rationality also argue from Quinean premises. It is to these challenges that I turn next.

2.2. Varieties of rationality

Kuhn, Rorty, and the advocates of social constructivism hold what we might call an irrationalist view of science: they deny that science has any particular claim on rationality. Defenders of naturalism, on the other hand, want to show that the social sciences can be good science precisely because the natural sciences are the paradigm of a rational investigation. Yet, "rationality" is a notoriously vague word; debates over scientific rationality are often unproductive, because it is not clear just what theses are at stake. Thus, before we can consider the arguments on either side, we need to clarify the issues. Below I distinguish multiple aspects of scientific rationality and use them to identify several different rationalist and irrationalist positions.⁶

When someone claims that science is rational, they may be making a claim about

current science versus science throughout history: Arguing that past science is fully rational is clearly a much stronger thesis than that "only" current science is.

scientific change versus stasis: During periods of scientific change, standards may be in dispute. Consequently, irrationalists may grant that ordinary science is rational but deny that transitions are.

science as evaluated by current standards versus as evaluated by standards at the time: On the simplest, most naive view there has been one homogeneous activity called science that began roughly with Galileo – since then science has progressed by applying the scientific method. Such a claim is much stronger than a view that allows scientific standards to change.

evidence construed narrowly as primarily empirical adequacy versus evidence broadly construed as "good reasons": Consistency with empirical data is the traditional empiricist requirement for rational scientific belief. Critics deny that empirical adequacy

⁶ Helpful work towards this end is found in Hellman (1983) and Doppelt (1988).

plays any such role in science; they allege that simplicity, consistency with metaphysical world views, cultural norms, and so on equally motivate the adoption of scientific theories.

the reasons actually used by scientists versus the reasons that could have been given at the time. Scientists may sometimes come to believe their theories for non-rational reasons, even though they could have reached those same views from evidence available at the time.

reasons that uniquely determine which theory to accept versus reasons that are relevant to a theory. The reasons for a scientific theory may be reasons in the sense that they favor only that theory; alternatively, scientists may adopt a theory for good reasons, though those reasons did not uniquely warrant that choice.

rationality as having reasons or evidence versus rationality as truth. Scientific realists want science to be rational in the sense that it is true, approximately true, or getting truer. They obviously want more than those who claim that current science has good evidence but nonetheless may not be true.

These dimensions do not exhaust the possible contrasts: they certainly illustrate just how complex questions about scientific rationality are. For our purposes, we need not specify all possible views in this conceptual space. Instead, let me distinguish some major positions according to how far they are willing to advance down the following list of theses:

- (1) Current science has good evidence by current standards.
- (2) Past science has good evidence by past standards.
- (3) Current science developed from past science because of good evidence by standards shared at the time.
- (4) The standards mentioned in (1) through (3) are identical, that is, there are universal standards of good science.
- (5) Evidence in (a) current or (b) current and past science is primarily empirical adequacy.
- (6) Current science is approximately true.

The first thesis makes a claim only about current science. Without some further constraints on what constitutes evidence, it is a quite weak claim. Theses (2) and (3) make assertions about the rationality of past science and scientific change, respectively. Thesis (4) asserts not just that current scientific norms arose through some rational process but that those standards hold for all times and places. The fifth claim requires that "evi-

dence" be interpreted in the narrower sense of observational adequacy. When added to the other five, the sixth thesis requires not just that science be a rational process but that it succeed in getting the truth. In short, it ties rationality to the doctrine of scientific realism. That transition is one opponents of realism like van Fraassen (1981) would reject.

Radical antirationalists deny all these claims. This, I shall argue, is really the position of the social constructivists: its advocates deny that evidence as traditionally understood plays any decisive role, either in scientific change or stasis, by current standards or past. Social causes, not reasons, underlie science. Moderate antirationalists like Kuhn accept the first two theses but go no further. Kuhn denies that shared standards ground scientific change. He likewise denies that empirical adequacy plays a primary role in scientific change⁷ and denies that science is getting closer to the truth. Radical rationalists of course want it all: not just good empirical evidence evaluated by shared standards but also truth. Kitcher (1993), for example, argues for this position.

To defend naturalism I obviously need to argue that the natural sciences are a paradigm of rational investigation. However, I do not need to defend a strong scientific realism to do so. We might never know that current science is largely true even though it has good evidence. Nonetheless, natural science could still seem our best shot at the truth, given the evidence we have. Consequently, naturalism about the social sciences would remain a highly significant and contentious issue – even if scientific realism were implausible. So a significant naturalism need not affirm thesis (6).

What then does a more moderate rationalism claim? Multiple versions are possible. "Moderate rationalism" would require theses (1) and (5) only. In short, it would demand that current natural science have good evidence by current standards, with evidence defined as primarily empirical adequacy. Without this latter restriction, rationalism would let nearly everything in, since even astrology and deconstructive literary criticism provide evidence in some sense. A stronger position thus holds not only (1) and (5) but also defends the rationality of scientific change – in short, some variant of (2) and (3).

In the rest of this chapter I shall straddle the fence between moderate

⁷ If this attribution seems unfair, consider the following: "Given a paradigm, interpretation of data is central to the enterprise that explores it. But that interpretive enterprise . . . can only articulate a paradigm, not correct it. Paradigms are not corrigible by normal science at all. . . . [Transition to a new paradigm is determined] not by deliberation and interpretation [of data], but by a relatively sudden and unstructured event like the gestalt switch" (Kuhn 1970, p. 122); the relative problem-solving ability of a paradigm is "neither individually or collectively compelling" (p. 155).

and strong rationalism. Research in the sociology of knowledge has not shown scientific evidence to be a mere social construct, nor has Kuhn shown that scientific standards do not change for rational reasons. So strong rationalism (in some form) is a potentially plausible position. Showing more than that is far too great a task for this book. Fortunately, I can straddle the fence in this way without interfering with my main task – defending naturalism about the social sciences. For, if by our best lights current science is a paradigm of rationality, then naturalism is an important issue no matter how we came to practice the science that we do. Moreover, I shall at least undercut irrationalist skepticism about those origins.

2.3 Kuhn and shifting standards

Thomas Kuhn took Quine seriously and looked empirically at the sciences. He found numerous phenomena that led him to reject the rationalist's position. Kuhn's main arguments turn on three basic theses: the incommensurability of meaning and standards, the theory-laden nature of data, and the malleable nature of scientific standards. Between them, these theses still constitute the most sophisticated challenge to rationalism. I want to discuss these rationales one by one. (Throughout my concern is with the Kuhn found in the original edition of *The Structure of Scientific Revolutions*; the later Kuhn [1970, Postscript] seems to have backed away from the original radical implications and is not my subject here.)

2.3.1 *Incommensurability*

Kuhn's most radical claim is that different paradigms cannot communicate. His basic argument is this: The meaning of terms in a theory is determined by their role in that theory – by their relations to other terms, the assertions of which they are a part, and so on. However, scientific change is revolutionary; the transitions from Aristotle to Newton and Newton to Einstein, for example, were radical changes in theory. Thus Newtonian “mass” and relativistic “mass” must have entirely different meanings, since they are parts of very different theories and theories determine meaning. In short, we cannot translate across paradigms.

Kuhn's argument is implausible. It both has absurd consequences and presupposes an implausible theory of meaning and translation. The absurdity comes in two forms. If different paradigms speak in entirely different languages, then they really never disagree. Since they share no meanings, they cannot assert what the other denies (Shapere 1964). Moreover, if meaning depends entirely on the overarching theory, then every difference in theory produces differences in meaning. So when any two individuals

have different beliefs about the world, meanings will differ as well. According to Kuhn, however, differences in meaning preclude successful translation. Those who do not share Kuhn's theory of science should be unable to understand him!

These absurdities point to a deeper problem confronting Kuhn. Kuhn's arguments do not result from a well-confirmed theory of meaning and translation. No one has such a theory, and Kuhn's armchair account is unpromising for two reasons: it ignores the apparent fact that meaning dependence is not an all-or-nothing matter and that sameness of *reference* rather than sense may suffice for translation.

A term or sentence's meaning apparently depends more on some connections than on others. If meaning is holistic in this more limited way, then not every change in theory is a change in meaning and changing the meaning of one term need not change the meaning of all. Room is thus left for different theories to share meanings. Newton and Einstein may have used “mass” differently, but not so differently that no plausible translation is possible. Even if we could not translate “mass,” that hardly means that Newton and Einstein used “telescope,” “material body,” “planet,” and other more observational terms in radically different ways. So data can be shared and theories tested. Kuhn can preclude this outcome, but only by arguing that any change in theory entails a radical change in meaning. Thus either translation is possible or Kuhn is stuck with the absurdities mentioned above.

Even if meanings do vary significantly across theories, incommensurability is not inevitable. For purposes of theory comparison, the important connections may be ones of *reference*, not meaning. Reference does not exhaust meaning, because terms with different meanings can have the same reference. Since that is the case, meaning may vary across theories without preventing one theory from seeing what objects the other theory is talking about. Kitcher (1978, 1993), among others, has defended this route around incommensurability. He has argued that we can identify the referent of “phlogiston,” for example, in terms of current chemical theory. Doing so will be a complex process, for “phlogiston” may not have a constant reference in all contexts. However, once we determine the referent of “phlogiston” in the relevant contexts, then we can understand and rationally evaluate earlier chemical theories. Similar approaches have been taken to “mass.” Kuhn's other favorite example of an incommensurable term (Pearce 1987). Hence meaning variance need not prevent translation.

Kuhn has another, more plausible reason for thinking that theories are incommensurable: they invoke different standards for good science. At various times infallible knowledge, inductivist methodology, mechanical explanations, and mathematical tractability, for example, have been sine

quans of good science. If, however, such basic standards change with changes in theory, then disputes between paradigms cannot be settled rationally. Each theory will appeal to its own conception of adequacy and no rational winner is possible.

Shifting standards do not make theory choice non-rational. Kuhn errs by treating theories as monolithic, undifferentiated wholes. More specifically, he thinks scientific change involves every aspect of science – goals, methods, standards, data, and theory. In analyzing historical change, this assumption means Kuhn “telescopes” the transition from one theory to another. If we take wide enough spans of time – say Aristotelian-Ptolemaic physics at one end and nineteenth century Newtonian physics at the other – and compare theories, then we may indeed see changes in every part of scientific practice. To see scientific change this way, however, is to collapse the historical process and ignore the details of change. Long-term changes may be radical. Yet the short-term changes that realize them may be only piecemeal. If change is piecemeal, we can use one part of a scientific practice to motivate change in others. In particular, we can rationally argue for change if we share other goals, standards, data, and theory. We can then learn that some standards do not promote our cognitive goals or are inconsistent with our best theories of the world. Looking only at the endpoints of scientific change, we miss this process. Newtonian physics, for example, started with the demand that science proceed only inductively. By the nineteenth century, physicists saw that such a stricture fit poorly with their best theories. Something similar happened with the early mechanical philosophy when it required science to eschew action at a distance and provide infallible knowledge.

While the prospects for piecemeal change are thus not mere idle speculation, I obviously cannot prove my case here. Doing so would require detailed historical studies. Much such work has been done by others.⁸ Equally important, Kuhn himself does not establish that his favorite examples of paradigm shifts were entirely monolithic. Instead, he simply focuses on the endpoints of long historical changes and argues that standards, methods and theories have all changed. Such evidence is entirely inconclusive. Since my argument only presupposes that current science is a paradigm of rationality, this conclusion will do. Kuhn has not demonstrated that scientific change is irrational because of shifting standards, and there is some good evidence to the contrary. On the other hand, I cannot accept the claim that choice between *current* theories is non-rational because of diverse standards. I will take up that issue in analyzing

another Kuhnian claim – namely, that theory choice rests on theory-laden data.

2.3.2 *Theory-laden data*

Data are theory laden when the data used to test a theory presuppose the very theory at issue. However, if the theory at issue determines the data, then we should worry (a) that every test of the theory is circular, since it presupposes its own truth in the process of testing, and (b) that competing theories will determine different data and therefore no data will be able to decide between them. Rational adjudication looks impossible.

Kuhn has several reasons for thinking that theory infects data. One rationale we have already seen and criticized, namely, meaning incommensurability. Other reasons turn on the fact that different theories ask different questions and on the holism of testing. Although Kuhn is surely right about these facts, they do not warrant irrationalist conclusions.

Different theories surely do ask different questions. Yet no drastic conclusions follow. After all, geology and physiology ask different questions as well, yet that is hardly reason for thinking that all testing is ultimately circular. Asking different questions does not preclude answering those questions with neutral data according to shared standards. Moreover, theories that differ on some questions may agree on others. Even though Aristotelians, for example, asked why motion continued and Galileo asked why it changed, there were other questions and evidence they shared. Galileo and subsequent physicists showed that focusing on change in motion was the best route to answering other questions that physicists did share. Finally, great theoretical differences – differences in “paradigms” – do not always mean great differences in the questions asked. For example, physicists in the 1960s debated a fundamental tenet of modern physics: general relativity’s claim that the gravitational constant is invariant (see Will 1986). The Brans-Dicke theory, the competitor, eventually lost out after an intense debate and experimental results. Though they differed on fundamental theory, the competitors did share a broad framework of questions, enough to convince Brans and Dicke to give up their theory.

Kuhn’s last argument for theory-laden data is probably his strongest. Testing, as Duhem and Quine taught us, is holistic. Every experiment requires background theory – what we earlier called “the theory of the test” – to interpret the experimental situation. But where does the background theory come from? Kuhn thought it came from the very paradigm being tested. Thus both in obtaining data and in determining what the data entails, we use the theory at issue. So testing is ultimately circular and relies on theory not shared by competitors. Data cannot rationally decide between theories.

⁸ See, for example, Gallison (1987), Zahar (1989), R. Laudan (1981), Franklin (1990), and Kitcher (1993).

Kuhn infers from the holistic nature of testing to the claim that every test is a test of the whole. That conclusion does not follow. The basic difficulty is one we have already seen: treating theories as monolithic blocks, ignoring internal complexity. "The theory" or "the paradigm" gives a single name to what is really a diverse batch of claims, methods, goals, and practices. Although every hypothesis relies on background theory, the theory of the experiment need not be the same as the theory being tested. Because theories or paradigms are composed of many different claims and methods, we can use one part to test another. If the theory we are testing is not used to determine the data, then no circularity threatens.

How do we tell if a theory is being used to test itself? Roughly the idea is this. The background theory and the hypothesis at issue should be (1) logically and (2) evidentially independent. If the presupposed background theory entails the hypothesis under scrutiny, then obviously the test is circular. Likewise, if the evidence for the hypothesis is the same as the evidence for the theory of the test, then we are testing only the whole. However, when conditions (1) and (2) hold, we can attribute blame or credit to just the hypothesis at issue.⁹

Let me discuss several examples that illustrate the arguments I have just made. A first example comes from electron micrographs (EMs). EMs are a central tool for testing hypotheses in cell biology. Of course, EMs rely extensively on background theory. Yet, this powerful tool depends on chemistry and physics, not on assumptions about cell biology itself. Physics tells us about the wavelength of electrons and how they scatter in different materials. Chemistry tells us about the reactions involved in staining samples. Not by the furthest stretch of the imagination, however, does quantum mechanics or the theory of chemical bonding rely essentially on cell biology for its evidence. So if these tests in cell biology depend on background theory, it is an independent theory with independent evidence.

A second example, due to Franklin (1984), comes from possible tests of Newtonian versus relativistic physics. Imagine a test of these two theories that involves creating a collision between two billiard balls with known velocity. After the collision, we measure the angle of scattering and the resultant velocities. Newtonian physics predicts the angle will be 90 degrees while relativistic physics does not. While the two theories will calculate the initial masses and momenta differently, they share an important range of background theory. Both accept the obvious sensory information — a ball is observed, balls collide, and so on — as well as the initial

velocity and the manner of measuring the scattering angle. Since the two theories make different predictions about these elements, data can decide between them. This example, like that of EMs, shows in practice what I defended above in theory: the holistic nature of testing does not entail that data is theory laden in any objectionable way.

2.3.3 *Ambiguous criteria*

So far we have answered Kuhnian challenges to scientific rationality based on the incommensurability of meanings, standards, and questions and on the theory-ladenness of data. I want now to consider one last irrationalist argument from Kuhn: the idea that criteria of theory choice are inherently ambiguous.

In an essay after *The Structure of Scientific Revolutions*, Kuhn (1974b) identifies five scientific virtues: accuracy, consistency, scope, simplicity, and fruitfulness. Though Kuhn believes these criteria are common to all or most science, he denies that they suffice for deciding between competing theories. According to Kuhn, "two men committed to the same list of criteria for choice may nonetheless reach different conclusions" (1974b, p. 324). They can do so because they can either interpret particular criteria differently or because they give individual criteria different weights. Simplicity, for example, can be read in many ways. Moreover, accuracy may be gained at the expense of simplicity, but there is no algorithmic way to make such tradeoffs. As a result, even shared scientific standards will not settle debates across paradigms.

Basic scientific virtues surely do admit of multiple interpretations. Here Kuhn has the holism of testing on his side. Claims about good science are empirical claims, ones that depend on our background knowledge for their interpretation and justification. Empirical adequacy, for example, may be shared in the abstract and yet mean different things for theories with very different views of the world — just as a slave holder and an abolitionist can agree on equality as a moral virtue and yet disagree on slavery, given different empirical assumptions about human nature. Kuhn is also surely right that competing theories can succeed and fail on different criteria, leaving the two at a standoff.

Nonetheless, Kuhn's irrationalist conclusions do not follow from these useful points. Consider first theories that interpret fundamental scientific virtues differently. The question is again whether these theories share enough other background theory to resolve the dispute, just as it was in the cases of shifting standards and theory-laden data. For example, Ptolemaic astronomers agreed with Galileo and Kepler that astronomy should save the phenomena; they had different ideas about how to achieve that goal. Ultimately, their rejection of the telescope lost out. For a Kuhnian

⁹ These ideas have been developed in much more careful detail by Kosso (1989).

See also Glymour (1980).

argument to succeed, it is not enough to show that these two schools understood empirical adequacy differently. It would also have to show that they did not share sufficient background to decide, after debate and investigation, whether the telescope provided reliable information.

Different theories can also succeed on different criteria without making theory choice a "conversion experience." Competing theories frequently have different virtues. The rational response is to look for theoretical developments and further data that will decide the issue – by showing, among other things, that one hypothesis is the best *no matter what criterion or data set is emphasized*. For example, in the recent debate over continental drift, defenders and critics of that theory arguably emphasized different criteria of scientific adequacy. Supporters pointed to a variety of confirming evidence and critics demanded novel predictions. Eventually, drift theory won out because it satisfied both criteria (Laudan and Laudan 1989). So divergent interpretations do not entail irresolvable disagreement.

We have now surveyed and rejected Kuhn's radical arguments. At most he shows that irrational factors are possible. We can sketch apparently standard scientific practices that suggest that such problems are not actual. As an attack on rationalism, Kuhn's views are unconvincing. However, leaving our assessment here would be to ignore Kuhn's many subtle insights into how science works. I want to end this section by sketching some of them.

The most obvious insight is that testing is holistic. Of course, Duhem and then Quine made this point earlier, but not with the kind of depth and scope found in Kuhn. He shows how standards of good science, metaphysical world views, interpretations of scientific virtues, skills, research strategies, and much more are involved in tying hypotheses to the world. Kuhn thus gives the idea that testing relies on "auxiliary assumptions" much more meat. Kuhn also shows that science involves much more than theories conceived of as sets of statements. Heuristics, problem-solving strategies, skills, implicit knowledge, and the like are part and parcel of scientific practice. Conversely, theories and models serve many roles other than describing reality. They are embedded in scientific practice; they thus also perform many non-explanatory and non-predictive functions.

Another Kuhnian lesson is that theories are ephemeral entities considered in isolation from the concrete investigations they motivate; we must understand paradigms via the concrete practices embodying them. So "the" paradigm or "the" methods characteristic of a particular science are a gross simplification; their real content comes in their application. Applications, of course, vary, so theories and methods will take on differ-

ent readings and different meanings in different contexts.¹⁰ This means that we must look at the actual practice of science rather than just its textbook formulations.

Finally, Kuhn's incommensurability doctrine does have a useful if less profound reading – namely, that different theories need not divide the world in the same way nor even divide them in ways that can be systematically matched onto each other. This is one way to understand the idea that scientists "live in different worlds" without extreme idealist implications (see Hacking 1993).

Of course these ideas are vague as they stand. Nonetheless, we will make good use of them when we discuss the symptoms of good science, causal explanation, the role of models in economics, the relation between macro- and microsociological kinds or descriptions, and other topics later in the book. They are useful antidotes to lurking positivist assumptions.

2.4 Social constructivism and post-modernist rhetoric

While Kuhn thought science was essentially a social process, he did not completely abolish the idea of evidence and data. Some of Kuhn's successors have not hesitated to take that further step. On their view, scientific belief, like all belief, is caused by natural processes, not by such mysterious entities as "the data" or "the scientific method." The only difference between science and any other activity is a sociological one.

In this section I look at some recent work from two approaches: (1) the broad movement I have labelled "social constructivism" and (2) an approach emphasizing the rhetoric of science. Advocates of constructivism and the rhetorical approach do not always agree among themselves. I will not attempt to trace all the ins and outs in these two movements. Instead, I shall discuss the work of Bloor and Latour on social constructivism and of Rorty on science as persuasion. These are seminal figures for these approaches. Seeing where they go wrong will give us a good basis to judge both positions.

Social constructivism comes from social scientists dissatisfied with past sociology of science. Traditionally, historians and sociologists of science had distinguished internal from external factors. Internal factors concerned the data, methods, reasons, and theories present in a particular scientific period. These factors were rational ones. The social scientists first sought to explain scientific belief internally. When such explanations failed, the historian then invoked external factors. External factors were generally the social causes of belief – for example, racial prejudice or theo-

¹⁰ For a further development of this idea, see Cartwright (1993).

logical assumptions. These external factors were non-rational and the explanation of last resort.

Constructivists like Bloor (1976) find this distinction odious. As sociologists of science, their task is to explain science scientifically. Looked at scientifically, however, beliefs are natural objects in the universe; *all* beliefs are caused, even those that we normally think of as motivated by the evidence. Bloor, for example, calls this the "causality principle" and makes it a central tenet. Scientific beliefs are to be explained by social processes; they are not somehow outside the causal network and produced by such abstract entities as "the data" and "the scientific method."

Not surprisingly, Bloor also rejects the idea that there are two kinds of factors in science, the rational and non-rational. We may identify some beliefs as true and others as false, but "the same type of cause would explain both" (1976, p. 5). The traditional history and sociology of science, Bloor thinks, assumed just the opposite – namely, that social factors explain the unjustified beliefs of scientists and non-social rational factors explain the rest. Bloor rejects this asymmetrical treatment, advocating instead the "symmetry principle": all scientific belief has the same type of cause. Rationality falls by the way.

Social constructivism mingles obviously justified complaints about past sociology of science with completely unjustified inferences about scientific rationality. Surely science is a social process and scientific belief does have social causes. Yet that fact does not *entail* that evidence, reasons, scientific method, and rationality do not ground science, for several reasons.

For at least one sense of scientific rationality, the causal origin of a belief is irrelevant. In Section 2.2 we distinguished between reasons that we can give for a theory versus reasons that actually lead to its adoption. Even if scientists have their beliefs for social reasons, those beliefs might nonetheless have good evidence. For example, if I have a dream that causes me to believe arithmetic is incomplete, that does not preclude there being a proof for the claim as well. Or, racism and political agendas may explain why IQ studies are interpreted as they are, yet we can still ask whether those interpretations are valid. Beliefs can be caused and still be reasonable. Social constructivists implicitly and explicitly make these judgments regularly, even though they are strictly speaking not entitled to them.

More important, reasons can be causes.¹¹ If reasons can be causes, then scientific belief can be caused, and be caused by a social process, and still result from evidence and scientific method. A rational process causes scientific beliefs if that process works through beliefs about the data and

about the best methods for pursuing the truth. This process can likewise be thoroughly social. Beliefs about the evidence and the best methods can come from interacting with others. So long as that process results from evidential beliefs, then social causation and rationality are not necessarily incompatible. Prestige and pecuniary gain may motivate scientists. Yet if the scientific community rewards the search for evidence and good theories with such commodities, then beliefs about evidence and so on can still drive scientific practice.¹² Science will not be a mere social process.¹³

So Bloor's symmetry principle – namely, that rational and irrational beliefs have the same kind of cause – is either compatible with scientific rationality or begs the question. If the principle covers any kind of cause, including beliefs about evidence, then rationalists will not disagree. If the principle requires "the same kind of cause" to be a non-rational cause, then the principle simply begs the question. For, it now says we are to explain all scientific beliefs in terms of *non-evidential* reasons. That assumption, of course, is just the question at issue.

Finally, even if social processes that do not work through evidence cause scientific beliefs, those causes might still promote rationality.¹⁴ Suppose that in disputes over evidence we find that the most powerful scientists always win. It might nonetheless still be the case that those in power are generally *right* – that the social hierarchy on the whole plays a positive epistemic role. Naturalized philosophy of science evaluates methods by how well they promote scientific goals. So if scientific beliefs are produced by a social process, we can still ask how well that process promotes truth.¹⁵ Or, to take a more accessible goal, we can ask with what frequency meth-

¹² For an abstract model showing how pursuit of gain can nonetheless lead to truth promotion, see Goldman and Shaked (1991).

¹³ Social constructivists of course do not like the adjective "mere" here, though it is not at all clear how they can deny it is appropriate and still maintain the radical relativism of the approach. When Andy Pickering, for example, explains why the description is inappropriate, he in effect backs away from the radical versions of the approach and allows that "reasoned static appraisals of the relations between theory and evidence [are] vital to science" (1990). Pickering denies that those appraisals are sufficient to decide scientific controversy and thus his view is much akin to that of Kuhn on ambiguous criteria.

¹⁴ This latter possibility is developed by Goldman (1987).

¹⁵ Of course this presupposes some independent evidence about what the truth is. However, that assumption only looks troublesome when we focus on the rationality of theories as a monolithic unit or if we are defending rationalism in the strong form that defends the claim that we know current science is largely true. Once we recognize the diversity of claims, processes, types of evidence, and so on that make up real science, the claim of circularity is much less obvious – as our discussion of Kuhn pointed out.

¹¹ The classic defense of this claim is Davidson (1980).

ods promote empirical success. The fact that science is a social process does not tell us the answer to that question.

This third alternative thus suggests that science as a collective activity can produce rational results even if its component practitioners are less than fully rational. In short, it may be that the scientific community as a whole produces rational outcomes even though individual beliefs are not based on a correct and careful assessment of the evidence and so on. Science is a collective process, and many traditional scientific virtues such as objectivity are perhaps best thought of as properties of the community as a whole. For example, biases pulling in opposite directions may produce a collectively unbiased outcome. In short, a rational scientific community need not be a community of rational scientists. Both social constructivists and defenders of rationalism generally miss this possibility. I suspect, because of hidden individualist assumptions (see Solomon 1994 for an important exception).

So beliefs *can* be caused without thereby being irrational. A full rationalist defense would involve weaving at least the three possibilities sketched above into a story about how real science works.¹⁶ That story would no doubt be complex, but pointing out that science is a social product in no way shows such a story impossible. Thus the pessimistic conclusions of the social constructivists are at best an empirical possibility. Obviously I cannot assess here the empirical merits of every study allegedly supporting constructivism.¹⁷ However, I do want to look briefly at some of the most interesting constructivist work, that of Latour.

Latour argues that the process of science is a social process of negotiation, one in which actors seek to build bigger and bigger networks. Establishing scientific fact is establishing an unassailable network. Appeals to evidence, rationality, and nature are really strategies for defending that network; they are honorific terms for the process of social negotiation that constructs science. In *Laboratory Life* Latour along with co-author Woolgar (1979) studied the "discovery" of TRH, a pituitary hormone. They concluded that there was no logic of science that dictated when TRH was found; it was social negotiation that turned hypothesis into fact. In *Science in Action* Latour (1987) describes how scientific networks are built using scientific literature, laboratories, machines, and experiments. Facts, evidence, and rationality drop out and are replaced by network building.

Latour's work is often very insightful, particularly when it comes to describing the minutiae of scientific practice. Yet its social constructivist

conclusions are hardly forced upon us by the data he cites. Latour assumes that either science is governed by some universal algorithmic scientific method or it is all a social process of negotiation. These are not the only two possibilities. Abstract methods are unlikely to be sufficient to decide scientific debates, and when they have any content, they are unlikely to be universal – this Quinean moral we drew earlier. Yet that does not mean evidence and rationality have no place. Canons of good science, interpreted according to context and combined with background knowledge, can still be a decisive force. Nothing Latour discusses precludes this. In fact much of the "network building" he describes has a natural rationalist reading: the work of others is cited because it is data, challenging other people's network requires doing one's own experiments, i.e., testing, and so on. What Latour describes as network building can be redescribed as a search for evidence and reasons; the network comes in because evidence is holistic. In short, Latour has not established his constructivist conclusions.

So far I have argued that the social constructivists have not made their case. However, their view faces another, even more serious problem. Is constructivism itself a species of scientific activity? Its advocates generally think so. However, isn't all scientific belief caused on its view – and caused in such a way as to exclude appeals to evidence and reasons? The implication then is obvious: we have no compelling reason to accept the constructivist conclusions, no evidence that would require us to believe them. At most social constructivism can influence our beliefs, but not because their beliefs are the most rational or best supported by the evidence. In short, the doctrine appears to be self-refuting. Constructivists are, of course, free to reject the demand for reasons, and then their views will no longer be directly self-refuting. Yet, few social scientists will want to take this step with them, for few researchers really think *their own work* is only a social construction with no claim on evidence and truth as traditionally understood.

I have criticized constructivism on numerous grounds. Yet I should emphasize that my target is its irrationalist conclusions, not the attempt to understand the social processes of science. Constructivist studies bring to light the much-ignored social side of science. Their inquiries can be enormously important for exposing bias and for judging the reliability of scientific practices. This is especially so for constructivist studies of the social sciences, where politics and morals so tightly intertwine with the pursuit of fact. Shorn of their irrationalist ideology, constructivist investigations have much to contribute to good science in the social sciences.

I want to turn next to another approach that argues for similar conclusions – namely the "rhetoric of science" movement associated with post-

¹⁶ See Kitcher (1993) for a start on such an account.

¹⁷ See Cole (1992), J. Brown (1989), Gallison (1987), and P. Rohr (1987) for doubts about the empirical adequacy of their studies.

modernism. I shall focus on its most philosophically sophisticated advocate, Richard Rorty.

Starting from Quinean premises, Rorty attacks traditional epistemology. Traditional epistemology holds that "the mind" or "reason" has a nature of its own, that discovery of this nature will give us a "method," and that following this method will enable us to penetrate beneath the appearances and see nature "in its own terms" (1982b, p. 192). Less figuratively, for traditional epistemology we have real knowledge only if we identify the *a priori* constraints on rationality that guarantee that our beliefs accurately mirror the world. Rorty thinks that Quine and fellow travelers Sellars and Davidson have shown that to be a hopeless enterprise. There are no *a priori* constraints and no universal essence of rationality that can serve as the criterion for knowledge. Justification is a holistic, pragmatic affair. Knowledge has no foundations over and above human practice. The idea that our mind "mirrors" reality – that the world has one right description – is incoherent. Truth as correspondence is an appealing but ultimately useless idea. Traditional epistemology is bankrupt.

If traditional epistemology is bankrupt, according to Rorty, then so are traditional conceptions of science. There is no "scientific method" and no special sense in which science is objective, and nothing "called 'scientific status' which is a desirable goal" (1987, p. 42). To be rational is "to simply discuss any topic . . . in a way that eschews dogmatism" (1987, p. 42). The only constraints on science are "conversational"; the only useful sense of objective truth is intersubjective agreement. Thus the only important difference between science and other social practices is one of persuasion versus force. Of course, many human endeavors achieve agreement by persuasion rather than force. So the debate over naturalism rests on a mistake: it falsely presupposes that there is something to be like, an essence of science or rationality.

Rorty's view is seductive. Not only does he raise legitimate criticisms, he appeals to democratic virtues that Western intellectuals admire. Ironically, Rorty's own positive view is both poorly argued for and potentially totalitarian. He mistakenly assumes that the only alternative to the positivist conception of science is an anything-goes sociological approach. Rorty is right that scientific method is not justified *a priori*, that good science is not simply a matter of having the right method, that there may be no single method characteristic of all science, and that the notion of truth as correspondence has proven extremely hard to clarify. However, from these claims it does not follow that science is just persuasion.

As we saw earlier, scientific method can be rationally justified even if it is neither *a priori* nor universal. Rorty assumes that scientific methods either are justified *a priori* or are only conversational constraints. He has

ignored a third possibility: that we learn the methods of good science from experience. In fact, we have compelling evidence that mere persuasive power, with no further constraints, is poor reason for belief. Moreover, since most philosophers are persuaded that Rorty is wrong, this is a view he ought to endorse! Of course, this reply assumes skepticism is implausible. If we have no knowledge, then maybe conversational constraints are the only constraints. But then Rorty's claim is only a disguised skepticism, showing his view to be of no independent interest.

Rorty is right that methods do vary across sciences and do not provide a foolproof, mechanical basis for choosing theories. However, as we saw in criticizing Kuhn, changing standards need not be irrational ones. We can have good reason to change our methods: we can learn that they do not promote our cognitive goals. And, as I suggested in Section 2.1, methods will generally not settle scientific disputes *all by themselves*. However, few scientific assumptions will settle disputes in isolation. Embodied in theories and background assumptions, methods can help rationally adjudicate scientific controversy. *A priori* eternal criteria and local rhetorical devices do not exhaust the options.

Rorty's complaints about truth are reasonable. Philosophers have given no adequate substantive account of truth.¹⁸ Assume that Rorty is right that truth as correspondence has no coherent formulation. What follows? We may have trouble showing that science is rational in the sense of getting closer to the truth (defined as correspondence). Nonetheless, as I pointed out in Section 2.2, there are other important rationalist theses. Reasons and evidence can motivate scientific change and theory choice, even if truth is not guaranteed. Likewise, the world, not just conversation, can constrain what we believe even if we have no coherent account of truth as correspondence. In short, empirical or observational evidence can still be the heart of good science. We can show science is rational in this sense without showing that current science is largely true. Rorty mistakenly equates rationalism with its strongest variants.

The social constructivists and Rorty thus point out important inadequacies in past pictures of science. Both err, however, in assuming that the sociological approach is the only alternative. It is not.

2.5 The subtle invasion of values

Lurking behind the irrationalists' views is another, older argument: that there is no fact-value distinction and thus that science is ulti-

¹⁸ I assume here that disquotational theories are not "substantive" – they are compatible with correspondence, coherentist, etc. accounts and are less than what strong scientific realism requires.

mately a kind of moral choice. If science does essentially involve moral or political choices, then there are serious doubts about scientific rationality. There are well-known reasons for thinking that moral beliefs are subjective. Even if moral beliefs can be rational, they are not likely to be so in the full-bodied sense that we standardly attribute to natural science. I thus turn in this section to consider challenges to the rationality of science based on fact and value.

Let me clarify some preliminary issues before turning to the main arguments:

- (1) The quest for a value-free social science is not a quest for science that presupposes no value judgments. Science essentially involves innumerable judgments about what is good and what ought to be done. However, value assumptions are problematic only if they are moral or political values – as distinct from epistemic values. Reliability, objectivity, fruitfulness, scope, and so on are important values in science, but they are epistemic values. The interesting thesis is thus that science presupposes *non-epistemic* values.
- (2) Any interesting argument for value-ladenness needs to show that science *inevitably* makes moral evaluations – that the scientist qua scientist makes moral judgments. Pointing to racial or sexist biases and the like is not enough.¹⁹
- (3) A value-free science does not require that morality is inherently subjective. In other words, the fact-value dichotomy and the objective-subjective dichotomy are not necessarily coextensive. Some or all moral statements might well be objective. So arguments that we can give reasons and evidence for moral claims are not directly to the point.

In sum, those who argue that science is value laden must show that science essentially presupposes non-epistemic values and that these values vitiate the prospects for a neutral, objective scientific enterprise. These may sound like tough requirements. However, they are requirements current commentators claim to meet.

Gunnar Myrdal (1970) argued – long before positivism had exhausted itself – that science did not proceed by taking entirely neutral data and applying a formal logic of confirmation to prove or disprove theories. Myrdal saw that this more complex picture of scientific practice made it

much harder to sharply separate scientific from moral concerns. Mainstream economic theorizing about underdevelopment, he argued, focused on equilibrium models; by doing so, it emphasized the virtues of free markets and ignored the failings of capitalist development. For Myrdal, this was just one example of a general and unavoidable problem: we develop theories and gather data only after we have decided what are the interesting questions. But, it is not a scientific fact which questions we should pursue. Thus science is inherently value laden. The best we can do is make our value assumptions explicit.

Myrdal's remarks are insightful and suggestive. Nonetheless, he has not shown that values permeate science. Theories do presuppose some prior decision about which questions to pursue. It does not follow, however, that every question we ask is a subjective matter. We pursue some questions because they result from our best theories. "What is the nature of genes and inheritance?" is not just something biologists found interesting; it was a question imposed in clarifying evolutionary processes. Moreover, even when extra-scientific considerations determine which questions we ask, our answers need not be subjective. Once we set the question, it may be a perfectly objective, non-normative matter what the answer is.

A different argument comes from Rudner (1953). Scientific evidence is always a matter of degree. At some point we have to decide if our evidence is enough, whether our predictions sufficiently correspond to the facts to warrant belief, and so on. Nothing in the scientific process, Rudner claims, tells us when the evidence is good enough. Instead, decisions to accept a hypothesis depend on what we would lose by being wrong and gain by being correct. Evidence that is clearly good enough in a circumstance when we have little to lose from being wrong may not be good enough when circumstances are reversed. Thus, accepting hypotheses is inherently value laden.

Rudner's argument turns on an important truth: accepting a hypothesis involves more than its scientific virtue. However, that fact – one emphasized by Kuhn and other post-positivist writers – does not entail that science is value laden. The reason is relatively simple. We can still distinguish scientific or epistemic reasons from normative ones.²⁰ As real people, we take both into consideration. No doubt when deciding on whether to accept, or act upon, or even entertain a hypothesis, pragmatic or normative consequences are relevant. Yet, their relevance is an independent matter. We can weigh the cost of error, but we must also know the *probability* of error. In short, the degree of evidence and the cost of being wrong are two separate factors, and the former can remain an entirely objective matter.

¹⁹ To be fair to the defender of value-ladenness, "unavoidable" here is something short of logical or conceptual necessity. If we can show that values are built in for all practical purposes, then the rationalist position looks defeated.

²⁰ This line of argument is advanced by McMullin (1982).

The arguments of Myrdal and Rudner ultimately turn on the holistic nature of belief. Longino (1990) has recently argued directly from the holistic nature of testing and theory to the conclusion that values are a necessary part of science. Following Quine and many others, Longino argues that data by themselves do not tell us what to believe, for their implications and even description depend on background assumptions. So no scientific method can guarantee that values are removed from the scientific enterprise. Values are ultimately built into science, because they are part and parcel of the assumptions we use to make science. Does this then make science merely a reflection of individual values? Longino does not think so. Science is essentially a social enterprise – scientific results are subject to public scrutiny and check. This assures us of an intersubjective constraint that removes subjective *individual* preference from science.

Longino nicely identifies ways in which values can infiltrate science. She also argues convincingly that some research on gender differences reflects various androcentric or sexist assumptions. She correctly points out that, given the holism of testing, value-free hypotheses tested against value-free data may still involve value-laden background assumptions. However, this does not show that *good* science can be value-laden science. Good science requires more than testing factual hypotheses against neutral data. It also requires that the theory of the test meet scientific standards as well. When statistical inference, for example, rests on unsubstantiated and biased assumptions about the relevant causal variables, linearity, and so on, bad science is at work. However, the formal rules of statistical inference do not prevent such sloppy research. That does not mean we must count such research as good science, only that we must evaluate it on substantive and empirical grounds. Background assumptions supported only by moral judgments make for bad science. However, if we look at the research Longino cites, biased or bad science is what we see: research based on shaky, often unacknowledged background assumptions that reflect androcentric stereotypes. Nothing here shows that values can or must permeate good science.²¹

So none of these arguments show that non-epistemic values inevitably influence science. Of course, if natural science is not inherently value laden, *social* science might be. We will consider some arguments for this idea in later chapters. However, even if values do not inevitably penetrate

²¹ Longino (1990, p. 100) has another argument based on the idea that values are built into "the constitution of the object of study." The idea seems to be that in defining a scientific domain non-epistemic values are inevitable. This raises issues too complex to deal with here, but I am unconvinced by this reasoning as well – see Kincaid (in press b), where I discuss such arguments in the context of economics.

social science, nonetheless *bias* is a much more troublesome problem in the social sciences than in the natural sciences. Longino and the others have helped me see just how subtly values can penetrate science.

2.6 The symptoms of good science

This chapter set out to do three things: answer irrationalist critics, sketch some post-Kuhnian philosophy of science, and identify features of good science. We have addressed directly the first two tasks. I want now in this final section to take up the third task. Doing so is crucial to defending the naturalist position advocated throughout the rest of the book.

To show that the social sciences can be scientific, we need some account of what being scientific comes to. Time and again I shall argue in later chapters that the social sciences can achieve the standards of good science and that some research actually has done so. However, attempts to identify the defining features of science have a long and disappointing history. Furthermore, the view of science and scientific method I used in answering the irrationalists seems *prima facie* incompatible with defining good science. If there is no purely formal, *a priori* scientific method but only contingent, historically specific empirical evaluations that must be embodied in concrete theories, how can we say anything useful about "the" methods of good science? How can we say that areas as diverse as molecular biology and sociology use the same methods or have the same virtues? Thus before discussing the traits of good science, I need to address these procedural questions.

Past attempts to define science – such as Popper's falsifiability criterion – have indeed failed. The positivists' goal, however, was much more ambitious than anything needed for the argument advanced here. I make no claim to have the individually necessary and jointly sufficient criteria that uniquely divide all inquiries into two groups, namely, the scientifically respectable and the pseudoscientific. Instead, the criteria I advance are the symptoms of good science; they are indicators of good science that make no claim to completeness or perfection. To argue for naturalism, I need not show that there is a sharp distinction between the scientific and pseudoscientific. Rather, I need to show that the social sciences share the basic virtues common to paradigm instances of good science – that the social sciences are on the scientific end of the science/non-science continuum, if there be such.

The naturalism I defend here also does not argue that the social sciences meet criteria that characterize good science throughout history. There might be no such criteria, even loosely construed. My argument seeks "only" to establish that the social sciences can be sciences according to criteria we *now* take to define good science. (The rationalism I am presup-

posing is a moderate one – recall the discussion of Section 2.2.) That conclusion, however, is still highly contentious. And if current criteria for good science do reflect universal standards, my arguments entail that the social sciences can meet them.

Furthermore, areas as different as molecular biology and sociology can share the same virtues. It is useful here to distinguish between *abstract* and *realized* virtues. Scope, simplicity, accuracy, falsifiability, and so on are abstractions in that they do not directly mention specific empirical presuppositions. Realized virtues, on the other hand, are the ways a specific science embodies abstract virtues in the relevant empirical detail. Thus simplicity in nineteenth century geology became a principle of uniformitarianism;²² falsifiability or objectivity in current clinical research is realized in double-blind experiments, for example.

Obviously, at the most concrete level, the social sciences do not and cannot share the virtues of the natural sciences. Yet, at the most concrete level even the disciplines of physics do not share the same methods either. Nonetheless, it can still be the case that (1) there are broad virtues at some level of abstraction that characterize both the natural and social sciences and (2) the social sciences can and do successfully embody or realize those virtues. Both propositions are essential. Abstracting broad virtues from their concrete applications allows us to identify the core of scientific rationality and to avoid concluding that there are as many scientific methods and virtues as there are scientists. Yet simply identifying abstract virtues like scope and fruitfulness is not enough. A convincing argument requires showing how those virtues get embodied in social science research, a lesson we also learned from Kuhn. This two-sided approach avoids a universal *a priori* logic of science without denying that there are general standards of scientific rationality;²³

Nonetheless, it might turn out that no unitary set of virtues emerges even after abstracting from contextual details. As we will see shortly, that may well be the case for non-evidential constraints on scientific rationality. Yet we can still defend naturalism. Naturalists can grant that there are multiple ways to do good science. However, they cannot allow that those methods are all beyond the reach of the social sciences. Nor can they allow that the social sciences have special methods that do not embody some basic scientific virtue guiding the natural sciences. One common substantive set of scientific virtues makes arguing for naturalism easier, but it is no prerequisite.

What then are the symptoms of good science? For our purposes it will

²² I take this point from Sober (1988).

²³ Here I thus disagree with Rouse (1987, p. 124).

be useful to break “science” into two components – the process of scientific inquiry and the products of that investigation. Traditional philosophy of science generally concentrated on the latter, namely, on theories and their relation to the evidence. However, Kuhn and the social constructivists have highlighted just how one-sided that emphasis is. Science is above all an activity, and evaluations of good science that ignore that fact are at best incomplete. Thus the symptoms of good science involve both the symptoms of good products and the symptoms of good processes.

Traditionally, philosophers of science have identified symptoms of a good scientific *product* that fall into three broad categories: (1) evidential, (2) explanatory, and (3) formal. Evidential virtues concern how well data support a theory. It is these virtues that philosophers of science emphasize. However, science aims for more than true or well-confirmed statements. A phone book contains numerous truths, but it lacks an essential scientific virtue: the ability to explain. Good science does not just describe, it explains.²⁴ Finally, it is sometimes thought that good science does more than just provide true or well-confirmed explanations – it produces *theories* with certain formal properties. Theories, on this view, are systems, where a “system” is usually equated with explicitly stated bodies of propositions, universal in scope, that can be axiomatized and deductively organized.

The symptoms of good scientific *practice* are those that indicate reliability – in short, they show that a given practice brings about good scientific products. Philosophers of science have ignored this side of science, so we have no standard typology. We can categorize good processes by the different components of reliability: a process is reliable if it produces good products in large quantities or efficiently or quickly and so on.²⁵ We can also classify virtues depending on whether we describe them epistemically or non-epistemically. An “impartial” process, for example, is one that is epistemically described, for it is tied to producing a good scientific product. A “competitive, interest-driven” process has no such connotation, though *a priori* it is not precluded from reliably producing good science as well.

Throughout this book I shall focus primarily on three kinds of symptoms: evidential and explanatory virtues of scientific products and the epistemically described virtues of social processes. I downplay other kinds of virtues for two reasons. Chapters 3 and 7 will argue that formal virtues

²⁴ Of course some philosophers of science – instrumentalists – may argue that good science is only empirically adequate, yet they too typically recognize that explanation is in some sense essential to good science and thus attempt to show how explanatory virtues can be collapsed into evidential ones.

²⁵ Here I am borrowing from Goldman (1986).

are relatively unimportant; good science does not require deductively systematized theories of universal scope. Thus I will not discuss those virtues any further here. I downplay the symptoms of good processes for a very different reason: we have no well-developed account of what processes produce good science, at least when those processes are not described epistemically. Psychology and sociology have much to learn about what factors make for good science, and I thus cannot draw on an established body of knowledge. That does not mean we can say nothing. Chapter 8 will make some limited claims about these topics when we discuss why the social sciences are not more effective. However, those judgments are tentative. Thus the main argument for naturalism will concentrate on showing that the social sciences can and sometimes do use good methods to produce well-confirmed explanatory accounts.

With these methodological preliminaries out of the way, let's turn now to look concretely at the symptoms of good evidence and good explanation. We can describe these virtues at different levels of abstraction. At the most abstract level are the standard, one-word virtues such as scope and fecundity. At the most concrete level, these symptoms become double-blind experiments, Durbin-Watson tests for autocorrelation of time series data, and other such domain-specific practices. Between these two extremes are generalizations about good science that hold across disciplines and yet are more informative than the simple list of scientific virtues. What, for example, are the general practices that specify empirical adequacy, fruitfulness and objectivity? In what follows I discuss the idea of good science at the first two levels; the third level will have to wait for later chapters when specific empirical work is at issue.

At the most abstract level, good science typically has at least the following *evidential* virtues:

Falsifiability. While falsifiability is far from the entirety of good science, it is the first line of empirical adequacy.

Empirical accuracy. The more predictive success, the better – measured in terms of both quantity and quality.

Scope. We want theories that not only repeatedly predict the same kind of phenomena with precision but also that predict a wide variety of different kinds of phenomena.

Coherence. A good theory coheres with our best information from other sciences. Logical consistency is a first start, tight evidential interconnections, the ideal.

Fruitfulness. Theory evaluation is not just evaluation at a time; it requires looking at past track record and future promise.

Objectivity. Standard lists of scientific virtues seldom explicitly mention objectivity, apparently assuming that it falls out from the above traits. But Kuhnian worries about incommensurability and doubts about value-free science should suggest that empirical success may not guarantee objectivity. Nonetheless, the ideal of an unbiased, disinterested pursuit of the truth is the hallmark of science. Science is objective when our beliefs reliably indicate the way the world is rather than the way we want the world to be.²⁶

These traits help separate good science from bad. Yet, they are also abstract and simplistic. By themselves, these standard criteria will seldom resolve disputes over scientific adequacy – because they admit of multiple interpretations, because we do not know how to measure them or trade one for another, because they are simplifications that hold only *ceteris paribus*. Hence we need to move to the more concrete practices that realize these virtues. In particular, I want to focus on three broad features of good empirical tests. Good science undergoes *fair tests*, *independent tests*, and *cross tests*. Before we spell out these tests in detail, let me explain why these traits are important.

On the most philosophical level, being tested against experience is the basic defensible idea from empiricism. If empiricist accounts of meaning, of demarcation, and of theory and its relation to observation ran into fatal criticism, the idea that knowledge requires testing beliefs against experience did not. Fair tests, cross tests, and independent tests are a concrete embodiment of that empiricist tenet. Testing virtues are also important because they bring in the active side of science. Focusing on the relation between observation statements and theoretical propositions, as so much philosophical confirmation theory has, violently abstracts from real science. Appeals to “the data” summarize a history of hypothesizing and manipulating the experimental conditions that lead to the end product. Focusing on testing virtues acknowledges that real confirmation is essentially a process and thus gives us virtues of scientific practices as well as scientific products.

Of course, the best rationale for focusing on these testing virtues would be a full empirical theory of methodology replete with detailed evidence. I clearly have no such theory nor is there one extant in the literature to

²⁶ This intuitive idea, which lies behind the more concrete traits that I discuss below, has its origin in psychological approaches in epistemology and is developed in connection with philosophy of science by Miller (1987).

draw upon. Nonetheless, the kinds of tests I describe here do loom large in several careful recent accounts of modern experimental physics, namely, those of Franklin (1986, 1990) and Galison (1987). Their work also makes some effort to show that appeals to such tests are not merely rhetorical devices in a social struggle between scientists. I can offer no further support here, though presumably readers will find these tests characteristic of the natural science they know best.

Let's look now at these testing virtues, starting with *independent* tests. An independent test is one which tests the hypothesis at issue rather than the entire theory to which it belongs. As we saw in Section 2.3, the holistic nature of theories does not preclude directing tests against specific hypotheses. To do so the theory of the test and the hypothesis at issue should be logically independent. The theory of the test includes all the background assumptions needed to generate data relevant to the hypothesis at issue. Those assumptions and the hypothesis being tested are logically independent when the falsity of either one would not entail the falsity of the other. In short, in an independent test the hypothesis does not presuppose the truth of the theory used in constructing the instruments, identifying the initial conditions, interpreting the data, and so on, nor do those theories presuppose the truth of the hypothesis.

Tests are independent in a second, equally important sense when the theory of the test has its own evidence. If we accept the theory of the test for the same reasons we accept the hypothesis at issue, then our test is compromised – especially if the test assumptions are improbable given what else we know. In other words, in the ideal test the hypothesis is both logically and evidentially independent of background theory. The physics that grounds electron microscopy does not depend on what we know, for example, about how cells ingest large proteins.²⁷ So when an electron micrograph shows a cell wall enclosing a particle, we have a quite independent test of cell biology. If, on the other hand, our only reason for believing in the data from EMs was the same as that for the endocytosis hypothesis, our test would not be independent.

Independent tests thus give content to the abstract virtues listed above. Theories without independent tests are hard to falsify. They also obviously lack objectivity, for the evidence for the hypothesis depends essentially on our background beliefs, not the world. Lack of independent tests also

indicates lack of coherence, for the theory being tested is only integrated with itself.

Good theories also survive *fair tests*. The basic idea behind a fair test is that a good theory rules out its competitors in an objective fashion. Fair tests must first be independent tests. Theories that presuppose their own truth are not winning fairly. But independent tests are not enough. Fair tests depend crucially on competitors sharing sufficient background to pick a winner. That requires sharing, first, language, and second, the assumptions essential to the theory of the test. Section 2.3 discussed how different theories can share language. Ideally, both theories share a range of terms used in constructing the test. When that is not possible, theories must at least be translatable – there must be some way, perhaps complex, to capture the reference of key terms. Beyond sharing language, a fair test requires that no assumptions are made that bias the test towards one hypothesis rather than another. Competing hypotheses need not share all background theory, but they do need to share enough to generate unbiased tests.

Thus theory evaluation is essentially comparative. Even if we had some metric for virtues like predictive accuracy and scope, those measures would not be good enough, because they do not ensure a fair test. Without common data and assumptions, simply counting up successful predictions would not tell us which theory was preferable. Of course, no individual test is by itself logically decisive. But, as Lakatos (1970) saw, repeated failures of fair tests tip the scales. Fair tests are thus the essence of scientific objectivity.

Finally, good science is *cross tested*. Cross testing involves numerous procedures – for example, triangulation, bootstrapping, and data-reliability analysis – that are part and parcel of good science. The root idea is roughly to make sure that it is the world, not our background assumptions, which leads us to accept a hypothesis. Just as we use Mill's methods of variation to isolate causes, so too varying the different components of a test helps show that our hypotheses have really been tested. Cross testing therefore involves varying some assumptions while holding others constant. Ideally we would like to test a particular hypothesis using diverse background assumptions. Similarly, we would like to see the same theory of the test employed to investigate a variety of different hypotheses. For similar reasons, multiple measures of the data are desirable. These processes put the holistic nature of theory and evidence to good use. Far from entailing that theory evaluation is subjective, the web of belief places numerous constraints on theory acceptance. As evidence, hypotheses, and background assumptions are varied and checked against each other, the odds diminish that successful tests happen just by chance.

²⁷ The situation is, however, more complex than this statement indicates. If I have independent evidence for endocytosis in a given cell line, then successful EM depiction of that process is evidence in turn for the reliability of EMs. Obviously, a more sophisticated account calls for identifying degrees of independence, primary evidence, and the like.

The modern theory of the gene nicely illustrates cross testing. Postulated to explain the facts of inheritance, the gene hypothesis is cross-checked by embedding it in what we know about cell functioning, biochemistry, and evolutionary change. Biochemistry of course makes innumerable predictions other than those about genes and is cross-tested against our background knowledge of both physics and cell functioning. Cell biology is in turn tested by what we know about chemistry and physics (as in the electron micrograph example), physiology, anatomy and the like. Finally, the gene postulate likewise generates specific predictions when combined with evolutionary theory. Of course, these areas are really large complexes of diverse claims and evidence, independent in varying degrees. As a result, the actual cross checking is enormous.

Like independent and fair tests, cross checking also gives substance to the more abstract goals of empirical accuracy, scope, and coherence. Replication and randomization are varieties of cross checking that try to vary *unknown* background assumptions. Diverse kinds of evidence are important, because different evidence requires varying the theory of the test. Integrated theories are important not simply or even primarily because we value organized wholes; more important, integration forces us to cross-check, upping the odds of truth. And, while not enough to ensure objectivity, cross checking certainly makes bias more difficult. When combined with independent and fair tests, cross checking is a powerful method.

These three kinds of tests give us a more concrete guide to good science. They will play a key role as we defend naturalism in later chapters. There I shall argue both that the social science research can provide such tests and that no current social research confirms without doing so. There can be and is good social science research – and it is the result of fair, independent, and cross testing.²⁸

I want to turn now to symptoms of good scientific explanations. Numerous such traits are cited by philosophers and scientists. Identifying mechanisms, providing relevant laws, citing unobservable entities, unifying diverse phenomena, and describing causal processes are mentioned most frequently. Unfortunately, all of the above constraints have been dis-

puted at one time or another – basically because philosophers disagree about what explanation itself is. Thus it is harder for us to give any non-controversial account of what promotes good explanation.

Most approaches to explanation do, however, fall into two broad camps: those that emphasize causation and those that emphasize unification.²⁹ On the latter view, explanation comes from showing how diverse phenomena fit into a common pattern. When we can account for phenomena by invoking fewer assumptions, we improve our ability to explain. Newton took the apparently diverse phenomena of tides, projectiles, and planetary motion and showed how they all followed from one set of principles. Darwin did the same for the many different traits and behaviors of plants and animals. Explanation seems to come from unification.

If explanation as unity emphasizes the global, the causal approach starts with more local facts. Science at its best gives us well-confirmed accounts of particular causal processes. It is those processes which the experimental method best illuminates and which tell us how the world works. Theories of broad scope may be possible, but their explanatory power comes from citing causes. Particular causal processes may be instances of general causal laws. Low-level causal generalizations may in turn hold because of more fundamental and general causal processes. So the causal approach can allow unifying laws, both fundamental and derived. Yet those laws are explanatory only because they summarize the causal facts.

Both approaches, I should note, can allow that pragmatic factors play a role in explanation. Philosophers have had some success in clarifying pragmatic factors in explanation (Achinstein 1980; Garfinkel 1981; van Fraassen 1981). One helpful suggestion is that explanations are answers to questions. Questions, however, do not carry their meaning on their face: the same question can call for different answers depending on the context. One contextual factor is the *contrast class*: “Why did Adam eat (rather than throw, etc.) the apple?” calls for a different answer than “Why did Adam (rather than Eve, the snake, etc.) eat the apple?” Philosophers have developed these ideas in some detail.

The causal and unification views incorporate these insights as helpful emendations. For the causal view, different questions and different background knowledge means focusing on different strands in the causal net.

²⁸ These three virtues have a natural Bayesian reading. Fair tests ensure that alternative explanations are used to calculate the probability of the data, among other things. Cross tests provide ways of determining whether $P(E/H)$ really rests on auxiliary assumptions, of assigning a low probability to the hypothesis that the data result from artifact, of raising $P(H)$ by showing it coheres with other work, and so on. Independent tests are ones where $P(H/A) = P(H)$, and A is the auxiliary test assumptions. It can be shown that with such an assumption the evidence will bear differently on the hypothesis and the theory of the test. See Howson and Urbach (1989, pp. 96–102).

²⁹ This division is not the only way to categorize different accounts: Wesley Salmon's (1984) epistemic vs. ontic classification, for example, is similar to but not the same as the unification-causation dichotomy, for there might be, for example, other real-world relations besides causation that explained. Nonetheless, my discussion of explanation here and elsewhere in the book draws heavily on Salmon's general approach.

When the highway engineer and the car manufacturer ask what explains the accident rate, they may want different answers. Those answers focus on different parts of a complex total cause, but they are still nonetheless *causal* explanations. The unification view can take a similar tack. Pragmatic factors may determine the connotation of questions and the relevant kind of answer. But an answer that explains is still one that provides the most unified story.

How do we go about defending naturalism, given these two different pictures of explanation? I can see three basic approaches: (1) The "separate but equal" strategy opts for a plurality of explanatory virtues. Causation is the key to explanation in some domains, unification the key in others. If they go this route, naturalists must then argue that the social sciences can and must provide good explanation by one of these two standards. (2) The "unity in diversity" strategy would argue that causation and unification are really two compatible aspects of good explanations.³⁰ Causation emphasizes that part of explanation that involves an objective relation in the world, what Salmon calls the ontic notion of explanation; unification focuses on the epistemic side of explanation – it clarifies the idea that explanations make something more familiar, expected, or understandable. Naturalists taking this route then would ideally like to show that the social sciences can and must produce explanations that are causal and unifying. (3) The "eliminativist" strategy would argue that one or the other virtues was really derivative – that causation only explained when it provided unity or vice versa. A naturalist who adopted this strategy would thus argue that the social sciences could and must have good explanations by whatever criterion was basic.

My defense of naturalism will combine these three strategies. My main concern will be to show that the social sciences can produce well-confirmed causal accounts. I will forestall worries that unification is the essence of explanation in several ways. First, I shall assume that citing causes does explain and thus that any theory of explanation must be consistent with this fact. Since defenders of unification typically grant this point, my assumption is a safe one. Secondly, I shall argue in the next chapter that unification is not *necessary* to explanation and probably not sufficient. So showing that the social sciences can give causal accounts is showing that they can explain.

What about the second naturalist thesis – namely, that the social sciences *must* proceed along natural science lines? I will adopt the "separate but equal" strategy and argue that all social science explanations involve

either unification or causation. The social sciences, in short, have no special basic routes to explanation. When social theories unify but do not cite causes, our assessment of them is tentative. If inquiry ultimately shows that explanation is unification, those theories will explain. If, as I would suggest, citing causes is the prototype for explanation, then those theories may serve other roles. However, whatever the outcome, the same conclusion will hold for the natural sciences, since unification of course plays an important role there. So explanation will draw no sharp divide between the natural and social sciences.

Lest it seem that I will be arguing for a radical monism, let me emphasize a point made in Chapter 1: basic scientific virtues common to all science are compatible with great diversity in scientific practice. The virtues described here are abstractions. Those abstractions may be embodied in many different specific explanatory strategies and confirmational methods. In fact, our Quinean philosophy of science expects such a plurality. For example, even if the social sciences always gave us causal explanations, it could nonetheless give us many different *kinds* of causal accounts. Explanations emphasizing structural causes, causal redundancy, micro-causal details, functional causes, and so on might well be appropriate, depending on our background knowledge. We can allow and indeed expect a similar diversity when fair tests, cross tests, and independent tests are embodied in real research. Thus nothing in my account rules out "methodological pluralism," properly understood.³¹ Naturalists deny only that those methods are outside the realm of natural science – that they are in no significant sense embodiments of basic standards of good science in the natural sciences. So arguing for naturalism requires seeing exactly how and if social scientific practice realizes these abstract virtues. It is to that task which we turn next.

³¹ As, for example, defended by P. Roth (1987).

³⁰ Salmon (1989, p. 180) has suggested that some such reconciliation might be possible.