

Naturalism without Scientism

P. KYLE STANFORD

Integrative Naturalism

Let us begin with naturalism itself. These days, most philosophers seem eager to regard themselves as naturalists, but this is possible only because there is so little agreement about what naturalism actually involves or requires. There is, of course, a straightforwardly (and self-consciously) scientific form of philosophical naturalism, which simply assumes that the claims of our best scientific theories should serve as both starting point and boundary conditions for any serious form of inquiry. Alexander Rosenberg, for instance, boldly announces that “Naturalism is the label for the thesis that the tools we should use in answering philosophical problems are the methods and findings of the mature sciences – from physics across to biology and increasingly neuroscience” (2014, 17). Likewise, according to the “radically naturalistic metaphysics” defended by James Ladyman and Don Ross (2007, 1), “Naturalism requires that, since scientific institutions are the instruments by which we investigate objective reality, their outputs should motivate all claims about this reality, including metaphysical ones” (2007, 30). Like Rosenberg, Ladyman and Ross cheerfully plead guilty to the charge of “scientism” (2007, 65), and they insist that metaphysical claims should not even be taken seriously unless they seek to unify the independent findings of our best scientific theories:

Any metaphysical claim that is to be taken seriously should be motivated by, and only by, the service it would perform, if true, in showing how two or more specific scientific hypotheses jointly explain more than the sum of what is explained by the two hypotheses taken separately, where a “scientific hypothesis” is understood as an hypothesis that is taken seriously by institutionally *bona fide* current science. (2007, 30)

Philosophers and other nonscientific inquirers are thus advised to reform their investigations so as to begin from the results and methods of institutionally *bona fide* scientific inquiry itself, or at the very least not stand in the way of making further progress in such

inquiry; as the song says, “get out of the way, if you can’t lend your hand, for the times they are a-changin’.”

Though combining such ardently scientific varieties of philosophical naturalism with resistance to scientific realism might indeed be incoherent or a contradiction in terms, I see no compelling reason to embrace their dogmatic starting point. It is, of course, true that any satisfying form of naturalism will have to make good sense of the dramatic and undeniable practical achievements of our best scientific theories, but a more appealing and systematic motivation for this demand can be found in W.V.O. Quine’s (1951/1980) suggestion that there is only a single, interconnected project of inquiry into the world and our own place within it, and his insistence that this single integrated inquiry is responsible for simultaneously making good sense out of *all* of the reliable evidence we have available to us at any given time. Such evidence will certainly include the impressive feats of prediction, control, and putative explanation achieved by our best scientific theories, but it will include much more in addition. This *integrative* form of naturalism will require that we also make sense of the similarly impressive practical and explanatory accomplishments of many past successful scientific theories that we now regard as offering quite profoundly mistaken accounts of the fundamental constitution of some part of the natural world, of the radical discontinuities between succeeding generations of empirically successful theories in the same scientific domain across the history of science, of our repeated failure to even conceive of scientifically serious alternative theoretical possibilities that were well confirmed by the evidence available at the time they remained unconceived, and of much else besides. A convincing integrated conception of ourselves and our place in the world must be no less responsive to such historical evidence than it is to the impressive experimental demonstrations, explanatory achievements, and other sorts of empirical evidence we have found to support contemporary scientific theories themselves.

As Quine suggests, the naturalistic philosopher pursuing such an inquiry

begins his reasoning within the inherited world theory as a going concern. He tentatively believes all of it, but believes also that some unidentified portions are wrong. He tries to improve, clarify, and understand the system from within. He is the busy sailor adrift on Neurath’s boat. (1975/1981, 72)

Such a Quinean naturalist will not *begin* her inquiry by privileging any part of her “inherited world theory” – her efforts to “improve, clarify, and understand the system from within” are responsible for making sense of the entire body of empirical evidence she takes herself to have at any given time. The only truly consequential distinction between *kinds* of evidence that such a naturalist will recognize is between *good* and *bad* evidence, and even that division will have to be earned from within the ongoing project of inquiry itself.

At the heart of any appealing naturalistic conception of inquiry, I am suggesting, is neither a dogmatic scientism nor even Quine’s own famously overzealous holism about testing or meaning, but instead simply the broadly Quinean insistence that *all* (good) evidence matters, and a consequent refusal to automatically elevate scientific, historical, philosophical, or any other particular form of evidence into a privileged position while simply neglecting the rest. On this view, evidence concerning the distinctive characteristics of our scientific theories and ourselves as theorizers must feed back to inform our view of what we are doing when we theorize about the world, of what we can and cannot reasonably expect such scientific theorizing to accomplish under various sorts of epistemic conditions,

and of the character, reach, and epistemic status of the scientific theories generated by this process. For such integrative naturalists, understanding what our best scientific theories are telling us about the world and understanding how we go about entheorizing that world in the first place are not distinct challenges: both are part of the overarching and more fundamental challenge of trying to *simultaneously* understand both the world and our own place within it.

I expect many philosophers of science to find this conception of naturalism congenial. After all, those of us who think of what we do as broadly continuous with the sciences often do so *because* we think that there is only a single project of inquiry and that we are engaged in it jointly with scientists themselves. That project has plenty of moving parts, from developing new statistical methods and quantitative measures of confirmation to building better microscopes, and from proposing and testing novel empirical hypotheses to interpreting the historical record of scientific inquiry, but these are all parts of a single interconnected inquiry into both the world and our own place in it, including the place of activities like science itself within that single project of systematic inquiry.¹

Indeed, this conception of naturalism might even seem too ecumenical, insofar as many paradigmatically non-naturalistic philosophers will presumably be happy to sign on to *some* version of the idea that all of our efforts to understand the world and our own place in it are simply parts or aspects of a single, integrated inquiry. It is therefore worth noting explicitly that this integrative naturalism will reject the persistent suggestion that there are distinctively philosophical modes of inquiry or characteristically philosophical routes to knowledge that are categorically distinct from those offered by our ordinary practices of empirical and scientific investigation. The suggestion that there are such distinctively philosophical forms of inquiry (sometimes alleged to have a stronger claim on our credence than ordinary empirical methods of investigation) has played a prominent role across the history of the discipline, whether we think of Plato's prescriptions for retrieving our souls' memories of the Forms carried from before our births, Descartes' clear and distinct perception of his own ideas or concepts, or Kant's transcendental inquiry into the conditions of possible experience, to name just a few historically influential examples.

To be sure, those who defend such distinctively philosophical forms of inquiry have usually offered some rationale for thinking that they can deliver important and otherwise inaccessible truths concerning nature and/or our own place within it. Descartes, for example, thought he had decisive reasons for regarding his concepts as the creations of an omnipotent, omniscient, and benevolent God who would (therefore?) not deceive him so long as he was careful to reason from and about those concepts judiciously. And Kant's "Copernican revolution" is widely seen as insisting that much of the structure of the external world as we experience it is imposed by the fundamental concepts we use to transform the raw materials available to us in sensation into genuine experience of such a world in the first place, and thus that the investigation of those concepts reveals how the world must be for us.²

1 This proposal closely parallels the much more detailed form of philosophical naturalism influentially defended by my colleague Penelope Maddy (see esp. 2007), and I am deeply indebted to her for many insights and useful conversations on this subject. Despite our agreement concerning the evidence to which an appealing form of naturalism must be responsive, however, she would sharply dissent from many of the conclusions I will go on to draw from that evidence.

2 There are, of course, other profound differences between Descartes and Kant, but the relevant similarity is that each is at least sometimes willing to set all empirical investigation aside and try to reach substantive conclusions about ourselves and our place in the world on the basis of some categorically distinct form of inquiry.

The integrative naturalist I have described will, of course, refuse to simply *presume* that there are any such independent routes to knowledge in advance of inquiry, and the presently available evidence will lead her to regard existing rationales for such distinctively philosophical forms of inquiry as manifestly unpromising. It is worth recalling that Kant's own case for the permanence of the fundamental Categories, for example, rested in large part on the fact that the logical apparatus from which they were generated had not changed substantially since the time of Aristotle – but Kant made this argument only a hundred years before Frege radically transformed the logic whose presumed permanence was the ground on which Kant thought he could identify such metaphysical fixed points in the ways we are capable of experiencing and/or cognizing the world. Though there may well be obligate features of the ways human beings are able to experience or conceptualize the world, empirical psychology now seems a far better bet for identifying whatever we can know about them than any distinctively philosophical form of transcendental inquiry. Of course, contemporary philosophers sometimes suggest that conceptual and/or linguistic analysis can ground such independent inquiry because it is only by satisfying our concepts and/or the meanings we assign to the associated terms that something can ultimately *count* as a “material object” or an “injustice.” But such conceptual and/or linguistic analysis seems considerably better suited to probing how we (presently) think and talk about the world than to providing substantive information about how the world itself is or must be (even how it is or must be *for us*). Though the specific intuitions, concepts, and linguistic judgments we find ourselves with at any particular time do, of course, themselves constitute a sort of empirical data, the balance of the evidence we have concerning the sources and plasticity of such intuitions, concepts, and judgments offers little reason to suppose that they offer distinctively philosophical routes to substantive knowledge of ourselves and the world around us that are independent of our broader empirical and scientific inquiries.

This integrative conception of philosophical naturalism clearly owes an enormous debt to Quine: at its heart are such central Quinean insights as the fact that there is only a single, integrated project of inquiry into the world and our place within it, that this inquiry must seek to improve the “inherited world theory as a going concern” from within without automatically privileging some reliable sources of evidence over others, and that it will neither presuppose nor conclude at present that there are any distinctively philosophical routes to substantive knowledge about ourselves or the world. If all this is right, however, Quine himself stumbles badly when he famously suggests that epistemology should simply be assimilated to or replaced with empirical psychology:

Naturalism does not repudiate epistemology, but assimilates it to empirical psychology. Science itself tells us that our information about the world is limited to irritations of our surfaces, and then the epistemological question is in turn a question within science: the question how we human animals can have managed to arrive at science from such limited information. (1975/1981, 72)

This Quinean impulse has been criticized in a wide variety of ways, but the most fundamental mistake here *concerning naturalism* would seem to be the suggestion that evidence concerning how sensory irritations generate scientific beliefs will be all the evidence that ultimately matters to epistemology. The central concern of epistemology is when and how we are *justified, warranted, or well-advised* to hold particular beliefs or sorts of beliefs in particular sorts of epistemic circumstances (and why), and the thoroughgoing integrative

naturalism I have described will see much more than just the evidence supporting the claims of our best psychological theories concerning the details of the causal chains between the “irritations of our surfaces” and scientific beliefs themselves as relevant to making that determination.³ More specifically, it will insist that such judgments must be responsive to evidence from the history of science, existing scientific practices, and much else that goes beyond the traditional purview of empirical psychology.

Epistemic Instrumentalism

Suppose now that we do replace the scientific naturalist’s exclusive focus on the impressive empirical achievements of our best scientific theories with a much wider conception of the evidence to which a satisfying integrated conception of the world and our own place within it must be responsive. Why will this make any difference to the view we take of the descriptions offered by those theories of how things stand in otherwise inaccessible domains of nature or challenge the scientific realist’s conviction that those descriptions must be at least probably and approximately true?

Critics of scientific realism have long noted that in most scientific fields contemporary theories are simply the latest in a succession of fundamentally distinct theoretical alternatives, many of which enjoyed impressive empirical successes of at least the same broad kinds as those enjoyed by their contemporary successors. Some critics (e.g., Laudan 1981) argue that this long history of empirically successful scientific theories now regarded as profoundly mistaken in their fundamental claims about the constitution of nature (including such famous examples as Newton’s mechanics, the phlogiston theory of chemistry, the caloric theory of heat, the wave theory of light and electromagnetism, Weismann’s theory of the germ-plasm, and many more besides) establishes that even the most dramatic forms of empirical success are simply not a reliable indicator that a scientific theory is probably and/or approximately true. Others (e.g., Stanford 2006) suggest that our repeated failures to even conceive of the fundamentally distinct and even more empirically impressive successors that would ultimately come to replace such theories shows that as inquirers we are not proficient at conceiving of the full range of theoretical alternatives that are well confirmed by a given body of empirical data, and thus that there are very likely alternatives to contemporary scientific theories that remain unconceived by us despite the fact that they are also reasonably well confirmed by the evidence we now possess. It is to this sort of historical evidence concerning the reach of our inquiry and our own capacities as inquirers that the integrative sort of naturalist I have described will insist that our conception of the world and our own place within it must be no less responsive than it is to the dramatic empirical achievements that convince scientific realists that successful contemporary scientific theories must be at least probably and/or approximately true.

I suggest, then, that even if such a naturalist begins her inquiry from within our “inherited world theory as a going concern,” she will ultimately come to see the *full range* of evidence

3 Quine himself denies that naturalizing epistemology in this way sacrifices its normative dimension (1981, 181; 1990, 19; 1995, 49; see also Maddy 2007, § I.6 for useful discussion). But what matters most here is why the integrative naturalist will resist the sort of assimilation Quine *appears* to endorse in this and other passages: the fundamental concern is that Quine’s own naturalism takes too narrow a view of the range of relevant evidence to which our epistemological judgments should be sensitive. I will go on to say more in later sections about how expanding the range of such evidence contributes to recovering the normative dimension of epistemology.

informing our integrated conception of the world and our own place within it as indicating that we lack sufficient warrant for regarding some of even the most successful contemporary scientific theories as probably and/or approximately true. She will resist the scientific realist's conception of such theories, despite the fact that she cannot propose any alternatives to them that she thinks have better or even equally good claims to represent the truth concerning the relevant otherwise-inaccessible domains of nature, and even if she cannot now say anything about what she ultimately expects their still more empirically impressive and fundamentally distinct successors to look like. But having reached that conclusion, what view will our integrative naturalist take of such theories instead? Where and how will they fit into the inherited world theory that she seeks to improve from within?

The most historically influential alternative to the realist's view that our best scientific theories are at least probably and/or approximately true has been the competing proposal that such theories are instead "mere instruments," which is to say useful cognitive and conceptual tools that we use to navigate our practical predicaments and to predict and intervene in the world around us as successfully as possible. On such a view, a scientific theory is no less a tool for being conceptual in character, or a "way of thinking" about some otherwise inaccessible domain of nature. Hilary Putnam famously pointed out (in a different context) that there are tools like a hammer, which require only one person to use, and tools like a steamship, whose effective use requires the coordinated efforts of many people. Here we might simply add that many of the tools we use, like mnemonic devices, work schedules, and recipes, are not physical objects at all.

In the middle decades of the 20th century, many defenders of instrumentalism famously embraced *semantic* versions of this proposal, variously holding that the apparently theoretical claims of science are really just claims about observable phenomena in disguise, that such claims do not actually assert anything at all but instead simply licensed inferences from one observable state of affairs to another, or that such claims are in principle eliminable from scientific discourse altogether. For good reason, each of these semantic proposals is now widely regarded as facing insurmountable obstacles (cf. Stanford 2005). But the fundamental idea that at least some of even our best scientific theories should be viewed merely as conceptual tools or instruments for guiding our practical engagement with the world in productive and instructive ways, rather than as accurate descriptions of how things stand in otherwise inaccessible domains of nature, depends *in no way* on the demand for a special semantic interpretation or translation of scientific claims. We might instead embrace an *epistemic* version of such instrumentalism, on which we see scientific theories as *saying* just what they seem to say about what the world is made up of and how it works, but nonetheless deny that we have sufficient justification for *believing* those claims when they are so understood. Such epistemic instrumentalism holds that scientific theories frequently can and do help us successfully navigate the world in productive and systematic ways, even when the claims they make about otherwise inaccessible domains of nature are not even approximately true. Accordingly, the epistemic instrumentalist thinks that we can and should *use* those theories to guide our pragmatic engagement with the world without *believing* what they say about how things stand in otherwise inaccessible domains of nature.

Such epistemic versions of instrumentalism face a persistent and powerful line of criticism, however, to the effect that there is no room for a genuine difference between making "merely" instrumental use of a scientific theory across the *full* range of diverse contexts in which such theories are put to work and simply believing such a theory to be true. Howard Stein argues, for example, that once we recognize that the instrumental uses of a theory

include not only calculating experimental outcomes but also adequately representing phenomena more generally and serving as “resources for inquiry; especially as sources of clues in...the search for good hypotheses” (1989, 52), there simply remains no room for a difference that *makes* any difference between such a thoroughgoing instrumentalism and a suitably modest form of scientific realism that gives up any claim to be telling us about how things stand independent of the ways it is possible for us to experience and/or conceive of them.

Simon Blackburn offers a similarly motivated but more detailed criticism of Bas van Fraassen’s influential “constructive empiricist” form of epistemic instrumentalism, which famously holds that instead of believing even our best scientific theories, we should merely “accept” them as “empirically adequate,” meaning we should believe all and only what they say about observable phenomena and remain agnostic about their further claims regarding the unobservable. But as Blackburn notes, the constructive empiricist nonetheless recommends that we “immerse” ourselves in our theories in order to take full advantage of them:

The constructive empiricist is of course entirely in favor of scientific theorizing. It is the essential method of reducing phenomena to order, producing fertile models, and doing all the things that science does. So we are counseled to *immerse* ourselves in successful theory...Immersion will include acceptance as empirically adequate, but it includes other things as well. In particular it includes having one’s dispositions and strategies of exploration, one’s space of what it is easy to foresee and what difficult, all shaped by the concepts of the theory. It is learning to speak the theory as a native language, and using it to structure one’s perceptions and expectations. It is the possession of habits of entry into the theoretical vocabulary, of manipulations of its sentences in making inferences, and of exiting to empirical prediction and control. Van Fraassen is quite explicit that all of this is absolutely legitimate, and indeed that the enormous empirical adequacy of science is an excellent argument for learning its language like a native... Immersion, then, is belief in empirical adequacy plus what we can call being “functionally organized” in terms of a theory. (2002, 117–119; emphasis in original)

But once we are “immersed in” or “animated by” our best scientific theories in order to take full instrumental advantage of them, Blackburn suggests, we have lost any room for a distinction between regarding our theories merely as useful cognitive instruments and simply believing them to be true:

The problem is that there is simply no difference between, for example, on the one hand being animated by the kinetic theory of gases, confidently expecting events to fall out in the light of its predictions, using it as a point of reference in predicting and controlling the future, and on the other hand believing that gases are composed of moving molecules. There is no difference between being animated by a theory according to which there once existed living trilobites and believing that there once existed living trilobites...What can we do but disdain the fake modesty: “I don’t really believe in trilobites; it is just that I structure all my thoughts about the fossil record by accepting that they existed”? (2002, 127–128)

Like Stein, Blackburn suggests that recognizing the full range of instrumental uses of a theory seems to eliminate any room we might have to distinguish “merely” instrumental reliance from belief itself. Thus, as we progressively sophisticate our grasp of the various ways in which we make use of our best scientific theories or put those theories to work, it becomes harder and harder to see how such systematic reliance on or use of a scientific

theory amounts to anything less than simply believing that same theory. The epistemic instrumentalism embraced by our integrative philosophical naturalist thus threatens to simply collapse into scientific realism itself!

We can see that something has gone badly wrong in this line of argument by noting that the idea of making systematic instrumental use of a scientific theory we do not believe to be even approximately true turns out to be one that the scientific realist needs no less than the epistemic instrumentalist – after all, this is precisely the view that the realist herself takes of a theory like Newtonian mechanics! That is, the scientific realist believes the claims of Newton's mechanics about the fundamental constitution and operation of nature to be profoundly mistaken: gravitational motion is not caused by forces exerted by massive bodies on one another, space and time are not absolute, and so on. But she knows perfectly well how to use Newton's mechanics rather than Einstein's to send rockets to the moon and in a wide array of other engineering applications in which we know that any differences in their empirical predictions are small enough to be safely disregarded. Thus, we might begin by asking what the *realist* means when *she* says that we make effective instrumental use of Newtonian mechanics to solve practical problems and challenges but we do not believe its fundamental claims about the constitution of nature to be even probably or approximately true.

The most natural way to describe the realist's attitude toward Newtonian mechanics would seem to be that though she does not believe the theory's claims about the fundamental constitution and operation of nature to be true, she nonetheless knows how to apply the theory as a true believer would to entities, events, and natural phenomena she thinks she can correctly characterize in a variety of other ways. That is, the realist accepts the existence of entities like rockets, rattlesnakes, and RNA molecules on the evidence of her senses and/or the descriptions offered of them by theories that she does strictly and literally believe, and she knows both how to represent such entities within a Newtonian framework and how to derive consequences and empirical predictions concerning their behavior when they are so represented. She knows how a Newtonian *would* characterize (in terms of forces, masses, etc.) the rockets, rattlesnakes, and RNA molecules familiar from her own experience and/or other theories that she takes to be probably and/or approximately true, and she can use those characterizations together with the dynamical apparatus of the theory itself to derive predictions, explanations, and recipes for intervention just as a Newtonian would concerning entities, events, and phenomena as she understands or conceives of them independently of Newtonian mechanics (or indeed any other theory toward which she herself adopts a similarly instrumentalist attitude).

Perhaps surprisingly, however, *this* conception of instrumentalism is equally available to the epistemological instrumentalist herself, so long as she recognizes at least some routes of epistemic access to entities, events, and phenomena in the world toward which she does not extend an instrumentalist attitude. Many classical instrumentalists have thought that objects and events *as they are given to us in sense perception* are those with respect to which scientific theories allow us to predict and intervene successfully. But the instrumentalist is by no means limited to the evidence of her senses in conceiving of the natural phenomena with respect to which she thinks a given scientific theory offers an effective instrumental tool for prediction, intervention, further investigation, or anything else: she is in a position to take full advantage of *whatever* routes of epistemic access she has to entities, events, and natural phenomena besides the descriptions of them offered in any scientific theories toward which she takes an instrumentalist attitude, no matter what those routes turn out to be. Indeed, this conception is available to the epistemic instrumentalist even if she embraces

Quine's further suggestion that *all* of our knowledge is theoretical in character and thus that even our understanding of the most familiar entities and events is given to us by a hypothesis of "the bodies of common sense" (1976, 250). So long as she does not adopt an instrumentalist attitude toward *this particular hypothesis*, the descriptions of entities and events it offers can serve perfectly well as the conception of those entities and events regarding which theories that are "mere instruments" will allow us to predict and intervene effectively. The instrumentalist, just like the realist, will simply be faced with a decision concerning *which* theories are those toward which she will adopt a realist attitude, and why all and only those.

This realization suggests that it was a mistake all along to think of the difference between the scientific realist and her instrumentalist opponent as a difference in the respective attitudes each takes toward "theories" or "theoretical knowledge" as such. Both realists and instrumentalists will take instrumentalist attitudes toward some theories (e.g., Newtonian mechanics) and realist attitudes toward others (e.g., the Quinean hypothesis of the bodies of common sense), and the crucial remaining difference between them is simply that thinkers who embrace the instrumentalist label are those who are prepared to treat a much *wider range* of theories merely as useful conceptual instruments for guiding our pragmatic engagement with and further investigation of nature than are their realist counterparts. In particular, the instrumentalist, but not the realist, will be prepared to adopt this attitude even toward theories that are clearly more empirically successful than any competitor we have yet managed to develop, even when she cannot yet say anything at all about what she expects their still more empirically successful and pragmatically powerful successors to look like.

Of course, scientific theories sometimes posit the existence of entities or events to which we have no route of epistemic access at all beyond the causal, dynamical, or other role ascribed to them by that very theory. Because contemporary particle physics does not allow quarks to be isolated, for example, it posits the existence of "gluons" to bind quarks within a proton, while contemporary cosmology posits the existence of dark energy to account for the accelerating expansion of the universe. In both cases, however, the ground for believing in these entities is simply the demand of the relevant theory that *something* play the role that each is hypothesized to fulfill. Someone taking an instrumentalist attitude toward contemporary particle physics (or cosmology) will therefore not believe in even the actual existence of gluons (or dark energy), much less any part of the theoretical description of them offered by the theory, but will *simply* make use of such descriptions and claims in seeking to successfully predict, intervene, further investigate, and so on with respect to *other* entities, events, and phenomena to which she does see herself as having independent epistemic access. That is, she will regard gluons or dark energy *simply* as instrumentally useful fictions, just as the realist so regards Newtonian gravitational forces or absolute space and time.

As this comparison suggests, plenty of room remains for reasonable disagreement concerning how far realist or instrumentalist attitudes should be extended to particular scientific theories. Though quantum mechanics is widely regarded as the most empirically successful theory in the history of science, for example, even many realists hesitate to embrace the idea that its descriptive claims about otherwise inaccessible domains of nature are at least probably and/or approximately true, in part because it seems so difficult to understand those claims *as* a description of what actually occurs in the physical world. In addition, quantum mechanics is straightforwardly inconsistent with other wildly successful physical theories in ways that it is not easy to see how we can avoid or overcome without making fundamental changes to one or more of those theories (cf. Barrett 2003).

More generally, there is ongoing disagreement concerning the right morals to draw from the historical record of scientific inquiry itself. Much controversy has centered on whether there are particular *characteristics* of some theories (such as their “maturity” or their ability to make successful predictions of novel phenomena) that should lead us to expect them to avoid the fate of their historical predecessors, and/or whether the historical record reveals particular *parts or aspects* of (sufficiently) successful scientific theories (such as their claims about the “structure” of nature, the entities they posit, the “working” posits to which they ascribe substantive causal roles, or those aspects regarded as best-confirmed by the relevant scientific communities themselves) to be invariably or at least reliably preserved by their historical successors, justifying our confidence that such theories, or at least their relevant parts or aspects, will persist throughout the course of all further inquiry. I argue elsewhere (Stanford 2006, chs. 6 and 7) that these rearguard maneuvers are unconvincing: the criteria used to pick out the relevant characteristics, parts, or aspects of theories either run afoul of notable counterexamples from the historical record or are framed in sufficiently vague and indeterminate ways as to provide little *prospective* guidance in picking out the relevant theories and/or parts or aspects of theories in which we can now safely invest our credence.

We need not resolve such disputes, however, to achieve the aim of the present chapter, which is simply to show how resistance to scientific realism *could* be combined with philosophical naturalism in an attractive and compelling way. We have already seen how widening the range of evidence to which our naturalistic inquiry must be responsive so as to include empirical evidence from the historical record and elsewhere might motivate the combination of an integrative philosophical naturalism with a thoroughgoing epistemic instrumentalism. Moreover, even if we ultimately decide that the real lesson of such historical evidence is simply that we should limit our belief to successful scientific theories that share some further characteristic (e.g., “maturity,” the ability to make novel predictions), or simply to some parts or aspects of such theories (e.g., their “structure” or their “working” posits), this would already suffice to illustrate how the sort of integrative naturalism I have proposed motivates a view of our epistemic entitlements that falls well short of traditional realist commitments.

Philosophy of Science as Constructive Interpretation

Suppose now that we have (1) abandoned the idea that there are distinctively philosophical methods of investigation that are independent from those of ordinary empirical inquiry, (2) accepted that realism and instrumentalism are not best understood as distinctive philosophical views of theories or theoretical knowledge as such, but instead as cognitive attitudes that realists and instrumentalists alike take toward some particular theories or bodies of information about the world (e.g., Newtonian mechanics) and not others (e.g., Quine’s hypothesis of the bodies of common sense), and (3) recognized that thoroughgoing naturalists will see the correct attitude to take toward any particular scientific theory as depending not only upon individualized consideration of the empirical achievements and further characteristics of that theory, but also on the historical record of our successive efforts to entheorize the world more generally, and perhaps further empirical evidence concerning ourselves as theorizers besides. By this point, it might seem that philosophical activity has simply been excised from, rather than integrated into, the resulting naturalistic conception of the world and our own place in it. I think the idea that embracing integrative

naturalism leaves philosophy or philosophers with nothing to contribute to the resulting project of inquiry is profoundly misguided, and this is neither simply because scientists themselves do not typically attend to what I have been suggesting is the full range of evidence relevant to deciding what cognitive attitude to adopt toward any particular theory, nor simply because there are many important questions that belong to no single scientific specialty (cf. Maddy 2007, 115ff.). I therefore wish to conclude by describing what seems to me a crucial aspect of an integrated naturalistic conception of the world and our own place within it to which I think philosophers of science are particularly well-positioned to contribute, and indeed have been contributing for quite some time, whether they realize it or not.

Let me begin with a long-standing tension in the self-conception of philosophers of science concerning whether ours is fundamentally a prescriptive or a descriptive enterprise. Many philosophers of science, I suspect, find themselves in an uneasy tension between competing impulses on this question, thinking that the job of the philosophy of science is to understand and describe the workings of science as we actually find it, but also that there is an important normative dimension to this enterprise and that it is sometimes appropriate for philosophers of science to draw conclusions from their own work concerning how scientific inquiry *should be* conducted in general or *should* proceed in a particular case. I think there is room within integrative naturalism to reconcile these apparently competing impulses in an attractive way.

Both impulses have, I think, been deservedly influential across the history of the philosophy of science itself, but each has also sometimes been *too* influential. When the prescriptive impulse has exerted an outsized influence, it has produced a variety of philosophy of science we might caricature as the “Epistemic Police.” The Epistemic Police suppose that the role of philosophers is to justify and protect the epistemic authority of science while correcting and/or improving the activity of scientists by telling them what the demands of rationality or of *genuinely* scientific inquiry are and when or how those demands have or have not successfully been met. This caricature captures some of the flavor of Karl Popper’s falsificationism or Imre Lakatos’ methodology of scientific research programs, as well as some improvident arguments in support of contemporary formal theories of confirmation, and of course some logical positivist and logical empiricist philosophers of science were sufficiently confident in this role to be willing to replace the actual history of science with “rational reconstructions” of that history in order to more clearly draw the appropriate normative morals from it.

Perhaps it was inevitable that this conception of the relationship between philosophy and science would eventually produce a backlash, but for whatever reason, more recent philosophy of science has tended to take a far more deferential attitude toward the details of actual scientific practice and the historical record of how scientific inquiry itself has actually proceeded: most contemporary philosophy of science simply does not take itself to be in the business of telling scientists how to do science. This otherwise welcome humility has, however, sometimes threatened to generate an almost purely descriptive form of inquiry, whose caricature might be called the “Anthropology of Science.” Anthropologists of Science reject the presumptuous penchant of the Epistemic Police for lecturing scientists concerning how to conform their practices more closely to the demands of rationality or the requirements of genuinely scientific inquiry, and seek instead to describe the scientific enterprise as we actually find it, as well as to draw instructive morals from the resulting more detailed and accurate conception of that enterprise (or some part of it) regarding how science is able to do what it does. Much of the work of the Anthropologists of Science consists of careful

field reports concerning the actual theoretical commitments, practices, views, and/or thinking of various groups of scientists, often preceded or followed by vague and hopeful gestures toward traditional philosophical questions about the nature of scientific inquiry that might be illuminated thereby.

I think there is an available conception of the philosophy of science that recovers the attractions of the Epistemic Police and the Anthropologists of Science while allowing us to dispense with them both. I will begin describing it by borrowing the idea of a constructive interpretation from the legal philosophy of Ronald Dworkin (1986, esp. ch. 2), who himself relies to some extent on the earlier work of Stanley Cavell. Dworkin introduces constructive interpretation by contrasting it with ordinary conversational interpretation, in which we seek to recover the actual intention or mental state of some speaker who has produced an utterance. Thus, if you say, "Would you like to go to the movies this weekend?" or "Are you planning to pay for that?", my goal in interpreting you will be to identify the mental state that actually led you to produce that utterance, and the right interpretation of the utterance will be one that correctly reconstructs that mental state. Dworkin argues, however, that such conversational interpretation is not the best way to understand how we interpret phenomena like works of art or social practices (and ultimately laws), which are better seen as targets of constructive interpretation. In such constructive interpretation, the interpreter's objective is not to reproduce the actual intention or mental state of any particular speaker or social actor, but instead to find the interpretation of the work of art or social practice that best fits the object of interpretation, while simultaneously seeking to make it, in Dworkin's repeated phrase, "the best it can be" (1986, 53). That is, constructive interpretations seek to simultaneously maximize two independent desiderata, which Dworkin calls "fit" and "acceptability": a constructive interpretation is better insofar as it better rationalizes or motivates more elements of the work of art or social practice as a whole, but also insofar as the rationale, point, purpose, aim, or value it ascribes to that work of art or social practice renders it more independently appealing, interesting, useful, attractive, significant, mind-blowing, and so on (by the interpreter's own lights and/or those of any community whom she hopes to convince to adopt or share her interpretation). For the constructive interpreter, the *right* reading(s) of *All My Sons* or "All Along the Watchtower" is simply that (or those) which manages to maximize both fit to the object of interpretation and independent attractiveness to the greatest possible degree (and, of course, the appropriate tradeoff of these virtues in any particular case is itself an interpretive decision).

Whether or not such constructive interpretation is the best option for interpreting works of art or the law, Dworkin's suggestion that it is the right approach to interpreting social practices offers a useful tool for reconceiving the philosophy of science without the Epistemic Police or the Anthropologists of Science. Philosophers of science can reconcile their competing prescriptive and descriptive impulses in an attractive way, I suggest, by seeing themselves as offering candidate constructive interpretations of the social practice of science itself. Such constructive interpretations invariably begin with some preinterpretive identification of the object of interpretation: the poem on the page, the play on the stage, or the social practice (in this case, scientific practice) as we find it. In seeking to maximize fit, a good constructive interpretation of science will seek to motivate, explain, and rationalize as much of what scientists themselves *actually* do (including much of what they say) as possible. But constructive interpretations also retain the freedom to ultimately exclude some parts of the preinterpretively identified object of interpretation as mistakes,

because fit is not all that matters. At some cost to the fit of their interpretations to the practice they are interpreting, constructive interpreters of science are free to reject parts or aspects of the *actual* practice of science as inconsistent with the most attractive constructive interpretation of most of or the rest of that practice; that is, with the constructive interpretation that makes the social practice of science the best it can be. On this view, rejecting a given part or aspect of scientific practice as mistaken or even just suboptimal is similar to rejecting a part of a play or a book, or a line in a song, as contributing nothing toward or even detracting from what is otherwise the most attractive constructive interpretation of the play or song's point, purpose, aim, value, rationale, or significance, and therefore as not actually being part of the play or book or song (as it is best interpreted) in the first place. In extreme cases, a sufficiently attractive constructive interpretation may lead us to reject a substantial part of the preinterpretively identified object of interpretation, as when some critics reject the disastrous final chapters of *The Adventures of Huckleberry Finn* as inconsistent with the best constructive interpretation of the book as a whole, even at considerable cost to the fit of that constructive interpretation with the full text of the published novel itself. For such interpreters, these final chapters of the novel are a *mistake*, in just the same way that even widespread scientific practices (of fraud, or dogmatic adherence to a theory, or poor statistical methodology) can be regarded as normatively mistaken insofar as they are, though familiar features of the practice of science as we preinterpretively identify it, nonetheless inconsistent with what we see as the independently most attractive conception of the rationale, point, purpose, aim, value, or significance of that practice.

This view of the matter allows us to understand the appeal of the Epistemic Police and the Anthropologists of Science, while dispensing with both. It sees each of these traditions as excessively fixated on just one of the evaluative dimensions of constructive interpretation (acceptability and fit, respectively) and insufficiently attentive to the other. Accordingly, it rationalizes the broad lines of criticism that rightly make each of these caricatures seem unappealing: the complaint that the Epistemic Police are seeking to characterize a philosophical fantasy rather than the actual scientific enterprise and the complaint that the Anthropologists of Science settle too quickly for mere description and prematurely abandon any substantive normative appraisal of whether and how features or elements of the scientific enterprise as we (preinterpretively) find it should be revised, rejected, resisted, or reformed as that practice is extended into the future.

Beyond respecting both the prescriptive and descriptive impulses motivating the Epistemic Police and the Anthropologists of Science, however, I suggest that seeing the philosophy of science as engaged in attempts to constructively interpret science or parts of science promises to rationalize much of the rest of what philosophers of science themselves say and do, including their efforts to develop and defend philosophical "theories" of explanation, confirmation, and the like. Bayesian confirmation theory, for example, does not seem a promising description of the process by which most scientists actually make confirmational judgments, nor do most Bayesians seem to be suggesting that we simply ignore scientists' own confirmational judgments and adopt the Bayesian apparatus as a way of making such judgments *instead*. Bayesians are best seen, I think, as offering constructive interpretations that seek to rationalize as much as possible of existing scientific judgments of comparative confirmation, of evidential virtues, and so on (though retaining the freedom to reject any particular scientific judgments or instance of that practice as mistaken), while also offering the most independently attractive conception of that practice (in part, for

familiar reasons concerning Dutch books and the like⁴). I will here limit myself, however, to detailed discussion of just one further example, in which I think the attractions of seeing the philosophy of science as aimed at constructive interpretation in this way are especially conspicuous.

Bas van Fraassen famously claims that the fundamental conflict between scientific realism and his constructive empiricist alternative concerns the *aims* of science. On a realist view, “[s]cience aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true” (van Fraassen 1980, 8). By contrast, van Fraassen’s constructive empiricist position holds that “[s]cience aims to give us theories which are empirically adequate, and acceptance of a theory involves as belief only that it is empirically adequate” (1980, 12), where empirical adequacy means simply that the theory’s claims about observable phenomena are true. But many philosophers of science, including perhaps most influentially Gideon Rosen (1994), have expressed puzzlement regarding just what van Fraassen is asserting. Van Fraassen is quite clear that he does not mean to suggest that the pursuit of empirically adequate theories is the aim of all or most individual scientists (an interpretation that Rosen rightly rejects as turning van Fraassen’s book into “a work of inept sociology” (1994, 147)), nor is he arguing that scientists *should* adopt the constructive empiricist’s aim, even if they presently do not.⁵ Indeed, van Fraassen explicitly distances the constructive empiricist’s perplexing claims about the aim of science from any view of what the aspirations of individual scientists are or should be with an analogy to chess, in which he says the aim of the game *itself* is checkmate, even if individual players are motivated by “fame, gold, or glory” (1980, 8). Rosen nevertheless insists that the manifest gap between constructive empiricist commitments and the (explicit or implicit) intentions of most individual scientists entails that constructive empiricism is false if taken “literally” (1994, 154) and is therefore instead “best seen as a fiction about science put forward not as true, but rather as adequate to the phenomena of scientific activity” (1994, 153).

Having proposed this fictionalist reading of constructive empiricism, however, Rosen insists that van Fraassen fails to provide any reason to adopt or even care about such a view that should appeal to anyone who does not *already* share his own commitments both to the rationality of scientific practice and to a broad form of empiricism which contends that experience is our only legitimate source of knowledge about the world. The most van Fraassen can hope to show, Rosen argues, is that constructive empiricism is the most stringent form of a broadly *empiricist* commitment to science, “which at the same time renders the phenomena of scientific activity fully intelligible and ratifiable” (1994, 162); he thinks van Fraassen says nothing that should encourage anyone who is not already and independently “committed both to the rationality of science and to empiricism in the broad sense” (1994, 163) to become a constructive empiricist in the first place.

But we need not and should not see van Fraassen as simply preaching to the choir of those antecedently devoted to both empiricism and the rationality of scientific practice.

4 The rational appeal of Bayesianism is sometimes thought to be captured in the recognition that any agent who does not assign and update probabilities in accordance with the Bayesian apparatus is open to accepting a “Dutch book”: a combination of wagers on which she will be guaranteed (not just likely) to lose money.

5 It is perhaps revealing that the two interpretations of van Fraassen Rosen considers and rejects as inadequate before proposing his own are that he means simply to describe the intentional state of most scientists (the Anthropology of Science) and that he means to prescribe what those intentional states should be (the Epistemology Police).

If we instead see his sermon as an attempt to constructively interpret science as a social practice, it is the need to generate a *convincing* interpretation of this sort that leads him to seek to rationalize as much of that practice as possible, rather than a prior and undefended commitment to the rationality of science itself. Similarly, we need not follow Rosen in seeing the further virtues van Fraassen claims for constructive empiricism simply as attempts to make the attractions of the view evident to anyone who already happens to share his own broad commitment to empiricism. Instead, when van Fraassen claims that constructive empiricism exposes us to fewer epistemic risks than realism (1980, 69; 1985, 246, 280–281), that it eliminates the need to see scientists as embracing beliefs about nature on the strength of comparative virtues that are pragmatic rather than epistemic (i.e., truth-tracking) in character (1980, 87–96; 1985, 264–265, 280–281), that it disdains the realist’s “empty strutting and posturing [and] display of courage not under fire...in embracing additional beliefs which will *ex hypothesi* never brave a more severe test” (1985, 255), and the like, we can see such claims as attempts to display the *independent* appeal of embracing a constructive empiricist interpretation of science itself. After all, one need not be an empiricist to want to avoid unnecessary or pointless epistemic risks, to want to limit the range of influences determining our scientific beliefs to characteristics of those beliefs that plausibly track the truth, or to find empty strutting and posturing and grandiose expressions of “courage not under fire” unappealing. Though van Fraassen clearly thinks that constructive empiricism *also* more effectively rationalizes more of the practice of science itself than does the realist alternative (1980, 73–87; see also 1994, 191), we may nonetheless see claims such as “constructive empiricism...makes better sense of science, and of scientific activity, than realism does” (1980, 73) as asserting the superiority of constructive empiricism to scientific realism on grounds of independent acceptability, as well as fit. That is, I think we can and should see constructive empiricism as offering what van Fraassen (rightly or wrongly!⁶) takes to be the most appealing conception of the point, purpose, value, or aim of the social practice of science (acceptability) that also best explains, rationalizes, and defends as much of the actual activity of science as possible (fit).⁷ That is, it offers a characterization of the aim of science *itself* (as opposed to the aims of all or any of its individual practitioners) that is simultaneously descriptive and normative in just the way that constructive interpretations invariably seek to be.

This proposal should also make clear why constructive empiricism is neither seeking to describe nor even particularly concerned with what most individual scientists take the aim of science to be. One of the grounds on which Dworkin argues that interpreters must construct social practices in the course of interpreting them is that there simply is no canonical participant or group of participants in such a practice whose individual intentions concerning the point, purpose, aim, or value of the practice itself *could* automatically or dispositively settle the issue (1986, ch. 2). Thus, any individual’s conception of the practice in which she herself participates simply represents a candidate constructive interpretation like

6 That is, while I think van Fraassen is best seen as attempting to constructively interpret science, I think a superior constructive interpretation is offered by the epistemic instrumentalism I sketched in the previous section. Moreover, I think van Fraassen is profoundly mistaken to see the need for such constructive interpretation as arising only when we step back from scientists’ own “immersion” in scientific investigation to consider the epistemic status of its products (1980, 80–82). Instead, I will go on to suggest that such constructive interpretation is (and always has been) a critical part of conducting scientific inquiry itself.

7 Indeed, I am prepared to defend this as the most attractive constructive interpretation of *constructive empiricism*, no matter what view of the matter van Fraassen himself takes.

any other – and we are more likely to think it the right one as it better maximizes both fit and acceptability, making that practice “the best it can be,” while simultaneously motivating or rationalizing as much of the practice as possible. Likewise, in the scientific case, there is no canonical person, group, or institution whose views on the aim of science could plausibly serve to simply settle the question authoritatively in the way that the actual intentions of a speaker do in the case of conversational interpretation. Thus, though scientists speak with especially detailed knowledge of their own practices, their own views about the point, purpose, aim, or goal of science itself do not automatically have any special standing or claim to authority. When scientists express such views, they themselves are acting as constructive interpreters of science: they propose views of what science *itself* has been aiming to do all along and of how to best extend that practice going forward in light of that aim. Such proposals must be evaluated in precisely the same ways as those offered by philosophers of science, just as empirical hypotheses proposed by philosophers should be evaluated in just the same ways as scientists’ own empirical hypotheses.

These two forms of evaluation are not the same, however. For too long, too many philosophers of science have thought of themselves as engaged in the sort of hypothesis-testing with which they see scientists and scientific inquiry as primarily concerned, using details of the practice or history of science as the “data” for which any philosophical “hypothesis” about the nature of science or some part of science is responsible. But philosophy of science is not science, nor are its central activities best assimilated to hypothesis-testing. Perhaps the most important difference is that scientists are not free to normatively evaluate and reject their data in the same way that philosophers of science normatively evaluate the practices of science itself. That is, while scientists can and certainly do evaluate and reject particular data as suspect or unreliable, they are not free to reject what they take to be genuine empirical results as *normatively inappropriate* or as *mistakes* simply because they conflict with what is otherwise the most exciting, impressive, interesting, mind-blowing, or in some other way desirable theoretical account of a given natural domain that would account for the rest of the data we have regarding it. By contrast, philosophers of science are free to propose constructive interpretations of science as we (preinterpretively) find it that reject particular aspects or instances of that practice as normatively inappropriate, misguided, and/or inconsistent with the best conception of the point, purpose, goal, aim, or rationale of scientific practice as a whole. In fact, I have suggested that doing so is part of the very point of conducting philosophy of science in the first place.

To recognize this crucial difference between science and the philosophy of science is not, however, to see the latter as seeking to investigate the world or our own place within it by means of some distinctively philosophical form of inquiry that is somehow independent of ordinary empirical investigation itself. It is instead to suggest that forming a conception of *what we ourselves are doing* when we use those empirical methods to investigate the world and our own place within it is *already* part and parcel of any such integrated naturalistic inquiry. Indeed, at least since the origin of modern science in the earlier intellectual traditions of natural philosophy and natural history, articulating such conceptions and contesting their adequacy has been a central concern of scientists themselves. The constructive interpretation of science is thus something that scientists themselves already do (and have always done) in trying to situate their own practices within a broader vision of the world and our own place within it. I suggest that any satisfying conception of our “inherited world theory as a going concern” will have to include not only our best scientific theories but also our best constructive interpretation of the point, purpose, value, or significance of

constructing, exploring, and testing such theories, including what kind(s) of knowledge of the world we think they can and cannot give us under various kinds of epistemic circumstances.

Of course, the suggestion that philosophy of science is in the business of the constructive interpretation of science is neither purely a description of everything that philosophers of science presently do (still less what they already think of themselves as doing), nor purely a recommendation regarding what I think they ought to be doing instead, because philosophy of science is itself a social practice, and therefore itself stands in need of constructive interpretation, which is to say, a simultaneously descriptive and prescriptive conception of its aim, purpose, value, point, goal, or significance and of how we might accordingly continue that collective practice into the future. By its own lights, the constructive interpretation of *the philosophy of science* I have just proposed will have to compete with others on the grounds of both fit and acceptability, and I suggest that this interpretation not only best fits or rationalizes what philosophers of science say and do but also offers an independently attractive conception of what we have been up to all along. When it comes to interpreting ourselves, it's turtles all the way down.

Acknowledgments

Thanks to my LPS colleagues Jeff Barrett, Jim Weatherall, and especially Penelope Maddy for many useful conversations regarding this material. Thanks also to generous audiences at the Universidad Autónoma Metropolitana in Mexico City and at a conference held at Lingnan University in Hong Kong. I am grateful to the University of Pittsburgh's Center for the Philosophy of Science, where some work on this chapter was completed while I was Senior Fellow there, and to the Australian National University, where I received helpful feedback on it as a Visiting Fellow.

References

- Barrett, J.A. (2003). Are Our Best Physical Theories Probably and/or Approximately True? *Philosophy of Science* 70: 1206–1218.
- Blackburn, S. (2002). Realism: Deconstructing the Debate. *Ratio* 15: 111–133.
- Dworkin, R. (1986). *Law's Empire*. Cambridge, MA: Belknap.
- Ladyman, J. and Ross, D. (with Spurrett, D. and Collier, J.) (2007). *Every Thing Must Go. Metaphysics Naturalised*. Oxford: Oxford University Press.
- Laudan, L. (1981). A Confutation of Convergent Realism. *Philosophy of Science* 48: 19–49.
- Maddy, P. (2007). *Second Philosophy: A Naturalistic Method*. Oxford: Oxford University Press.
- Quine, W.V.O. (1951/1980). Two Dogmas of Empiricism. In *From a Logical Point of View*, 2nd edn. Cambridge, MA: Harvard University Press.
- Quine, W.V.O. (1975/1981). Five Milestones of Empiricism. In *Theories and Things*. Cambridge, MA: Harvard University Press.
- Quine, W.V.O. (1976). Posits and Reality. In *The Ways of Paradox and Other Essays*, 2nd edn. Cambridge, MA: Harvard University Press.
- Quine, W.V.O. (1981). *Theories and Things*. Cambridge, MA: Harvard University Press.
- Quine, W.V.O. (1990). *Pursuit of Truth*. Cambridge, MA: Harvard University Press.
- Quine, W.V.O. (1995). *From Stimulus to Science*. Cambridge, MA: Harvard University Press.

- Rosen, G. (1994). What is Constructive Empiricism? *Philosophical Studies* 74: 143–178.
- Rosenberg, A. (2014). Disenchanted Naturalism. In *Contemporary Philosophical Naturalism and its Implications*, edited by Bana Bashour and Hans D. Muller, pp. 17–36. New York: Routledge.
- Stanford, P.K. (2005). Instrumentalism. In *The Philosophy of Science: An Encyclopedia*, edited by Sahotra Sarkar and Jessica Pfeifer, pp. 400–405. New York: Routledge.
- Stanford, P.K. (2006). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. New York: Oxford University Press.
- Stein, H. (1989). Yes, But... Some Skeptical Remarks on Realism and Anti-Realism. *Dialectica* 43: 47–65.
- van Fraassen, B. (1980). *The Scientific Image*. Oxford: Clarendon.
- van Fraassen, B. (1985). Empiricism in the Philosophy of Science. In *Images of Science: Essays on Realism and Empiricism*, edited by Paul Churchland and Clifford Hooker, pp. 246–308. Chicago: University of Chicago Press.
- van Fraassen, B. (1994). Gideon Rosen on Constructive Empiricism. *Philosophical Studies* 74: 179–192.

UNCORRECTED PROOFS