

TRUTH AND THE SCIENCES

Anjan Chakravartty

Institute for the History and Philosophy of Science and Technology, and Department of Philosophy
University of Toronto

Abstract. Conceptions of truth in relation to the sciences vary extensively along two dimensions. The first concerns the *applicability* of the notion of truth to scientific knowledge, resulting in a number of contentions: that scientific knowledge is wholly, or in some important sense, independent of truth; that truth is, in some form, a proper aim of science; and that truths are, in fact, illuminated by means of scientific investigation. The second dimension concerns the particular *theory* of truth one might think applicable, and here one finds a variety of preferences, including: truth as coherence, especially suited to historicist and sociological approaches to science; truth as utility, described by pragmatist approaches to science; and truth via correspondence or truth-making, especially in the context of various forms of “realism” in connection with scientific knowledge. In this essay, I travel along the first dimension in the direction of increasing commitment to the applicability of truth, and in each case explore the second dimension: how different views of scientific knowledge appeal, explicitly or implicitly, to different theories of truth.

1. Introduction

What is the relationship of truth to the sciences? Ideally, as a prelude to answering this question, one might begin by specifying what, in this context, truth is generally thought to be. Unfortunately, however, though many philosophers of science believe that the concept of truth is important to a consideration of the practice and outcomes of scientific work, there is no consensus regarding *which* concept of truth, among the several philosophers have developed more generally, is most appropriate here. Indeed, there is no consensus even regarding the question of *whether* the concept is relevant to philosophical considerations of the sciences, and if so, how. In light of this diversity of opinion, perhaps not surprisingly, the task of getting to grips with the relationship of truth to the sciences is not entirely straightforward: it requires an examination of the different approaches to scientific investigation and knowledge that generate these different commitments. In this essay, I explore the many different ways in which concepts of truth have been brought to bear in connection with these different approaches.

The various attitudes towards truth found in the philosophy of science have correlates in a number of what one might view as philosophically pre-reflective attitudes exemplified in everyday discussions and debates about science. The different philosophical approaches towards the sciences I have alluded to may thus be thought of as more careful or rigorous philosophical extensions of these everyday attitudes. By way of introduction, then, it may be helpful to enumerate these pre-reflective attitudes before digging into their philosophical counterparts. On the one hand,

there are more or less “optimistic” attitudes regarding the connection between truth and the sciences – those that view the sciences as producing truths concerning the natural and social worlds in which we live, or as arbiters with respect to such truths, or at the very least, our best bets for truth production and arbitration. On the other hand, there are more or less “pessimistic” attitudes, which regard the connection as relatively weak or in some cases, illusory. Let us consider this range of possibilities, briefly, in turn.

Almost everyone, I suspect, would agree that the sciences are in the knowledge production business – scientific work produces theories, models, predictions, and explanations which are offered, *prima facie*, as candidates for knowledge. The optimism and pessimism I have gestured towards concern the extent to which these candidates can or should be thought of in terms of truth. Optimists commonly hold that our best scientific theories and models can be thought of as yielding truths regarding their subject matter, or as furnishing descriptions that are, over time, approaching the truth (in a sense to be specified, as we shall see in sections 3 and 5). This is the implicit attitude exemplified in everyday contexts in which scientists are presented as authorities regarding everything from the common cold to genetic manipulation to climate change. This generally positive view of the connection between science and truth has its most pronounced philosophical extension in varieties of “realism” with respect to scientific knowledge. Realists of different sorts typically defend the proposition that our best scientific theories provide true or approximately true descriptions of the world, but in different ways: some defend realism rather generally, and others only with respect to certain classes of ontological or other claims which they assess as having greater epistemic warrant than others.

One need not be a realist, however, in order to adopt a positive attitude towards the connection between science and truth. A certain degree of pessimism with respect to truth is compatible with a degree of optimism. There are those, for example, who are somewhat suspicious of the often definitive- and precise-sounding claims of scientists – particularly in connection with things very far removed from our abilities to see or otherwise register information in the absence of elaborate instrumentation. Examples of such things may include the very small, such as viruses and subatomic particles, or the very spatiotemporally distant, such as events in the immediate aftermath of the Big Bang or those leading to the extinction of the dinosaurs. Some sceptics nevertheless hold that scientists are likely on to something, even if they are suspicious about their furthest reaching claims. This everyday attitude also has well developed philosophical extensions. A number of philosophical positions, including certain varieties of empiricism, regard the sciences as

producing truths, but of a restricted subset of scientific claims comprising descriptions of observable phenomena (as opposed to other, less directly detectable phenomena).

Of course, it is possible to be even more strongly sceptical. One might doubt that the sciences produce truths at all, or that while they might do so on occasion, there is little reason to think that they are particularly successful in this regard. Everyday scepticism about the veracity of scientific knowledge has its source in a multiplicity of motivations, from the epistemological to the religious, social, and political. In the philosophy of science, the motivations for this stronger sceptical attitude are invariably epistemological, but in some cases these views do have important social and political dimensions. Some who ascribe to strongly historicist conceptions of knowledge doubt that truth is especially relevant to an assessment or understanding of scientific knowledge. Social constructivist and feminist critiques of the notion of truth and associated concepts, such as objectivity, engage head-on with the socio-political implications of these ideas, thereby injecting a consideration of them into the epistemology of science.

In what follows, I will outline the range of views that have emerged in the philosophy of science regarding the notion of truth. I will begin in section 2 with views that ascribe no role, or a limited role, to truth in connection with scientific practice and knowledge. The primary driver of scepticism about the role of truth here is most commonly manifested as a dissatisfaction with the correspondence theory of truth, according to which truth is a matter of corresponding, in the right sort of way, to a mind-independent reality. (The correspondence theory holds that a belief (statement, proposition, etc.) is true if it corresponds to the way this reality is, or to the facts or states of affairs that make it up; for a detailed treatment of the view in this volume, see David.) In some of these cases, an account of truth in terms of coherence – that is, in terms of the consistency of a set of beliefs (statements, propositions, etc.) – may better accord with the relevant approaches to science. (The coherence theory holds that a belief is true if it is part of a coherent system of beliefs; on this view, see Walker.) In section 3, I will examine positions that view truth in some form as an appropriate aim of science, and in section 4, I will consider the even stronger commitment to truth adopted by positions that regard it not merely as an aim, but as an achievement of scientific investigation. In the latter case, truth is generally conceived in terms of some sort of correspondence, or truth-making, or both.

One may think of considerations of truth in connection with the sciences by means of two analytical dimensions: the first concerning the degree to which one may think the notion is applicable to the sciences; and the second concerning the precise theory of truth one may think applies. I will consider, in turn, views that are increasingly interested in truth, and for each, identify

the theory or theories of truth that are most appropriate. In section 5, I will conclude with an overview of some challenges regarding truth to emerge in recent work, particularly in connection with scientific realism, the most truth-indulgent philosophy of science.

2. Truth as Orthogonal to Scientific Knowledge

2.1 Historicism and practice-oriented philosophy of science

Logical positivism (or logical empiricism), the founding movement of the philosophy of science as a recognizably distinct sub-discipline of philosophy, dominated the field for much of the twentieth century. In the 1960s, the demise of positivism was hastened by a historical turn in the philosophy of science, associated with authors such as Thomas Kuhn, Paul Feyerabend, and Norwood Russell Hanson, which set the stage for many of the issues discussed in the philosophy of science to this day. In particular, Kuhn's extraordinarily influential book, *The Structure of Scientific Revolutions*, played a large role in establishing a form of historicism about scientific knowledge, and an allied (though strictly separable) preoccupation with the actual practice of science by scientists, both of which have significant consequences for thinking about truth. I will take Kuhn's views as exemplary in this regard for present purposes, though it is important to note that in a number of respects, the morals I will draw concerning truth apply in partially overlapping ways to some of his positivist ancestors, and most certainly to a number of his contemporaries.

The underlying commitment of the approach to science taken in the historical turn was to treat the history of the sciences and their constituent practices seriously as a means by which to investigate the nature of scientific knowledge. This was a self-conscious attempt at *descriptive* philosophy – describing the nature of scientific knowledge as it is actually found – as opposed to a normative or more abstract exercise in idealized epistemological speculation. Kuhn presented the history of the sciences as displaying a recurring pattern of development: periods of so-called normal science punctuated by scientific revolutions, which overturned the established order and thereby laid the foundations for a new period of normal science. One might regard classical physics, for example, as spanning a period of normal science, which was ultimately replaced by relativistic and quantum physics, which in part constitute the normal science of our own time. The key implications for truth on this picture derive from Kuhn's characterization of knowledge on either side of a revolutionary divide. Two different periods of normal science, he held, are

incommensurable with one another, and indeed, exemplify a phenomenon that is now commonly referred to as *world change*. Let us consider these ideas, briefly, in turn.¹

Kuhn held that each period of normal science is typified by a shared commitment to a ‘disciplinary matrix’ or ‘paradigm’, consisting of a number of elements: symbolic generalizations (such as the mathematical formalism of scientific laws); metaphysical beliefs (for instance, the scholastic commitment to the existence of essential natures); values (such as those favouring certain theoretical virtues, like simplicity); and exemplars (standard problem solving algorithms or techniques). Scientific theories falling under the rubric of different paradigms are incommensurable with one another – a metaphor derived from the origins of the term in the Greek geometrical concept of having ‘no common measure’. If two theories are incommensurable with one another, they are not incomparable *per se*, but they are not comparable in way that would allow one, for example, to judge that one is true and the other false, or to determine that one is closer to the truth, relatively speaking. This incommensurability may take the form of a difference in methods or standards operative in science, differences in perception on the part of scientists whose observations are differently “theory laden”, and most importantly for Kuhn, a difference in the very meanings of our words. This semantic incommensurability, which he later analyzed in terms of the failure of meaning preserving translation (Kuhn 1983), problematizes assessments of truth across paradigms.

Kuhn’s theory of meaning is a species of meaning holism, or meaning contextualism: it suggests that concepts are learned in groups, such as those constituting the sets of interconnected beliefs composing paradigms. As a result, when some of these concepts change, the meanings of terms throughout the network are invariably altered. Thus, the term ‘mass’ as used in classical physics simply does not have the same meaning as the term ‘mass’ as used in relativistic physics, and to suggest that one theory’s characterization of mass is closer to the truth is importantly confused – it is to equivocate with respect to two very different concepts which cannot be understood except from within the paradigms in which they occur. Indeed, this holism or contextualism is so pronounced that Kuhn maintained that there is strong sense in which, after a scientific revolution, scientists live in a different world. In the latter stages of his career Kuhn gave a neo-Kantian gloss to this notoriously cryptic remark. One may think of different paradigms as different lenses one must employ so as to engage in and indeed, create the reality of scientific phenomena. Perhaps not surprisingly, this is a challenging idea to articulate and defend.

¹ For just a few of the many detailed studies of these topics, see Horwich 1993, Hoyningen-Huene 1993, Sankey 1994, and Bird 2000.

The pursuit of truth plays an important role, on this picture, during normal science. *Within* the scope of a given paradigm, scientists solve problems and generate facts that may be described in terms of truth. But *across* paradigms, such talk becomes incoherent. One cannot judge one scientific theory to be closer to the truth than its predecessor if the very comparison is rendered ineffectual by semantic incommensurability. And if one takes the idea of world change seriously, the problem is compounded, for there is no paradigm-transcendent, God's-eye point of view from which to judge. There are only paradigmatic judgements; empirical reality is in part structured by our paradigms. One might adopt a coherence theory of truth with respect to the intra-paradigmatic context, since knowledge generated within a paradigm should cohere with it. This will not suffice to allow for anything like scientific progress across paradigms, however, in the way one might hope given a correspondence theory of truth. It is for this reason that those moved by historicist approaches to scientific knowledge are not generally interested in the notion of truth. It is of limited applicability, and entirely inapplicable in the case of correspondence truth. Those inspired by this approach to engage primarily with descriptive features of scientific practice have also often found truth to be an unhelpful or irrelevant concept.

2.2 The sociology of scientific knowledge

Among other innovations, the historical turn in the philosophy of science turned a spotlight on the social context of the sciences. Descriptions of actual scientific practice revealed that complex social interactions are intimately connected to the generation of scientific knowledge – from the training of students by mentors, to the collaborative and competitive dynamics of the hierarchies of public and private institutions in which scientific work is done, to the role of economic, political, and other factors driving research by funding and other means, and so on. It was perhaps inevitable, then, that in the wake of the historical turn some should turn to sociology as a means of understanding the sciences. The sociology of scientific knowledge (SSK) investigates the social contexts of the sciences, and in this by itself, it is philosophically neutral with respect to truth. Just as those who are primarily interested in descriptive features of scientific practice might find the notion of truth surplus to analytic requirements, those primarily interested in social aspects might do likewise. In practice, however, most accounts of science inspired by SSK have significant philosophical consequences. Many suggest that once one appreciates the role that social factors play in generating scientific knowledge, a philosophical commitment to some form of *social constructivism* is inevitable, and this once again problematizes the notion of correspondence truth.

The term ‘social construction’ is used somewhat differently by different authors, but let us proceed here with a general understanding of it, according to which it applies to any process wherein what counts as scientific fact (that is, a claim having the status of scientific knowledge) has been shaped and determined in substantive ways by social factors, and in which different social factors (likely) would have produced facts that are inconsistent with those that were in fact produced. The substantive influences of the social on the scientific are of course numerous (consider, for example, the directions and methodologies of research permitted, encouraged, and funded), but this does not by itself establish the counterfactual dimension of the idea of social construction – the notion that had such influences been otherwise, one would have ended up with facts that are inconsistent with those currently accepted. In order to motivate the counterfactual, scholars in this vein have typically engaged in case studies of scientific work which aim to show how contingent decisions in the workplace are both determined by social factors and could have gone otherwise, resulting in different and potentially conflicting facts. Additionally, some have offered more global reasons for thinking that such contingency is inevitable. Chief among the proponents of the latter approach are supporters of the so-called “Strong Program” in SSK (also known as the “Edinburgh School”).²

Two of the central tenets of the Strong Program are particularly interesting in connection with the notion of truth. The first is the idea of *symmetry*, which suggests that in the scientific domain *inter alia*, beliefs that are taken to be true are on a par with those taken to be false in that both have the same causes, and thus should be given symmetrical explanations. This stands at odds with the not uncommon, asymmetrical practice of explaining the generation of true beliefs in terms of scientific investigation successfully latching on to some facet of a mind-independent world, while explaining false beliefs in terms of something going wrong, whether it be instrument malfunction, human error, corruption, or what have you. Putting *all* scientific belief on a par with respect to their causes opens the door to a substantively social account of both truth and falsehood. A second tenet of interest is the idea of *reflexivity*, which suggests that the claims of SSK are no less subject to sociological analysis (including the symmetry principle) than any putatively factual claims. This entails a wide-ranging commitment to relativism: truth and falsehood (even with respect to SSK) are defined only in the context of a social community, and have no community-transcendent

² For a mature statement of the Strong Program, see Barnes, Bloor & Henry 1996. For a sampling of different approaches to social constructivism, see Knorr-Cetina 1981, Pickering 1984, Shapin & Schaffer 1985, Latour & Woolgar 1986, and Collins & Pinch 1993.

meaning. As with many relativist positions, there is ongoing debate as to whether this feature of the view renders it self-undermining.

Later incarnations of the Strong Program adopt a theory of language acquisition, *meaning finitism*, which also has profound implications for truth. Finitism can be viewed as an elaboration of the later work of Kuhn, and especially as deriving from the philosophy of the later Wittgenstein. Kuhn emphasized the idea that one learns the meanings of concepts by mastering the use of exemplars that incorporate them, and this point can be extended to training, education, and acculturation more broadly. The ensuing claim of finitism is that meanings are social institutions – the meaning of a term is simply the concatenation of ways in which it can be used successfully in communication within a linguistic community. This relegation of meaning to social facts can be exploited so as to derive a forceful case for the idea that such terms need have no fixed or determinate meanings at all. For on this view, the meaning of a term is constituted by the social acceptance of its use, and this is something that can change subject to social negotiation, reflecting collective decisions about how to go on using language, which may themselves not be universally accepted. This once again appears to favour coherence as an analysis of truth regarding propositions entertained within a scientific community. Whether finitist observations regarding language acquisition successfully entail the associated view of meaning, however, is certainly contestable.

2.3 Feminist critiques of science

Just as the historical turn in the philosophy of science suggested and facilitated the emergence of SSK and the development of a number of forms of social constructivism, these latter positions facilitated the development of feminist critiques of science, often in conversation with them. There is significant overlap between SSK and feminist approaches in the acknowledgement of the role of social factors as determinants of scientific fact, but feminist scholars have extended this analysis in more specific ways, reflecting their more specific motivations. Another striking difference is the near universality with which feminist approaches are normative, offering corrective prescriptions for understanding concepts such as objectivity and knowledge (not to mention proposals for reforming scientific practice itself) which, as we shall see momentarily, have implications for how one conceives of truth.

Several prominent positions have emerged in this literature, and it will be useful to mention them in passing. *Feminist empiricism* emphasizes the possibility of warranted scientific belief in a

communal setting, but crucially, only where biases which enter into research – stemming from gender, ethnicity, socio-economic or political status, and so on – are transparent and appropriately considered. *Standpoint theory* explores the contention that differences in gender, ethnicity, socio-economic and political status, etc. define perspectives to which knowledge is inextricably indexed. *Feminist postmodernism* rejects traditional conceptions of universal or absolute objectivity or truth.³ Taken together, these positions offer a number of considerations which might push one in the direction of one theory of truth or another, or indeed, towards a rejection of any substantive theory at all. Foremost among these are considerations of objectivity and perspective. Let us examine them together, as they are interestingly connected.

The term ‘objective’, when used traditionally in connection with knowledge, has a number of stereotypical connotations. Most important in this context are connotations of disinterest (detachment, lack of bias), and of universality (independence of any particular angle of approach or perspective). Knowledge of this sort arguably comprises propositions that are candidates for describing a mind-independent world, and thus, is consistent with a correspondence theory of truth. There is near consensus in feminist critiques of science that objectivity in the sense of disinterest is regularly violated in science. Case studies are adduced to demonstrate how the presence of, for instance, androcentric bias exemplified by a community of scientists can lead to the establishment of one scientific fact at the expense of another, mutually inconsistent possibility (for example, see Longino 1990 for detailed accounts of two cases: one concerning explanations of the evolution of modern human biological characteristics in paleoanthropology; and another concerning the influence of hormones on biological characteristics and behaviour in neuroendocrinology). Many also reject the idea of objectivity in the sense of universality or perspective-independence, and as I will suggest, it is this latter thesis that is required in order to problematize the notion of correspondence truth.

Why is the presence of bias in scientific work not by itself sufficient to rule out correspondence? There are two reasons for this. One is that the presence of bias may not be, in a given case, epistemically significant. That is to say, if regardless of one’s bias one would have assessed the evidence for a given hypotheses in the same way, one might thereby end up with true beliefs concerning that hypothesis in the correspondence sense quite independently of one’s bias. Of course, most feminist case studies are not of this sort – they are cases in which bias *does* make a difference. A second reason that the presence of bias (taken by itself) is consistent with the idea of

³ For a sample of a number of influential versions of these views, see Keller 1985, Harding 1986, Haraway 1988, Longino 1990, Alcoff & Potter 1993, and Nelson & Nelson 1996.

correspondence truth, is that if the institutional infrastructure of science were arranged in such a way as to make epistemically significant bias transparent, then it should be possible in principle to expose such bias and correct for it where necessary. It is for this reason that many feminist philosophers of science advocate for institutional structures that would expose sources of bias. Consider, for example, the idea of effective peer review, where “peers” are drawn from across the space of possible biases exemplified by human agents. The contention that in many and perhaps most domains, the sciences do not yet instantiate such structures, is one motivation for the normative character of most feminist critiques.

If epistemically significant bias is a fact of the matter about much contemporary science, then the prospect of scientific knowledge of a mind-independent world is thereby diminished in practice if not in principle. A more in-principle concern emerges from the idea of perspectives. For just as the Strong Program in SSK embraces a thoroughgoing relativism, where truth is defined only within the confines of a context and not universally, some standpoint theory and certainly all postmodernism is likewise dismissive of the idea of standpoint- or perspective-independent truth. It is important to note here, however, that some would reject this implication. Some standpoint theorists, reflecting a Marxist inspiration, claim that certain perspectives are epistemically privileged over others – typically, subjugated perspectives are privileged over dominant ones as having deeper insight into the subject matter, in just the way that the proletariat might have deeper knowledge of human potential than the superficial knowledge possessed by those in positions of power. In the absence of such epistemic privilege, however, one is left with an irreducibly fractured picture of warranted scientific belief, where (at best) coherence within perspectives, not correspondence of beliefs from across perspectives, might serve as an analysis of truth.

3. Truth as an Aim of Science

3.1 The ideal end of inquiry

Having canvassed some reasons for dismissing the relevance of correspondence conceptions of truth in the context of the sciences (in some cases, in favour of a notion of coherence), or for dispensing with concerns about truth altogether, let us turn now to views in which truth is offered explicitly as a goal of scientific investigation. One of the most influential themes to emerge in the philosophy of science in this regard represents scientific inquiry as converging, ultimately, on the truth. The idea of convergence has two features of particular

relevance here. One is that convergence is something that can be realized by a process – in this case, the development of scientific knowledge – quite independently of whether the goal itself is ever realized. (Consider the analogy of a mathematical limit that can be approached but never quite realized, except at infinity.) Another feature of interest here is the broad compatibility of the notion of convergence with different views of science. Some scientific realists, for example, whom we will consider in section 4, hold that scientific knowledge is converging on truths in the correspondence sense. But one need not be a realist of this or any other sort to think that truth is the ultimate aim of science.

An important case in point of non-realist endorsement of convergence is associated with the pragmatist tradition in philosophy. The terminology can easily mislead here, since some self-avowed pragmatists also refer to themselves as realists, but as we shall see, their realism is generally not what goes by the name ‘scientific realism’ more specifically. One way of generating the distinction, conveniently for present purposes, is to pay attention to the theories of truth typically endorsed by these camps. While scientific realists generally opt for some version of the correspondence theory – or if they dislike classical accounts of truth, some version of truth-maker theory – pragmatists generally do not. (This is a traditional characterization of pragmatism; for a more nuanced treatment, see Misak in this volume.) In order to appreciate the pragmatist’s conception of convergence, let us digress momentarily to consider her connected notions of meaning, and truth.⁴

As is the case in any grand philosophical movement, not all participants share precisely the same view, but one tenet shared by many if not all pragmatists is a commitment to a criterion of meaning, whereby the content of a proposition is given, as C. S. Peirce put it in his essay ‘How to Make Our Ideas Clear’, by clarifying its ‘practical consequences’. Such consequences are to be understood in terms of human experience; that is, as having implications for actual or possible empirical observations. In the context of the sciences, this concerns everything from successful prediction, retrodiction, and communication, to heuristic use and problem solving, where theoretical claims are employed as a basis for action. In the work of William James in particular, this analysis of meaning was parlayed into a theory of truth, according to which positive utility of the sort indicated by useful practical consequences is the marker of truth. Given the fallible nature of scientific knowledge, however, the application of this concept is conceived of by many pragmatists as something of an ideal: truth is whatever it is that will be agreed upon in scientific investigation in

⁴ Here I will focus on some general pragmatist principles drawn primarily from James 1979/1907 and Peirce 1998/1992, though readily recognizable in the work of many who have identified with pragmatism since.

the ideal limit of inquiry. For both Peirce and James, the truth in this sense exhausts our conception of reality, which is made by us, not mind-independent. Peirce in particular held that through a process analogous to evolution, scientific methods will converge on a particular body of knowledge. In this way, truth may be conceived as the ultimate aim of science.

The evolutionary analogy here is an evocative one, but its appropriateness in the context of scientific knowledge may be questioned in a number of ways, including some we have already considered. If strongly historicist readings of scientific knowledge are compelling, for example, then the very idea of convergence is problematic, since in this case it is difficult to understand what it could mean for scientific theories across different periods of history to be moving in any particular direction. Indeed, Kuhn himself was attracted to an evolutionary picture of the development of scientific knowledge, but of a rather different sort. On our modern understanding of the analogue – evolution by natural selection – there is no sense in which the adaptations of organisms are “moving” in the direction of anything resembling a definitive ideal. Adaptations in the biological case are more or less optimal only relative to particular conditions of existence represented by particular ecological settings, and the same sort of thing (*mutatis mutandis*) might be thought true of knowledge in historical settings. It is for this reason that many regard the evolutionary analogy as properly suggesting only a notion of progress *from* a particular state of knowledge, as opposed to progress *towards* an ideal, and without the latter, there can be no convergence. Similar appraisals can be generated from the point of view of some social constructivist and feminist positions, and these are contested, as we shall see, by some varieties of realism, which like pragmatism subscribe to the idea of convergence.

3.2 Constructive empiricism and one formulation of realism

The pragmatist notion of truth in the ideal limit of inquiry is an example of one manner of thinking about truth as an aim of science. I have described this manner above in terms of conceptions of the *ultimate* aim of scientific investigation. One might think of truth as a goal, however, in rather different terms, as a *proximate* aim. That is to say, one might think that the sciences aim at the truth in the present, not merely as an imagined outcome in a limit. Indeed, it is possible to hold that truth is both a proximate and an ultimate aim of science in precisely these senses – that by aiming at truth in the present, one might thereby realize it in the limit of inquiry, whether this limit is held to be an idealization (as suggested by Peirce), or something one might think is achievable in a finite time, or even now with respect to some scientific facts, as many (if not

all) realists contend. This combination of proximate and ultimate aims would seem to be consistent with any position that takes scientific knowledge to be converging on the truth, in some way, shape, or form. Having considered one version of the idea that science aims at truth in the limit, let us turn to the idea that it aims at truth in the here and now.

Bas van Fraassen (1980) has been influential in suggesting that different views of the nature of scientific knowledge can be described in terms of how they view the proximate aim of science. This, I submit, is a tricky business, for it is unclear why *any* view of scientific knowledge should take it to have any one particular or primary aim, given the many candidates that seem to be transparently evident in scientific practice. Based on the activities of scientists, one might plausibly contend that the sciences aim to describe, to predict, to explain, and much more besides. Given this plethora of aims, a question arises as to what philosophical argument could establish the primacy of any one in particular. One possibility here would be to argue that some particular, over-arching, epistemic aim makes good sense of scientific practice and its various other aims, by accounting for the latter aims as subsidiary to the former. Pronouncements regarding “the (proximate) aim of science” issuing from philosophical accounts of scientific knowledge generally do not engage in such argument, however, and for this reason, any claim of this sort must be viewed as harbouring a significant promissory note.

Van Fraassen offers a characterization of scientific knowledge in terms of aims on behalf of two positions: his own view, which he calls ‘constructive empiricism’; and the view with which he contrasts it, scientific realism. The constructive empiricist holds that the aim of science is *empirical adequacy*, where this is defined in the following way: ‘a theory is empirically adequate exactly if what it says about the observable things and events in the world, is true’ (1980, p. 12).⁵ In contrast, realism is described as the view according to which the aim of science is truth *simpliciter* – not merely regarding scientific claims about observable things and events, but also regarding claims about unobservable things and events. In both cases, truth may be understood in a correspondence sense, and the distinction between the observable and the unobservable is drawn in terms of human sensory capabilities. Whatever is detectable by the unaided senses (such as planets, alpacas, and metals) is observable, and whatever is not so detectable (such as dark matter, genes, and

⁵ This definition suffices for present purposes, but it is worth noting that van Fraassen goes on to define empirical adequacy again in more technical terms: a theory is empirically adequate if the observable content of scientific investigation – ‘the structures which can be described in experimental and measurement reports’ (1980, p. 64) – are isomorphic to empirical substructures of some model or models of the theory, where a model is understood to be any structure of which the axioms and theorems associated with the theory are true. The relevant empirical substructures can then be taken to represent observable things and events.

neutrons) is unobservable. While this particular opposition regarding the putative aim of science has attracted significant attention, one could not imagine other possible epistemological stances with respect to scientific knowledge, each defined by a different proposal for the principal aim of scientific investigation.

There is arguably good reason, however, to be dissatisfied with this strategy for defining epistemological positions such as constructive empiricism and scientific realism. To define an epistemological position purely aspirationally – that is, in terms of aims – may strike one as too weak. For so defined, these positions are consistent with never actually achieving the relevant aim; indeed, they are consistent with the *impossibility* of achieving the relevant aim. Defined aspirationally, constructive empiricism is consistent with the belief that no scientific theory is empirically adequate, and with the belief that no theory will ever or could ever be empirically adequate. The same moral would seem to apply to realism, in connection with a more comprehensive truth of theories, covering both the observable and the unobservable. This separation of views concerning what science aims to do from those concerning what it actually achieves, or is capable of achieving, is perfectly consistent insofar as one is interested in aims alone. But from the point of view of epistemology, it is too weak. Many realists, for example, believe not only that the sciences aspire to truth, but that under certain conditions, they achieve this aim, in a manner to be specified. Let us move on now, to consider philosophical approaches to scientific knowledge that explicitly adopt an achievement view of truth.

4. Truth as an Achievement of Science

4.1 The many forms of realism

There are many analyses of scientific knowledge that are or could be, if interpreted in a favourable manner, consistent with the idea that scientific theories are true, to some extent or other. This last qualification, ‘to some extent or other’, introduces a complication which I have thus far ignored, but must now address. Those who believe that the sciences yield truths are often very careful about how they describe this commitment to truth. Some scientific claims (for example, certain claims about what sorts of things exist), they might say, are likely true *simpliciter*, but others (such as descriptions of laws of nature, or theories taken as wholes) are often merely *approximately true* – that is, close to the truth, in some sense to be explicated. Convergentsists in this camp then commonly hold that scientific theories are becoming increasingly approximately true over time,

with the development of science. I will return to the idea of approximate truth in section 5, but let it suffice here to say that when I describe views that take truth to be an achievement of the sciences, I am including those which subscribe to the idea of approximate truth as an appropriate surrogate for truth *simpliciter*, as and where appropriate.

A further qualification is in order here, concerning the breadth of positions that might, on a favourable interpretation, qualify as admitting scientific truths. Many so-called antirealist positions in the philosophy of science are consistent with scientific truth, but only regarding a very specific domain: the observable. *Instrumentalism*, for example, is the view that scientific theories are simply instruments for predicting observable phenomena, or for systematizing observation reports. On a traditional reading, this view also holds that claims pertaining to unobservable things, by themselves, have no meaning at all – they are merely instruments, not candidates for truth or falsity. Less traditionally, positions are sometimes described as instrumentalist if they admit epistemic access only to truths about observables; claims about unobservables here may be true or false, but the relevant truth value cannot be known. In either case, it is open to an instrumentalist to believe that the sciences yield truths about observables in the correspondence sense. Others, such as some logical empiricists, might agree, but only insofar as truth is understood in terms of coherence, in accordance with some or other form of neo-Kantianism. Others, such as Kuhn, might extend the remit of the neo-Kantian-coherence-type approach to truth to the unobservable, but only within periods of normal science. As should be clear, the possible combinations of restriction on the domain of truth, and theory of truth applicable, are many!⁶

The most unequivocal commitment to truth talk in connection with the sciences, however, is exemplified by versions of scientific realism. Most realists suggest, either explicitly or implicitly, that our best scientific theories are true or approximately true in some appropriate sense of correspondence with a mind-independent world.⁷ Understood this way, scientific realism is the view that scientific theories correctly describe both observable and unobservable features of the world. The main consideration offered in support of realism is an old idea, resonant with common sense, to the effect that the very empirical success of science – in prediction and

⁶ For a formative influence on instrumentalism, see Duhem 1954/1906. For discussion of the neo-Kantian reading of logical empiricism, see Richardson 1998 and Friedman 1999, and for a reading more compatible with realism, see Psillos forthcoming. The neo-Kantian theme is developed differently under the label ‘internal realism’ in Putnam 1981. For a consideration of fictionalism, which holds that things in the world are and behave as if our best scientific theories are true (thus admitting truths about observables), see Vaihinger 1923/1911 and Fine 1993.

⁷ There are exceptions, however. For arguments in favour of separating realism from the correspondence theory of truth, see Ellis 1988, Leeds 2007, and Devitt 2010 (especially chapters 2 and 4). I will return to this idea in another form in section 5.

manipulation of the natural world, not to mention technological applications – indicates that its theories are true. In recent times this idea is commonly referred to as the ‘miracle argument’ (or ‘no-miracles argument’), after Hilary Putnam’s (1975, p. 73) contention that realism ‘is the only philosophy that doesn’t make the success of science a miracle’. In other words, it would be miraculous if the sciences were as empirically successful as they are, and yet their theories were untrue, thus suggesting the greater credibility of an explanation of success that appeals to truth. As we shall see momentarily, however, the miracle argument is highly contested.

As it happens, there are many varieties of scientific realism, some developed with the intention of responding to forms of antirealist scepticism such as that concerning the miracle argument. The most general form of realism subscribes to the formula I have already presented: a general endorsement of the truth of our best scientific theories. Of the variations on this formula to emerge, two in particular have come to prominence in recent philosophy of science, both of which offer a more specific prescription for realist commitment. *Entity realism* is the view that under conditions in which one has significant causal knowledge of a putative (unobservable) entity, allowing one (*ex hypothesi*) to manipulate it and use it to intervene in other phenomena, one has good reason to think it exists, and consequently, that claims about the existence of such entities are true. The negative implication of this commitment is that the theories that describe these entities need not be true, and this raises a question about whether an antirealism about theories is ultimately compatible with a realism about the entities they describe. This view also appears to entail that one may continue to believe in the existence of an entity despite radical changes in one’s (theoretical) understanding of what that entity is like. This in particular might seem problematic from a historicist point of view, or indeed, from any point of view that takes accurate description to be an important component of successful reference to an entity.

A second more specific version of realism is *structural realism*, which holds that insofar as scientific theories offer true descriptions, they do not tell us about the natures of things like unobservable entities. Instead, they correctly describe the structure of the unobservable realm. There are two main sorts of structural realist positions: those that take the distinction between structure and nature to be an epistemic distinction; and those that take it to be an ontological distinction. The former suggest that knowledge of the entities that participate in the structural relations described by theories (for example, by mathematical laws) is simply beyond our grasp, but that theories furnish a kind of structural knowledge nonetheless. The latter suggest that structural knowledge is the most a realist can hope for, because there are in fact no entities that stand in those relations, or that if there are such entities, they are in some sense emergent from or

dependent on the structures in which they seem to appear. Like entity realism, structural realism faces interesting challenges. One might wonder, for example, about how clear the distinction between structure and nature is, or whether this distinction is tenable in connection with the concrete systems of things investigated by the sciences. The ontological version specifically faces the challenge of making the relevant ideas of emergence and dependence of entities intelligible.⁸

4.2 Scepticism about realism

Entity realism and structural realism each face challenges aimed at the coherence of their specific formulations of realist commitment. Additionally, it is important to note that all the reasons we encountered earlier for being suspicious of the notion of correspondence truth, in connection with different views regarding the nature of scientific knowledge, are also ultimately challenges to scientific realism as it is usually conceived. Often, in the background of these worries, are more general sceptical concerns about scientific knowledge that any realism must face if it is to be defensible. Indeed, the advent of both entity realism and structural realism was intended, in part, to offer responses to some of these more general concerns. I will not here consider whether or to what extent they are successful in this regard, but given the importance of these concerns to motivating (sometimes explicitly but often implicitly) a number of the antirealist positions we encountered earlier, let us consider them briefly. This more general antirealist scepticism stems from three notable issues: the ubiquitous use in science of a form of reasoning called inference to the best explanation; the so-called underdetermination of theories by data; and discontinuities in scientific theories over time.

Scientific inferences are generally inductive. One is rarely if ever in a position to generate scientific conclusions, laws, theories, and so on, by means of arguments that are deductively valid. Instead, one marshals evidence by means of observation and experiment, and reasons on the basis of this and other commitments to conclusions that are, ideally, inductively strong. In the sciences, a particular form of inductive inference is prevalent, *viz.* inference to the best explanation, according to which one infers hypotheses that, if true, would provide the best explanation for whatever it is that one is attempting to explain. Natural though this pattern of inference may be, antirealists have

⁸ For a defence of scientific realism, see Psillos 1999. For a discussion of novel predictions, often invoked in characterizing what it might mean for something to be one of our best theories, see Leplin 1997. Cartwright 1983 and Hacking 1983 are definitional texts for entity realism, as are Worrall 1989, French 1998, and Ladyman & Ross 2007 for structural realism. For a discussion of the relationships between scientific realism, entity realism, and structural realism, see Chakravartty 2007.

been quick to point out two potential difficulties with it in connection with truth. First, in order that one infer the truth, one must have the capacity to rank rival hypotheses reliably with respect to their likelihood of being true. But what assurance is there that the sorts of criteria typically employed in the sciences for purposes of ranking – the simplicity of a hypothesis, its consistency with other hypotheses one may like, its heuristic value, and so on – are indicative of truth? And even if one were to grant that the criteria for choice employed in scientific practice are truth-tracking, in order for inference to the best explanation to result in true belief, it is imperative that the true hypothesis be among those one is considering. But this, arguably, is not generally something one can know in advance.⁹

Moving on now to our second sceptical concern, the idea of the underdetermination of theory by data, derived from the work of Pierre Duhem, runs this way.¹⁰ Scientific hypotheses do not yield predictions all by themselves. In order to derive predictions, they must be conjoined with background theories, other theories, theories about instruments and measurements, and so on. Now, if observation and experiment result in data that do not agree with one's predictions, what part of one's scientific belief set should one adjust accordingly? Naively, one might think that the hypothesis under test should now be viewed suspiciously, but Duhem points out that given all of the things one must assume in order to derive the faulty prediction in the first place, it is not so obvious which part of this apparatus deserves suspicion. Different adjustments to one's belief set – different overall conjunctions of hypotheses and theories – will be consistent with the data. Therefore, there is always more than one overall set of beliefs that is consistent with the data. In more contemporary discussions, underdetermination is usually presented slightly differently, in terms of the idea that every scientific theory has empirically equivalent rivals: ones that agree with respect to the observable, but differ elsewhere. This then serves as the basis for a sceptical argument regarding the truth of any particular rival chosen for belief. It is precisely this sort of scepticism that underwrites the worries about correspondence truth commonly found in SSK and feminist critiques of science.

Finally, let us turn to an argument that is often simply referred to as the 'pessimistic induction', or the 'pessimistic meta-induction'. The history of the sciences is replete with discontinuities in the knowledge they ostensibly provide. Over time, older theories recede into the past as newer theories are proposed and accepted, resulting in changes in belief with respect to the

⁹ For a thorough exploration of inference to the best explanation, see Lipton 2004, and for an influential critique of it, see van Fraassen 1989, chapter 6.

¹⁰ The idea of underdetermination is commonly traced to Duhem 1954/1906, chapter 6, and many see affinities here with the later "confirmational holism" of W. V. O. Quine (1953). For this reason, it is sometimes called the 'Duhem-Quine thesis'.

entities, their properties and relations, and laws described by these theories. Thus, from the point of view of any given moment in time, most past theories must be considered false, strictly speaking. By induction, then, is it not likely that most present-day theories are also false? The ways in which they will, in future, be revealed as false may not yet be clear to us, but the inductive challenge remains. Given a knowledge of the history of theorizing in any particular domain of scientific investigation, realism may thus appear a rather too optimistic position regarding scientific knowledge. This general sceptical worry can be posed in a number of different ways, focusing on the falsity of past theories, unsuccessful reference on the part of their central terms, and in other ways besides.¹¹

Worries about inference to the best explanation, the underdetermination of theory by data, and the pessimistic induction have been highly influential in the philosophy of science. It would be irresponsible not to mention, however, that realists of various stripes have responded to these concerns in a number of ways, and the ensuing debates are very much alive. A thorough canvassing of the details of these debates lies far beyond the scope of this essay – indeed, these details comprise a large proportion of the field itself. I hope, however, that the ways in which the notion of truth bears on the sciences is now substantially clearer than when we began. My aims in this regard now largely fulfilled, let us turn in the next and final section to a few important considerations involving truth to emerge in areas of very recent interest in the philosophy of science.

5. Truth and scientific models

5.1 Abstraction and idealization

One of the most important developments in the philosophy of science over the past two decades has been an increasing awareness and study of the role of scientific models in discussions of scientific knowledge. In keeping with a general movement towards closer attention to scientific practice in the wake of the historical turn, a traditional emphasis on knowledge encapsulated by scientific theories has been supplemented (and perhaps, in some authors' estimations, replaced) by an emphasis on the ways in which models associated with these theories are used to represent target systems in the world. One reason for this, no doubt, is that 'scientific theory' is a somewhat loosely defined concept – its referents vary significantly depending on the context of its use. One

¹¹ Canonical formulations of the pessimistic induction are most often drawn from Laudan 1981. For a related and more recent version of history-induced scepticism, see Stanford 2006.

and the same theory may be identified in some cases with a set of equations, and in other cases with various models applied to different classes of phenomena putatively described by those equations, but in different ways. These different adaptations of mathematical descriptions to fit particular sorts of cases – sometimes even in ways that are mutually inconsistent – suggests that in many situations, it is the models themselves that are the primary representational tools of scientific description. Thus, a spotlight is turned onto the relationship between modelling and truth.

The term ‘scientific model’ covers many things here. It includes: abstract entities, such as theoretical models about which scientists may have some shared conception; concrete objects, such as graphs, diagrams, other illustrations, and three-dimensional physical representations, such as scale models; and processes, such as computer simulations. Though the forms of these representations are many, the question of how they can be thought to represent their targets veridically, if at all, has become a central topic of debate, not least because, as it happens, scientific models are often constructed so as to *deviate* from the truth. There are two primary ways in which this occurs, and I will call these practices *abstraction* and *idealization*. Although these are terms of art, used in slightly different ways by a growing number of philosophers, there is something of a consensus regarding the basic distinction. For purposes of illustration, I will sketch a representative version of this distinction below.¹²

‘Abstraction’ and ‘idealization’ can be thought of as labelling processes – ones by means of which models are constructed – as well as the resultant models themselves. Thus, an abstract model is the result of a process of abstraction, a process in which only some of the potentially many factors that are relevant to the nature or behaviour of something in the world are represented in a model of it. In the process of abstraction, other factors are ignored, often intentionally, in order to allow for the construction of models that are mathematically or otherwise tractable; and often relevant factors are unintentionally ignored, because their relevance is unknown. Consider, for example, the model of a simple pendulum. Many simplifying assumptions are made in constructing this model, such as the omission of frictional resistance due to the air in which the pendulum swings. Such assumptions are commonly held to compromise the truth of models not only because theoretically important aspects of their target systems in the world are left out, but also because, as a result, the predictions of these models differ from that which one finds in reality. One might argue that in both of these ways, abstraction introduces falsehood into practices of scientific

¹² For influential discussions in this area, see McMullin 1985, Cartwright 1989 (chapter 5), and Suppe 1989 (pp. 82-83, 94-99). For a comprehensive treatment, see Jones 2005. My take on the distinction here is abridged from Chakravartty 2010a, which considers the implications for truth in some detail.

representation. By the same token, however, the notion of abstraction suggests one way of making sense of the idea that one representation can be more approximately true than another: by incorporating a greater number of factors that are causally (or otherwise) relevant to their natures or behaviours.

Similarly, an idealized model is the result of a process of idealization, a process in which at least one of the features of the target system incorporated into the model is represented in a distorted or simplified manner. This process differs from abstraction in that one is not excluding factors, but rather incorporating them in such a manner as to represent them in ways they are not. (In order to sharpen the distinction between abstraction and idealization further, I view the latter in terms of representing things in ways they *could not* be, given the laws of nature. This is sometimes but not always true of abstract models: some abstract descriptions may be true of other systems in the world in which the omitted factors happen not to be present, which is something scientists often contrive in laboratory settings. In any case, given that abstraction and idealization are not mutually exclusive processes, some models may be both abstract and idealized.) Idealized models are, in a stronger sense than in the case of abstractions, false representations of things in the world. It is an implicit assumption of classical physics, for example, that the masses of bodies are concentrated at extensionless points at their centres of gravity, but this is something we clearly know to be false, notwithstanding the predictive and explanatory efficacy of models in classical physics. In such cases, understanding how one model can be more approximately true than another is more complex than in the case of abstraction, turning on how “degrees” of distortion may be conceived in particular cases.

5.2 Approximate truth and deflation

The notion of approximate truth is generally viewed as important to scientific realism, versions of which, as we have seen, are commonly more interested to characterize scientific knowledge in terms of truth than any other philosophical assessment of the sciences. The apparent importance of approximate truth here is plain, given that even realists must acknowledge that many if not all scientific theories contain falsehoods, and are thus, strictly speaking, false. The task of reconciling the realist’s positive appraisal of scientific knowledge in terms of truth with such obviously recalcitrant evidence of falsity provides strong motivation for thinking about how one theory can be “better off” than another with respect to truth, even when both are false.

The concept of approximate truth, and others such as ‘verisimilitude’ and ‘truthlikeness’ (sometimes used synonymously with ‘approximate truth’ and in other cases not, reflecting the intentions of different authors), have become subjects of detailed and technically striking work in their own right. I will not consider these developments here (but another essay in this volume does; see Oddie). Instead, my goal in this last section is to complete an examination of the implications of the recent interest in scientific modelling for considerations of approximate truth, and for truth more generally. For some advocates of the modelling conception of science hold that once one appreciates how models are used to represent things in the world, all substantive concerns about truth in this context are dissolved. These concerns arise in the face of two challenges: first, the challenge of explicating the notion of correspondence truth; and second, the challenge of explicating the notion of approximate truth. Given the difficulties inherent in both of these challenges, if it were the case that paying attention to practices of scientific modelling could obviate the need to meet them, this would be a significant result. Let us consider the prospects for doing so.

The most straightforward exposition of the idea that considerations of scientific modelling dissolve substantive issues of truth for the realist is found in the work of Ronald Giere, on whose view a scientific theory is constituted by ‘two elements: (1) a population of models, and (2) various hypotheses linking those models with systems in the real world’ (1988, p. 85).¹³ Hypotheses assert relationships of similarity between aspects of models and their target systems, and are true or false depending on whether these relationships obtain. It is the fact that models are the tools by means of which these epistemic commitments are expressed that dissolves concerns about truth, or so the argument goes. On this approach, Giere contends, there is no question of attempting to forge any sort of correspondence between linguistic entities and the world, and consequently, the challenge to explicate the notion of correspondence truth does not even arise. All that is required are assertions of similarity between models and the world, and for this a deflationary account of truth will suffice.

Deflationism about truth holds that there is nothing metaphysically substantive (such as the notion of correspondence) to the idea of truth (see Azzouni in this volume). One influential version of the view suggests that the predicate ‘is true’ in ‘ α is true’ is redundant; it adds nothing to the meaning or the assertion of ‘ α ’. Thus, to say that it is true that a two-body Newtonian gravitational model is similar in various ways to the system comprising the earth and moon, for example, is

¹³ Giere refines this portrait of scientific theories in subsequent work, but the refinements are immaterial for present purposes. For his most developed view of scientific knowledge borne of modelling, see Giere 2006, and for an appraisal of this sort of approach, see Chakravartty 2010b. For a longer discussion of modelling, correspondence, and approximate truth, see Chakravartty 2007, sections 7.4-7.5.

simply to say that these similarities obtain – the truth predicate is redundant. And once scientific claims take the form of asserted model similarities, clarified in terms of respects and degrees of similarity, the need for an overarching account of approximate truth is likewise superfluous, or so one might argue.

The modelling approach facilitates many important insights into the nature of scientific knowledge, but it does not, I suspect, make the issue of correspondence disappear for the realist. Whether one prefers to think of correspondence in terms of a property shared by all truth-bearers, or in terms of the existence of truth-makers for scientific claims, *some* form of correspondence is a requirement of realism. As a quick demonstration of the insufficiency of asserted model similarities in this context, consider the claim that the double-helical structure of Watson and Crick's physical model of the DNA molecule is similar to the double-helical structure of DNA molecules. Most scientific realists would endorse this claim, but *so too* would most traditional instrumentalists. The assertion of similarity here does not by itself distinguish the different ontological significance these different parties attach to it. For the realist, the assertion of similarity is interpreted as truly describing the properties of an unobservable entity, whereas for the instrumentalist, it is interpreted as asserting something that is neither true nor false, but useful as a vehicle for the derivation of observable predictions. It is only by admitting the relevance of considerations of correspondence, in one form or another, that allows one to distinguish these interpretations.

Neither does the possibility of describing similarities between scientific models and their target systems eliminate the yearning for an account of approximate truth. Though similarities of this sort are the bread and butter of scientific representation, there are undoubtedly other epistemic contexts, of a broader and grander sort, in which scientists and philosophers alike find themselves, understandably, reflecting on the notion of truth. Stepping back from the minutiae of particular representations to the world of scientific knowledge more broadly conceived, there are contexts within which it is entirely appropriate to contend that relativistic physics is closer to the truth than classical physics, or that the modern synthesis in evolutionary biology is closer to the truth of the nature of organisms and their populations than biology before Darwin. It is precisely this sense of progress that is celebrated by many realists, and in different ways by other convergentists, and disputed by varieties of historicists and social constructivists. It is this sense of truth that is the subject of so many of our most important and productive philosophical disputes about the relationship of truth to the sciences.

Acknowledgments

For helpful comments on an earlier version of this essay, I am grateful to Agnes Bolinska and Michael Glanzberg.

References

- Alcoff, L. & E. Potter (eds.) 1993: *Feminist Epistemologies*, London: Routledge.
- Azzouni, J. 2011: 'Deflationism', in M. Glanzberg (ed.), *The Oxford Handbook of Truth*. Oxford: Oxford University Press.
- Barnes, B., B. Bloor & J. Henry 1996: *Scientific Knowledge*, London: Athlone.
- Bird, A. 2000: *Thomas Kuhn*. Chesham: Acumen.
- Cartwright, N. 1983: *How the Laws of Physics Lie*. Oxford: Clarendon Press.
- Cartwright, N. 1989: *Nature's Capacities and Their Measurement*. Oxford: Clarendon.
- Chakravartty, A. 2007: *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge: Cambridge University Press.
- Chakravartty, A. 2010a: 'Truth and Representation in Science: Two Inspirations from Art', in R. Frigg & M. Hunter (eds.), *Beyond Mimesis and Convention: Representation in Art and Science*, pp. 33-50. Boston Studies in the Philosophy of Science. Dordrecht: Springer.
- Chakravartty, A. 2010b: 'Perspectivism, Inconsistent Models, and Contrastive Explanation', *Studies in History and Philosophy of Science* 41: 405-412.
- Collins, H. & T. J. Pinch 1993: *The Golem*. Cambridge: Cambridge University Press.
- David, M. 2011: 'The Correspondence Theory of Truth', in M. Glanzberg (ed.), *The Oxford Handbook of Truth*. Oxford: Oxford University Press.
- Devitt, M. 2010: *Putting Metaphysics First: Essays on Metaphysics and Epistemology*. Oxford: Oxford University Press.
- Duhem, P. M. M. 1954/1906: *The Aim and Structure of Physical Theory*, P. P. Wiener (tr.). Princeton: Princeton University Press.
- Ellis, B. 1988: 'Internal Realism', *Synthese* 76: 409-434.
- Fine, A. 1993: 'Fictionalism', *Midwest Studies in Philosophy* 18: 1-18.
- van Fraassen, B. C. 1980: *The Scientific Image*. Oxford: Clarendon.

- van Fraassen, B. C. 1989: *Laws and Symmetry*. Oxford: Clarendon.
- French, S. 1998: 'On the Withering Away of Physical Objects', in E. Castellani (ed.), *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, pp. 93-113. Princeton: Princeton University Press.
- Friedman, M. 1999: *Reconsidering Logical Positivism*. Cambridge: Cambridge University Press.
- Giere, R. N. 1988: *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.
- Giere, R. N. 2006: *Scientific Perspectivism*. Chicago: University of Chicago Press.
- Hacking, I. 1983: *Representing and Intervening*. Cambridge: Cambridge University Press.
- Harding, S. 1986: *The Science Question in Feminism*. Ithaca: Cornell University Press.
- Haraway, D. 1988. 'Situated Knowledges', *Feminist Studies* 14: 575-600.
- Horwich, P. (ed.) 1993: *World Changes: Thomas Kuhn and the Nature of Science*. Cambridge, MA: MIT Press.
- Hoyningen-Huene, P. 1993: *Reconstructing Scientific Revolutions: The Philosophy of Science of Thomas S. Kuhn*. Chicago: University of Chicago Press.
- James, W. 1979/1907: *Pragmatism*. Cambridge, MA: Harvard University Press.
- Jones, M. R. 2005: 'Idealization and Abstraction: A Framework', in M. R. Jones & N. Cartwright (eds.), *Idealization XII: Correcting the Model, Poznań Studies in the Philosophy of the Sciences and the Humanities* 86: 173-217.
- Keller, E. F. 1985: *Reflections on Gender and Science*. New Haven: Yale University Press.
- Knorr-Cetina, K. 1981: *The Manufacture of Knowledge*. Oxford: Pergamon.
- Kuhn, T. 1970/1962: *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kuhn, T. S. 1983: 'Commensurability, Comparability, Communicability', *PSA* 1982: 669-688.
- Ladyman, J. & D. Ross 2007: *Everything Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Latour, B & S. Woolgar 1986: *Laboratory Life: The Construction of Scientific Facts* (2nd ed.). Princeton: Princeton University Press.
- Laudan, L. 1981: 'A Confutation of Convergent Realism', *Philosophy of Science* 48: 19-48.
- Leeds, S. 2007: 'Correspondence Truth and Scientific Realism', *Synthese* 159: 1-21.
- Leplin, J. 1997: *A Novel Defence of Scientific Realism*. Oxford: Oxford University Press.

- Lipton, P. 2004: *Inference to the Best Explanation*, 2nd edition. London: Routledge
- Longino, H. 1990: *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press.
- Misak, C. 2011: 'Pragmatist Theories of Truth', in M. Glanzberg (ed.), *The Oxford Handbook of Truth*. Oxford: Oxford University Press.
- McMullin, E. 1985: 'Galilean Idealization', *Studies in History and Philosophy of Science* 16: 247-273.
- Nelson, L. H. & J. Nelson (eds.) (1996): *Feminism, Science, and the Philosophy of Science*. Dordrecht: Kluwer.
- Oddie, G. 2011: 'Truth and Truthlikeness', in M. Glanzberg (ed.), *The Oxford Handbook of Truth*. Oxford: Oxford University Press.
- Peirce, C. S. 1998/1992: *The Essential Peirce*, N. Houser, C. Kloesel, & the Peirce Edition Project (eds.). Bloomington: Indiana University Press.
- Pickering, A. 1984: *Constructing Quarks: A Sociological History of Particle Physics*. Edinburgh: Edinburgh University Press.
- Psillos, S. 1999: *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- Psillos, S. forthcoming: 'Choosing the Realist Framework', *Synthese*.
- Putnam, H. 1975: *Mathematics, Matter and Method*. Cambridge: Cambridge University Press.
- Putnam, H. 1981: *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- Quine, W. 1953: 'Two Dogmas of Empiricism', in *From a Logical Point of View*, pp. 20-46. Cambridge, MA: Harvard University Press.
- Richardson, A. W. 1998: *Carnap's Construction of the World*. Cambridge: Cambridge University Press.
- Sankey, H. 1994: *The Incommensurability Thesis*. London: Ashgate.
- Shapin, S. & S. Schaffer: 1985. *Leviathan and the Air Pump*. Princeton: Princeton University Press.
- Stanford, P. K. 2006: *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.
- Vaihinger, H. 1923/1911: *The Philosophy of 'As If'*, C.K. Ogden (tr.). London: Kegan Paul.
- Walker, R. 2011: 'The Coherence Theory of Truth', in M. Glanzberg (ed.), *The Oxford Handbook of Truth*. Oxford: Oxford University Press.
- Worrall, J. 1989: 'Structural Realism: The Best of Both Worlds?', *Dialectica* 43: 99-124. Reprinted in D. Papineau (ed.) 1996: *The Philosophy of Science*, pp. 139-165. Oxford: Oxford University Press.